Quantum Complexity Theory
Summer 2019, University of Paderborn

Sevag Gharibian

Lecture 7: Quantum-Classical Merlin Arthur (QCMA) and Ground State Connectivity

"I have called this principle, by which each slight variation, if useful, is preserved, by the term of Natural Selection."
— Charles Darwin

Contents

1 Quantum-Classical Merlin Arthur (QCMA) 1

2 Ground State Connectivity 2

2.1 QCMA-completeness of GSCON 4

2.1.1 Proof of Traversal Lemma 7

2.1.2 Tightness of the Traversal Lemma 8

Introduction. In Lecture 5, we introduced Quantum Merlin Arthur (QMA) as the *de facto* quantum generalization of NP, which verified a quantum proof $|\psi\rangle$ with a quantum verifier. It is not clear at all, however, whether a *quantum* proof is required to capture the full power of QMA. For even though an arbitrary quantum proof $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}$ can be a "complicated" quantum state, a QMA verifier is restricted to be a *polynomial-size* quantum circuit. Can such a limited verifier even "distinguish" between "complicated" proofs $|\psi\rangle$ and "simpler" approximations $|\widetilde{\psi}\rangle$ thereof (where by "simpler" we roughly mean that unlike $|\psi\rangle$, $|\widetilde{\psi}\rangle$ has a succinct classical description)? In other words, is a classical proof as good as a quantum one for the purposes of polynomial-time quantum verification?

In this lecture, we explore this question via the complexity class QCMA, which is QMA but with a classical proof. Whereas QCMA $\subseteq$ QMA holds trivially, it is not at all clear whether the reverse containment should hold. In other words, in line with the opening quote of this lecture, we do not yet know whether the "variation" of allowing proofs to be quantum yields something "useful" (meaning more verification power). In this sense, "Natural Selection" is yet to play its hand in the study of "quantum NP".

This lecture is organized as follows. We begin in Section 1 by defining QCMA. Section 2 studies the QCMA-complete problem of *Ground State Connectivity*, which gives some insight into the types of classical proofs which may be useful to quantum verifiers. The proof of QCMA-completeness (Section 2.1) will again utilize the history state construction of the previous lecture, and its soundness analysis requires a tool known as the Traversal Lemma (Section 2.1.1). The tightness of this lemma is discussed in the closing section of this lecture, Section 2.1.2.

1 Quantum-Classical Merlin Arthur (QCMA)

Sandwiched between PromiseMA and QMA is QCMA (more accurately, PromiseQCMA) (i.e. PromiseMA $\subseteq$ QCMA $\subseteq$ QMA), a natural complexity class which is arguably less well-understood than QMA. We begin with the definition of QCMA.

Definition 1 (Quantum-Classical Merlin Arthur (QCMA))**.** *A promise problem $\mathbb{A} = (A_{\text{yes}}, A_{\text{no}}, A_{\text{inv}})$ is in QCMA if there exists a P-uniform quantum circuit family $\{Q_n\}$ and polynomials $p, q : \mathbb{N} \mapsto \mathbb{N}$ satisfying the following properties. For any input $x \in \{0,1\}^n$, Q_n takes in $n + p(n) + q(n)$ qubits as input, consisting of* *the input x on register A, $p(n)$ qubits initialized to a "classical proof" $|y\rangle \in \{0,1\}^{p(n)}$ on register B, and $q(n)$ ancilla qubits initialized to $|0\rangle$ on register C. The first qubit of register C, denoted C_1, is the designated output qubit, a measurement of which in the standard basis after applying Q_n yields the following:*

- *(Completeness/YES case) If $x \in A_{\text{yes}}$, there exists proof $y \in \{0,1\}^{p(n)}$, such that Q_n accepts with probability at least $2/3$.*
- *(Soundness/NO case) If $x \in A_{\text{no}}$, then for all proofs $y \in \{0,1\}^{p(n)}$, Q_n accepts with probability at most $1/3$.*
- *(Invalid case) If $x \in A_{\text{inv}}$, Q_n may accept or reject arbitrarily.*

Exercise. What is the only difference between QMA and QCMA?

As with BQP and QMA, the completeness and soundness errors can be made exponentially small via parallel repetition of the verification protocol and a majority vote. However, as with PromiseMA, it turns out that without loss of generality, one may assume QCMA has *perfect* completeness, i.e. in the YES case, there exists a proof y accepted with certainty. An analogous statement for QMA remains an open question.

Exercise. For QMA, we distinguished between weak and strong error reduction. Why does strong error reduction for QCMA hold trivially?

Exercise. Note that the proof register B in Definition 1 is expected to contain a standard basis state $|y\rangle$ for $y \in \{0,1\}^{p(n)}$. Since Q_n is a *quantum* circuit, B is a register of $p(n)$ *qubits*. This means that in the NO case, a cheating prover can in principle send an arbitrary *entangled* state $|\psi\rangle$ in register B. Why does this not ruin the soundness property of Definition 1? (Hint: Without loss of generality, we may assume that before running the actual verification, Q_n makes a certain measurement. Which measurement should Q_n make, and how can Q_n simulate this measurement unitarily?)

2 Ground State Connectivity

We now study a physically motivated complete problem for QCMA, which gives some insight into what type of *classical* proof might be useful to a *quantum* verifier. As with the Quantum Cook-Levin theorem, the problem arises in the setting of ground space properties of local Hamiltonians, $H = \sum_i H_i$. In contrast to k-LH, however, we shift our attention away from the *ground state energy* of H and to the structural properties of the *ground space* of H. For inspiration, we look to the classical study of *reconfiguration problems.*

Reconfiguration problems. Given a 3-SAT formula $\phi : \{0,1\}^n \mapsto \{0,1\}$, computing different properties of ϕ can have different complexities. For example, we know from the Cook-Levin theorem that deciding whether the solution space for ϕ is *non-empty* (i.e. does ϕ have a satisfying assignment?) is NP-complete. If we instead wish to count the *size* of the solution space (i.e. the number of satisfying assignments to ϕ), this is much harder; it is #P-complete. We may also ask about the *structure* of the solution space — for example, is it *connected*? This turns out to be either in P, NP-complete or PSPACE-complete, depending on how the question is phrased.

Let us formalize what we mean by "connected" (we state it using ket notation to highlight the generalization to the quantum setting later). Given as input a Boolean formula $\phi : \{0,1\}^n \mapsto \{0,1\}$, two satisfying assignments $x, y \in \{0,1\}^n$, and length parameter 1^m for $m \in \mathbb{N}$, we say $|x\rangle$ and $|y\rangle$ are *connected* with respect to ϕ if there exists a sequence of length at most $N \leq m$ bit flips $(X_{i_1}, X_{i_2}, \ldots, X_{i_N})$ for $i_k \in [m]$ (where Pauli X_i is applied to qubit i) satisfying two properties:

1. (Intermediate states are in solution space) For all $k \in [N]$ and intermediate states $|x_k\rangle := X_{i_k} \cdots X_{i_1}|x\rangle$, $\phi(x_k) = 1$.

2. (Final state is target state) $X_{i_N} \cdots X_{i_1}|x\rangle = |y\rangle$.

In other words, is there a sequence of at most m bit flips we can apply to map x to y, such that each intermediate state attained is also a solution to ϕ?

Exercise. Let $\phi = (x_1 \vee \overline{x_2})$ and $x = 00$ and $y = 11$. Are x and y connected with respect to ϕ?

Exercise. Give a Boolean formula ϕ (not necessarily in CNF form) and solutions x, y which are *not* connected with respect to ϕ.

It is important to note that the number of bit flips m needed for mapping x to y in this manner need *not* be at most n; in fact, it can scale as $O(2^n)$. This is due to property 1 above; the naive greedy sequence of bit flips mapping x to y might take us temporarily *out* of the solution space. For this reason, if we drop the upper bound m in the input, the problem of determining if a 3-SAT formula is connected is PSPACE-complete. In the formulation above (i.e. with parameter m), however, the problem is NP-complete.

Exercise. Why is deciding if a 3-SAT formula is connected according to our definition above in NP?

Reconfiguration in the quantum setting. By generalizing the reconfiguration problem for 3-SAT to the quantum setting, we arrive at the main problem to be studied in this section. As before, we are interested in the structure of the solution space, where "solution space" now refers to the ground space of a local Hamiltonian.

Definition 2 (Ground State Connectivity (GSCON))**.** *Fix a polynomial* $\Delta : \mathbb{N} \mapsto \mathbb{R}+$.

- *Input parameters:*
  1. *k-local Hamiltonian $H = \sum_i H_i$ acting on n qubits with $H_i \in \mathrm{Herm}(\mathbb{C}^2)^{\otimes k}$ satisfying $\| H_i \|_\infty \leq 1$.*
  2. *Thresholds $\eta_1, \eta_2, \eta_3, \eta_4 \in \mathbb{R}$ such that $\eta_2 - \eta_1 \geq \Delta$ and $\eta_4 - \eta_3 \geq \Delta$, and 1^m for $m \in \mathbb{N}$.*
  3. *Polynomial size quantum circuits U_ψ and U_ϕ generating "starting" and "target" states $|\psi\rangle$ and $|\phi\rangle$ (starting from $|0\rangle^{\otimes n}$), respectively, satisfying $\langle\psi|H|\psi\rangle \leq \eta_1$ and $\langle\phi|H|\phi\rangle \leq \eta_1$.*
- *Output:*
  1. *If there exists a sequence of 2-qubit unitaries $(U_i)_{i=1}^m \in \mathrm{U}\left(\mathbb{C}^2\right)^{\times m}$ such that:*
     (a) (Intermediate states remain in low energy space) For all $i \in [m]$ and intermediate states $|\psi_i\rangle := U_i \cdots U_2 U_1 |\psi\rangle$, one has $\langle\psi_i|H|\psi_i\rangle \leq \eta_1$, and
     (b) (Final state close to target state) $\| U_m \cdots U_1|\psi\rangle - |\phi\rangle \|_2 \leq \eta_3$,
     then output YES.
  2. *If for all 2-qubit sequences of unitaries $(U_i)_{i=1}^m \in \mathrm{U}\left(\mathbb{C}^2\right)^{\times m}$, either:*
     (a) (Intermediate state obtains high energy) There exists $i \in [m]$ and an intermediate state $|\psi_i\rangle := U_i \cdots U_2 U_1|\psi\rangle$, such that $\langle\psi_i|H|\psi_i\rangle \geq \eta_2$, or
     (b) (Final state far from target state) $\| U_m \cdots U_1|\psi\rangle - |\phi\rangle \|_2 \geq \eta_4$,
     then output NO.

Intuition. Roughly, GSCON says: Given two ground states[1] $|\psi\rangle$ and $|\phi\rangle$ of a k-local Hamiltonian H, are $|\psi\rangle$ and $|\phi\rangle$ *connected* through the ground space of H? In other words, is there a sequence of 2-qubit gates mapping $|\psi\rangle$ to $|\phi\rangle$, such that all intermediate states encountered are also ground states? This turns out to

[1] More accurately, the definition of GSCON discusses low-energy states, which need not be ground states. However, by using the fact that QCMA satisfies perfect completeness without loss of generality, one can show QCMA-completeness of GSCON even when all states involved are ground states of H.

have an important physical motivation in the quantum setting — it asks whether the ground space of H has an *energy barrier* preventing one ground state from being mapped to another via short circuits (while remaining in the low energy space throughout the computation).

2.1 QCMA-completeness of GSCON

Theorem 3. *There exists a polynomial r such that* GSCON *is* QCMA*-complete for $m \in O(r(n))$ and $k \geq 7$, where n denotes the number of qubits H acts on.*

Proof. For brevity, we sketch containment in QCMA. We give a full proof of QCMA-hardness.

Containment in QCMA. Containment in QCMA holds for any $m \in O(\text{poly}(n))$, and is intuitively straightforward (hence we only sketch it here) — the prover sends[2] a classical description of the polynomially many 2-qubit gates $U_1, \ldots, U_m$ to the verifier as a proof. Since the verifier can prepare both the start and final states $|\psi\rangle$ and $|\phi\rangle$ efficiently via U_ψ and U_ϕ (also given as input), it can check the energy of each intermediate state $|\psi_i\rangle$ against H using the protocol from the proof of the Cook-Levin theorem. Finally, to check if $|\psi_m\rangle \approx |\phi\rangle$, the verifier prepares both states and applies the SWAP test[3].

QCMA-hardness. Let $A = (A_{\text{yes}}, A_{\text{no}}, A_{\text{inv}})$ denote a QCMA promise problem. Let $x \in \{0,1\}^n$ be an input, with corresponding QCMA verifier $V = V_N \cdots V_1 \in (\mathbb{C}^2)^{\otimes n} \otimes (\mathbb{C}^2)^{\otimes p(n)} \otimes (\mathbb{C}^2)^{\otimes q(n)}$. Recall V is a uniformly generated quantum verification circuit consisting of 1- and 2-qubit unitary gates, acting on registers A (n qubits containing the input x), B ($p(n)$ qubits containing the proof $|\psi\rangle$), and C ($q(n)$ ancilla qubits initialized to all zeroes). We make two assumptions about V without loss of generality: (1) The completeness and soundness parameters for V are $1-\epsilon$ and ϵ, so that $1 - 2\epsilon \in \Omega(1/\operatorname{poly}(n))$, and (2) V begins by measuring its proof in the standard basis (this trick was alluded to in a previous exercise; it forces a cheating prover to send a classical proof, and can be simulated unitarily via the principle of deferred measurement). Our goal is to construct an instance $(H, \eta_1, \eta_2, \eta_3\, \eta_4, m, U_\psi, U_\phi)$ with 7-local Hamiltonian H such that, if $x \in A_{\text{yes}}$, then there exists a sequence $U_1, \ldots, U_m$ of 2-qubit gates mapping $|\psi\rangle$ to $|\phi\rangle$ through the low energy space of H, and if $x \in A_{\text{no}}$, then any such unitary sequence must either leave the low-energy space of H at some point or map $|\psi\rangle$ "far" from $|\phi\rangle$.

The construction. We now state the construction, which at first glance does not seem to do anything interesting. The YES case in the correctness analysis will reveal the intuition as to why this works.

Let H^{CL} denote the 5-local Hamiltonian obtained from V the Quantum Cook-Levin Theorem's circuit-to-Hamiltonian construction. We define H to act on a *Hamiltonian* register denoted h (consisting of subregisters A, B, C, and D, where D is the unary-encoded clock register) and 3-qubit *GO* register denoted G. Specifically, $H \in \text{Herm}((\mathbb{C}^2)_h^{\otimes(n+p(n)+q(n)+N)} \otimes (\mathbb{C}^2)_G^{\otimes 3})$.

Exercise. What are the N qubits in the h register used for? Why are there N of them?

Define

$$H := H_h^{\text{CL}} \otimes P_G \qquad \text{for} \qquad P := I - |000\rangle\langle 000| - |111\rangle\langle 111|. \tag{1}$$

[2] More accurately, since arbitrary 2-qubit gates cannot be specified exactly to finite precision, the prover sends elements from an appropriate notion of an "ϵ-net" over 2-qubit unitaries.

[3] For brevity, we omit an in-depth discussion of the SWAP test, but it is a tool worth remembering: Given physical copies of states $|\psi\rangle$ and $|\phi\rangle$, the SWAP test outputs 0 with probability $(1 + |\langle\psi|\phi\rangle|^2)/2$, thus allowing us to estimate $\|\, \psi - \phi\,\|_2$. The circuit description for the SWAP test is simple — it adds an ancilla in the $|+\rangle$ state, and conditioned on the ancilla being $|1\rangle$, swaps the registers containing $|\psi\rangle$ and $|\phi\rangle$. We then perform a Hadamard on the ancilla and measure in the standard basis.

Exercise. Note that P as written is 3-local. Show how to write P equivalently as a 2-local Hamiltonian.

By the exercise above, we have that H is 7-local. Define the initial and final states as

$$|\psi\rangle := |0\rangle^{\otimes(n+p(n)+q(n)+N)}|0\rangle^{\otimes 3} \qquad \text{and} \qquad |\phi\rangle := |0\rangle^{\otimes(n+p(n)+q(n)+N)}|1\rangle^{\otimes 3}, \tag{2}$$

which have trivial poly-size preparation circuits U_ψ and U_ϕ. Finally, let W denote a unitary circuit of size $|W|$ which prepares the history state of H given classical proof y. Define $m := 2(p(n) + |W| + 1)$.
To complete the construction, set $\eta_3 = 0$, $\eta_4 = 1/4$, $\eta_1 = \alpha$, and $\eta_2 = \beta/(16N^2)$, where $\alpha := \epsilon/(m+1)$ and $\beta := \pi^2(1-\sqrt{\epsilon})/(2(N+1)^3)$ come from the Quantum Cook-Levin theorem's circuit-to-Hamiltonian construction. Note that if we apply weak error reduction to V so that ϵ is exponentially close to zero, then the gap $\Delta \in \Omega(1/N^5)$.

Correctness for YES case. Suppose $x \in A_{\text{yes}}$, i.e. there exists a proof $y \in \{0,1\}^{p(n)}$ accepted by V. We demonstrate a sequence $(U_i)_{i=1}^m$ of 2-qubit unitaries mapping $|\psi\rangle$ to $|\phi\rangle$ through the ground space of H.

Intuition. To see the intuition, we first need an exercise.

Exercise. Prove that $|\psi\rangle$ and $|\phi\rangle$ lie in the null space of H, i.e. $H|\psi\rangle = H|\phi\rangle$. Use the fact that $H^{\text{CL}} \succeq 0$ to conclude that $|\psi\rangle$ and $|\phi\rangle$ hence lie in the ground space of H, and in particular have energy at most α.

Thus, $|\psi\rangle$ and $|\phi\rangle$ both lie in the null space of H, and are identical except for the pesky 3 qubits in the GO register, which are set to $|000\rangle$ and $|111\rangle$, respectively. To map $|\psi\rangle$ to $|\phi\rangle$ via 2-local gates, the obvious idea is hence to flip the GO qubits from 000 to 111. The problem is that we cannot flip more than two qubits at a time — so after flipping (say) the first two GO qubits, the G register reads $|110\rangle$, and now we are in the support of P_G in the definition of H. This means that H^{CL} is now "turned on" and checks the h register for a history state. Since we are in the YES case, there *is* a good history state we can use, and moreover, since we are dealing with QCMA, this history state can be prepared from the all zeroes initial state via a polynomial-size (though not necessarily P-uniformly generated) circuit.

Exercise. Given a classical proof $y \in \{0,1\}^{p(n)}$, give a polynomial-size sequence of 2-qubit gates which maps the all-zeroes state to the history state $|\psi_{\text{hist}}\rangle$ in time polynomial in n. Why does this not necessarily work for QMA, i.e. given a copy of a quantum proof $|\eta\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$ in place of y?

The honest prover's actions. Recall the h register is broken up into four subregisters, A (input), B (proof), C (ancilla), and D (clock). The sequence of 2-qubit gates $(U_i)_{i=1}^m$ is as follows:

1. Apply Pauli X gates to h_B to prepare classical proof y, i.e., map $|0\rangle^{\otimes p(n)}$ to $|y\rangle$.
2. Apply W to h to prepare the history state $|\text{hist}_y\rangle$ of H^{CL}.
3. Apply $(X \otimes X \otimes I)_G$ to "initiate" checking of $|\text{hist}_y\rangle$.
4. Apply $(I \otimes I \otimes X)_G$ to "complete" checking of $|\text{hist}_y\rangle$.
5. Apply $W^\dagger$ to h to uncompute $|\text{hist}_y\rangle$.
6. Apply X gates to h_B to map the initial proof $|y\rangle$ back to $|0\rangle^{\otimes p(n)}$.

Exercise. Why is the length of the sequence above at most $m = 2(p(n) + |W| + 1)$, as desired?

Exercise. Verify that the sequence (U_i) correctly maps $|\psi\rangle$ to $|\phi\rangle$.

Exercise. As in Definition 2, recall the ith intermediate state is defined $|\psi_i\rangle := U_i \cdots U_1|\psi\rangle$. There is precise only one $i \in [6]$ above such that $|\psi_i\rangle$ is *not* in the null space of H — which i is this? Prove that for this i, $\langle\psi_i|H|\psi_i\rangle \leq \eta_1$, as desired.

Proof of correctness for NO case. Suppose $x \in A_{\text{no}}$, i.e. all proofs $y \in \{0,1\}^{p(n)}$ are rejected by V with probability at least $1-\epsilon$. Intuitively, we cannot proceed as in the YES case now because H^{CL} does *not* have a low-energy history state (indeed, all its eigenvalues are at least β). Thus, the moment we "switch on" the H^{CL} check in Step 3, we are in trouble. The only way a cheating prover can try to bypass this problem is to somehow try to switch all GO qubits from 000 to 111 *without* having significant support on the orthogonal space spanned by $\{|001\rangle, |010\rangle, |011\rangle, |100\rangle, |101\rangle, |110\rangle\}$; note that this statement is non-trivial because the prover is not restricted to simply performing Pauli X gates. Thus, our main task for the NO case is to prove that this is impossible. The main tool for this is the following lemma, whose proof is given in Section 2.1.1.

Lemma 4 (Traversal Lemma). *Let $S, T \subseteq (\mathbb{C}^2)^{\otimes n}$ be k-orthogonal subspaces. Fix arbitrary states $|v\rangle \in S$ and $|w\rangle \in T$, and consider a sequence of k-qubit unitaries $(U_i)_{i=1}^m$ such that*

$$\| \, |w\rangle - U_m \cdots U_1|v\rangle \, \|_2 \leq \delta$$

for some $0 \leq \delta < 1/2$. Define $|v_i\rangle := U_i \cdots U_1|v\rangle$ and $P := I - \Pi_S - \Pi_T$. Then, there exists $i \in [m]$ such that

$$\langle v_i|P|v_i\rangle \geq \left(\frac{1-2\delta}{2m}\right)^2.$$

Intuitively, this lemma says basically what we are looking for — that for certain types of subspaces S and T, any local unitary mapping from S to T must at some point "leave" $S \oplus T$. The "type" of subspaces this applies to are defined next.

Definition 5 (k-orthogonal states and subspaces). *For $k \geq 1$, a pair of states $|v\rangle, |w\rangle \in (\mathbb{C}^d)^{\otimes n}$ is k-*orthogonal *if for all k-qubit unitaries U, we have $\langle w|U|v\rangle = 0$. We call subspaces $S, T \subseteq (\mathbb{C}^2)^{\otimes n}$ k-*orthogonal *if any pair of vectors $|v\rangle \in S$ and $|w\rangle \in T$ are k-orthogonal.*

Exercise. Prove that $|000\rangle$ and $|111\rangle$ are 2-orthogonal. Are they 3-orthogonal?

We can now complete the proof of correctness for the NO case. We know the smallest eigenvalue of H^{CL} is at least β. Let S and T denote the $+1$ eigenspaces of $I_h \otimes |000\rangle\langle 000|_G$ and $I_h \otimes |111\rangle\langle 111|_G$, respectively.

Exercise. Characterize S and T.

Exercise. Show that S and T are 2-orthogonal subspaces, and that $|\psi\rangle \in S$ and $|\phi\rangle \in T$.

By the exercises above, for any sequence of two-qubit unitaries $(U_i)_{i=1}^m$, either $\| \, |\psi_m\rangle - |\phi\rangle \, \|_2 \geq 1/4 = \eta_4$ (in which case we have a NO instance of GSCON and we are done), or we can apply the Traversal Lemma (Lemma 4) with $\delta = 1/4$ to conclude that there exists an $i \in [m]$ such that

$$\langle\psi_i|P'|\psi_i\rangle \geq \left(\frac{1}{4m}\right)^2 = \frac{\eta_2}{\beta},$$

where recall $|\psi_i\rangle := U_i \cdots U_1|\psi\rangle$ and we define $P' = I - \Pi_S - \Pi_T = I_h \otimes P$. We conclude that

$$\langle\psi_i|H|\psi_i\rangle = \langle\psi_i|H^{\text{CL}} \otimes P|\psi_i\rangle \geq \beta\langle\psi_i|I_h \otimes P|\psi_i\rangle = \beta\langle\psi_i|P'|\psi_i\rangle \geq \eta_2,$$

where the first inequality follows since $H^{\text{CL}} \succeq \beta I$.

□

2.1.1 Proof of Traversal Lemma

To conclude the proof of Theorem 3, it remains to show the Traversal Lemma, which we reproduce below for convenience.

Lemma 4. *Let $S, T \subseteq (\mathbb{C}^2)^{\otimes n}$ be k-orthogonal subspaces. Fix arbitrary states $|v\rangle \in S$ and $|w\rangle \in T$, and consider a sequence of k-qubit unitaries $(U_i)_{i=1}^m$ such that*

$$\| \, |w\rangle - U_m \cdots U_1 |v\rangle \, \|_2 \leq \delta$$

for some $0 \leq \delta < 1/2$. Define $|v_i\rangle := U_i \cdots U_1 |v\rangle$ and $P := I - \Pi_S - \Pi_T$. Then, there exists $i \in [m]$ such that

$$\langle v_i | P | v_i \rangle \geq \left(\frac{1-2\delta}{2m} \right)^2 .$$

The proof of Lemma 4 requires the well-known "Gentle Measurement Lemma", which quantifies an intuitively expected behavior: If a measurement outcome Π has high probability of occurring for state ρ, then we expect the postmeasurement state (proportional to) $\Pi\rho\Pi$ to be approximately ρ (i.e. the measurement should not disturb the state much). We state this lemma below first, and then prove Lemma 4.

Lemma 6 (Gentle Measurement Lemma)**.** *Let $\rho \in \mathrm{L}\left(\mathbb{C}^d\right)$ be a density operator and $O \preceq \Pi \preceq I$ a projective measurement operator for $\Pi \in \mathrm{L}\left(\mathbb{C}^d\right)$, such that $\mathrm{Tr}(\Pi\rho) \geq 1 - \epsilon$. Then, $\| \, \rho - \Pi\rho\Pi \, \|_{\mathrm{tr}} \leq 2\sqrt{\epsilon}$.*

Note that $\Pi\rho\Pi$ above is *not* necessarily normalized.

Proof of Lemma 4. We give a proof by contradiction. Suppose that for all $i \in [m]$, the expectations satisfy $\langle v_i | P | v_i \rangle < \kappa := [(1-2\delta)/(2m)]^2$. Consider the following thought experiment inspired by the quantum Zeno effect[4]. Imagine that after each U_i is applied, we measure $|v_i\rangle$ using the projective measurement $(\Pi, I - \Pi)$ for $\Pi := I - P$, and postselect on obtaining outcome Π. Define the following two sequences:

- $|v_i'\rangle := \Pi |v_i\rangle$ for $i \in [m]$,
- $|v_1''\rangle := |v_1'\rangle$ and $|v_i''\rangle := \Pi \, U_i |v_{i-1}''\rangle$ for $i \in \{2, \ldots, m\}$.

Note that $|v_i'\rangle$ and $|v_i''\rangle$ are not necessarily normalized.

To set up our contradiction, we first prove by induction on i that

$$\| \, |v_i\rangle\langle v_i| - |v_i''\rangle\langle v_i''| \, \|_{\mathrm{tr}} < 2i\sqrt{\kappa}. \tag{3}$$

For the base case $i = 1$, we have $|v_1''\rangle = |v_1'\rangle$. Then, since $\langle v_1 | P | v_1 \rangle < \kappa$, we know that $\mathrm{Tr}(\Pi |v_1\rangle\langle v_1|) > 1 - \kappa$.

Exercise. Use the Gentle Measurement Lemma (Lemma 6) to conclude

$$\| \, |v_1\rangle\langle v_1| - |v_1''\rangle\langle v_1''| \, \|_{\mathrm{tr}} < 2\sqrt{\kappa}, \tag{4}$$

as required for the base case.

For the inductive case, assume Equation (3) holds for $1 \leq i \leq j - 1$. We prove it holds for $i = j$. Specifically,

$$\begin{aligned}
\left\| \, |v_j\rangle\langle v_j| - |v_j''\rangle\langle v_j''| \, \right\|_{\mathrm{tr}} &\leq \left\| \, |v_j\rangle\langle v_j| - |v_j'\rangle\langle v_j'| \, \right\|_{\mathrm{tr}} + \left\| \, |v_j'\rangle\langle v_j'| - |v_j''\rangle\langle v_j''| \, \right\|_{\mathrm{tr}} \\
&< 2\sqrt{\kappa} + \left\| \, |v_j'\rangle\langle v_j'| - |v_j''\rangle\langle v_j''| \, \right\|_{\mathrm{tr}} \\
&= 2\sqrt{\kappa} + \left\| \Pi U_j \left(|v_{j-1}\rangle\langle v_{j-1}| - |v_{j-1}''\rangle\langle v_{j-1}''| \right) U_j^\dagger \Pi \right\|_{\mathrm{tr}} \\
&\leq 2\sqrt{\kappa} + \left\| \, |v_{j-1}\rangle\langle v_{j-1}| - |v_{j-1}''\rangle\langle v_{j-1}''| \, \right\|_{\mathrm{tr}} \\
&< 2\sqrt{\kappa} + 2(j-1)\sqrt{\kappa} \\
&= 2j\sqrt{\kappa},
\end{aligned} \tag{5}$$

[4] Roughly, the quantum Zeno effect is the phenomenon that a quantum system which is continuously observed never evolves.

where the first statement follows from the triangle inequality, the second from the Gentle Measurement Lemma, the fourth from the facts that the Schatten p-norms are invariant under isometries and that $\| ABC \|_p \leq \| A \|_\infty \| B \|_p \| C \|_\infty$, and the fifth from the induction hypothesis. This establishes Equality (3).

Exercise. Use the fact that $\| |v\rangle\langle v| - |w\rangle\langle w| \|_{\text{tr}} \leq 2 \| |v\rangle - |w\rangle \|_2$ for unit vectors $|v\rangle, |w\rangle$ to conclude that

$$\| |v''_m\rangle\langle v''_m| - |w\rangle\langle w| \|_{\text{tr}} < 1, \tag{6}$$

We are now ready to obtain the desired contradiction. To do so, observe that since $|v\rangle \in S$, and since S and T are k-orthogonal subspaces, we have that for all $i \in [m]$, $|v''_i\rangle \in S$ (i.e., if S is 1-dimensional, this is the Zeno effect). Thus, we have $\langle v''_m|w\rangle = 0$, implying that

$$\| |v''_m\rangle\langle v''_m| - |w\rangle\langle w| \|_{\text{tr}} = 1 + \| |v''_m\rangle \|_2 \geq 1.$$

This contradicts Equation (6), as desired. □

2.1.2 Tightness of the Traversal Lemma

The proof of QCMA-completeness of GSCON relied crucially on the Traversal Lemma, and so it is natural to ask whether the lemma is tight. Namely, in Lemma 4, the lower bound on $\langle v_i|P|v_i\rangle$ scales as $\Theta(1/m^2)$ (for m the number of unitaries and for fixed δ). This suggests that one can better "avoid" the subspace P projects onto if one uses a longer sequence of local unitaries. Indeed, it turns out that, at least in some cases, this behavior is possible; thus, a dependence on m is necessary in Lemma 4.

Theorem 7. *We assume the notation of Lemma 4. Fix any $0 < \Delta < 1/2$, and consider 2-orthogonal states $|v\rangle = |000\rangle$ and $|w\rangle = |111\rangle$, with $P := I - |v\rangle\langle v| - |w\rangle\langle w|$. Then, there exists a sequence of m 2-local unitary operations mapping $|v\rangle$ to $|w\rangle$ through intermediate states $|v_i\rangle$, each of which satisfy $\langle v_i|P|v_i\rangle \leq \Delta$, and where $m \in O(1/\Delta^2)$.*

Proof intuition. We omit the proof for brevity, but the idea behind it is based on the following rough analogy: Suppose one wishes to map the point $(1, 1)$ (corresponding to $|000\rangle$) in the 2D Euclidean plane to $(-1, -1)$ (corresponding to $|111\rangle$) via a sequence of moves with the following two restrictions: (1) For each current point (x, y), the next move must leave precisely one of x or y invariant (analogous to 2-local unitaries acting on a 3-qubit state), and (2) the Euclidean distance between (x, y) and the line through $(1, 1)$ and $(-1, -1)$ never exceeds Δ (analogous to the overlap with P not exceeding Δ). In other words, we wish to stay close to a diagonal line while making only horizontal and vertical moves. This can be achieved by making a sequence of "small" moves resembling a "staircase". The smaller the size of each "step" in the staircase, the better we approximate the line, at the expense of requiring more moves (analogous to increasing the number of unitaries, m). This is the basic premise behind the proof of Theorem 7 — giving a formal proof takes some work, as the back-and-forth shuffling of amplitude with the application of each local gate U_i needs to be carefully managed.