Lecture Notes - Proof of the NLTS Theorem

Rio Weil

This document was typeset on November 23, 2022

Abstract:
This set of lecture notes is for a presentation/literature review given to the UBC Quantum Information Group in November of 2022. I go through the motivation and a sketch of the proof of the No Low-energy Trivial States (NLTS) theorem, as proven by Anurag Anshu, Nikolas P. Breuckmann, and Chinmay Nirkhe. I will discuss the background and give a sketch for the proof presented in [1], following closely the discussion in Nirkhe's PhD thesis [2] (with many statements taken verbatim). In developing these notes I have taken inspiration from a presentation about the result given by Nirkhe in [3].

Contents

1 Motivation - Breaking Intuition in Mathematics and Physics

When we think about proofs in mathematics, we think of them as rigorous, step-by-step procedures which must be taken in their entirety to ensure that they are correct. However, the celebrated *Probabalistically Checkable Proofs*, or PCP theorem tells us that we can be confident about the correctness of the proof under much less strict conditions - namely, a randomized polynomial-time verifier who only reads a constant number of (randomized) characters of the proof can verify it, up to a constant probability of error [4].

Much like the classical PCP theorem upends our expectations for mathematical proof, the unproven quantum PCP conjecture has implications which upend our expectations for quantum systems. One such implication is the *No low-energy trivial lying states*, or NLTS Theorem - which has been recently proven by Anshu, Breuckmann, and Nirkhe in [1]. Interpreted physically, this theorem tells us that there exists quantum systems for which entanglement is robust (and exotic); our intution as physicists is that thermal interactions should destroy quantum effects at higher temperatures, but in fact there exist NLTS Hamiltonians for which long-range entanglement persists in this regime.

One might expect that such Hamiltonians have to be very engineered, and indeed this is the case - they have been found by considering Hamiltonians that correspond to good quantum error-correcting codes. Let us now begin the journey of constructing such Hamiltonians.

2 Background - Classical and Quantum Complexity Theory

2.1 NP and CSPs

Definition: NP and NP-Hard/Complete

The complexity class NP consists of all decision problems for which a proof can be verified efficiently. In this case, we take a proof to be a string of n bits, and efficient verification to mean that a verifier (such as a Turning machine) is able to check the correctness in $\text{poly}(n)$ time.
A problem is NP-hard if it can simulate any other NP-problem, i.e. any NP-problem can be transformed into it in polynomial time. A problem is NP-complete if it is NP-hard and NP.

Definition: Constraint Satisfaction Problem (CSP)

A k-CSP C is a formula on n variables over $\{0,1\}$ (though the definition generalizes to an alphabet Σ) composed of m k-variable clauses C_i on subsets of the n variables. The value of the k-CSP C for any assignment $x \in \{0,1\}^n$ is:

$$C(x) := \sum_{i=1}^{m} C_i(x). \tag{2.1}$$

A CSP is satisfiable if $\exists x \in \{0,1\}^n$ such that $C(x) = 0$.

Note that determining a CSP is satisfiable means that one must decide one of two things:

1. $\exists x \in \{0,1\}^n$ such that $C(x) = 0$ (satisfied).
2. $\forall x \in \{0,1\}^n$, $C(x) \geq 1$ (not satisfied).

Theorem: Cook-Levin

A CSP instance being satisfiable is NP-complete.

2.2 The PCP Theorem

The Probalistically checkable proofs, or PCP theorem can be viewed as a strengthening of Cook-Levin.

Theorem: PCP (Gap Amplification)

Given a CSP C, is NP-hard to decide if:

1. $\exists x, C(x) = 0$ (satisfiable)
2. $\forall x, C(x) \geq \frac{m}{2}$ (at least half the local checks in C are violated for all potential solutions x). The $\frac{1}{2}$ present here is fairly arbitrary - one can replace this with some other constant fraction of clauses being guaranteed to be violated.

In other words, even when we are promised that any x will either satisfy all clauses or violate some constant fraction, the problem remains equally as hard as before! The above statement of the PCP theorem is known as the Gap-Amplification version as the statement - it sounds quite different from the one given in the introduction, which we state below:

Theorem: PCP (Proof Checking)

Any NP problem (any proof) can be converted into a form (in polynomial time) such that only a randomized, $O(1)$ (constant) number of bits need to be read to be 99% confident in its validity.

The two forms of the theorem seem quite distinct, but one way to see equivalence (at least in one direction) is that given a proposed solution x it only takes a constant number of checks to be 99% confident that you are either dealing with a case where the CSP is satisfied or whether half of the local checks are violated. As discussed in the introduction, this second formulation of the PCP theorem has profound implications for our understanding of proofs - we can check them quickly with randomized, global checks to be confident about their correctness. It also implies that noisy proofs can suffice, as (for example) any x with $C(x) \leq \frac{m}{4}$ (i.e. one that satisfies most of the constraints) can be probabilistically verified with a constant number of checks (it only takes a constant number of checks to be confident that $C(x) \geq \frac{m}{2}$ or $C(x) < \frac{m}{4}$) - so our notion of proof can be modified from something that satisfies all constrains to most of the constraints.

2.3 QMA and Local Hamiltonians

We now step out of the classical and into the quantum realm. To this end, we replace our classical objects with quantum ones; our classical proofs (bit strings) become quantum states on n qubits, and our verifiers no longer become classical Boolean machines but quantum verifiers running on quantum computers. Let us consider the generalization of NP in this setting.

Definition: QMA

Quantum Merlin-Arthur, or QMA are decision problems for which when the answer is YES there exists a polynomial-size quantum proof/state which can efficiently (polynomial-time) convince a quantum verifier (quantum circuit) with high probability. If the answer is NO, every quantum state is rejected with high probability[a].

[a] Usually these probabilities are taken as 2/3 and 1/3 respectively - but the actual values are arbitrary, and they only be separated by a universal constant.

Much like the CSP-satisfication problem gave us a canonical NP-complete problem, we have an analog for QMA:

Definition: Local Hamiltonian Problem

A k-local hamiltonian H is an operator acting on n-qubits composed of m k-local Hamiltonian terms h_i. Each term h_i is a Hermitian operator that acts on k-qubits, and we form H by taking their sum:

$$H = \sum_{i=1}^{m} h_i \tag{2.2}$$

where we have an identity operation on all sites which h_i does not act on. The ground state energy of the Hamiltonian is the minimum eigenvalue of H, denoted as $\lambda_{\min}(H)$, which can be expressed as $\min_{|\psi\rangle} \langle\psi|H|\psi\rangle$.

We are also given two numbers a, b such that $\Gamma = b - a > 1/\text{poly}(n)$. Γ is called the absolute promise gap, and $\gamma = \Gamma/m$ the relative promise gap.

The local Hamiltonian problem is then the following - is $\lambda_{\min}(H) \leq a$, or are all eigenvalues of H larger than b? Equivalently, it is deciding:

1. $\exists|\psi\rangle, \langle\psi|H|\psi\rangle \leq a$
2. $\forall|\psi\rangle, \langle\psi|H|\psi\rangle \geq b$.

Theorem: Kitaev

The local Hamiltonian problem is `QMA`-complete.

The above Local Hamiltonian problem can be reformulated in a less technical way to say that it is hard to estimate the ground state energy of an n-qubit local Hamiltonian with precision $1/\text{poly}(n)$. Note the close analogy with the CSP-satisfaction problem here; in both cases, we take a local phenomenon (local constraints/local Hamiltonian terms) and then sum them into a global problem, whose solution is something computationally universal.

Before moving to the quantum PCP conjecture, we note that it is widely believed that `NP` $\neq$ `QMA`, and so a classical witness for a `QMA` proof would be superpolynomial in the number of bits. In the Hamiltonian language, this means that there is in general no efficient (i.e. polynomial) classical description of the ground state Hamiltonian.

2.4 The QPCP Conjecture

While we have a PCP theorem, we only have a quantum conjecture - the problem has been open for the last two decades, and is arguably one of the main open problems in quantum complexity today.

Conjecture: QPCP (Gap Amplification) - Formal

The LH problem with a constant relative promise gap γ is `QMA`-hard under quantum polynomial time reductions.

For the sake of explanation, let us consider a slightly less formal statement:

Conjecture: QPCP (Gap Amplification) - Simplified

It is `QMA`-complete to decide whether a local Hamiltonian H on n qubits and $m = \Theta(n)$ terms has $\lambda_{\min}(H) \leq \epsilon m/2$ (yes) or $\lambda_{\min}(H) \geq \epsilon m$ (where ϵ is some small constant, say $1/5$), even when promised that one of the cases holds.

There is also an equivalent statement that involves the checking of proofs:

Conjecture: QPCP (Proof Checking)

For any decision problem in QMA there exists a polynomial time quantum verifier such that the verifier accesses only $O(1)$ (constant) qubits from the a given proof and decides on acceptance/rejection with constant error probability.

Much like in the classical case where noisy proofs which violated some small constant fraction of checks were viable, if the QPCP conjecture holds, then the same is true of quantum states. Namely, proofs/quantum states that are not the ground state (that is, low-energy states), but are some small energy above the ground state are valid (in the simplified statement above, proofs/states that have energy $< \epsilon m/2$ would be convincing to a verifier as correct (that is to say, these are sufficient to prove that a Hamiltonian has a low energy ground state, even if not the ground state itself), even if these are not the ground state!)

This has the implication that (low-energy) thermal states of the local Hamiltonian must be exotic - the argument is as follows. If the thermal states have sufficiently fast decrease in entanglement such that they are well-approximated by a product state, but nevertheless is able to convince a verifier of the LH problem, this implies that there is a classical proof and classical computation that could decide the LH problem, and therefore $\mathtt{QMA} = \mathtt{NP}$! Hence, if we assume $\mathtt{QMA} \neq \mathtt{NP}$, the QPCP conjecture implies all low-energy states (and not just the ground state)! are far from trivial - This comes back to the physically unintuitive result as normally we expect entanglement to decrease and thermal states to be much simpler than ground states.

2.5 To the NLTS Theorem

Generalizing slightly from the above example of thermal states, any low-energy state which is the output of a constant depth circuit is a trivial/low-depth state with a classical description, and there exists an efficient classical algorithm for computing the energy of a trivial state with respect to a local Hamiltonian.

Let us prove this. Consider a quantum circuit U of depth t which is composed of a tensor product of disjoint 2-qubit unitaries U_j at each level. So, $U = \prod_{j=t}^{1} U_j = \prod_{j=t}^{1} \otimes_i U_{ji}$. We then have the following Lemma concerning causality and lightcones:

Lemma: Lightcones

We define the lightcone of an operator A as the set of qubits on which $UAU^\dagger$ acts nontrivially, and the lightcome of a qubit i as the union of lightcones over all operators A supported on qubit i.
For a circuit U of depth t, the size of the lightcome of a given qubit i is $\leq 2^t$.

To see this, we consider that $U_t \dots U_1 A U_1^\dagger \dots U_t^\dagger$ acts non-trivially on at most 2^t qubits, since each U_i is a tensor product of 2-qubit unitaries.

A low-depth circuit U as described above has an efficient classical description. We consider a state $|\psi\rangle = U|0\rangle^{\otimes n}$ prepared by a circuit; such states will be trivial, as we can efficiently compute the energy of a local Hamiltonian w.r.t. such a state:

$$\langle\psi|H|\psi\rangle = \langle\psi|\sum_{i=1}^{m} h_i|\psi\rangle = \sum_{i=1}^{m}\langle\psi|h_i|\psi\rangle = \sum_{i=1}^{m}\langle 0|^{\otimes n} U^\dagger h_i U|0\rangle^{\otimes n} \tag{2.3}$$

where we have used the linearity of the Hamiltonian and expectation. Now, due to the above Lightcone Lemma, $U^\dagger h_i U$ acts non-trivially on a constant number $O(2^t)$ of qubits - so we are able to brute force the $O(2^{2^t}) = O(1)$ computation on a classical device. The point is that for a low-depth state, we only need to classically describe the circuit, which is efficient. So, such states are trivial in this context.

So, if the QPCP conjecture is correct and $\mathtt{QMA} \neq \mathtt{NP}$, all low-energy states cannot be classically described, and therefore cannot be low-depth states by our arugment above; hence we have the implication of the No Low-Energy Trivial States, or NLTS theorem:

Theorem: NLTS

There exist a family of Hamiltonians such that every low-energy state cannot be generated by a constant depth circuit.

Which we will now discuss the ingredients for proving! A small comment before diving in - I'd like to make an analogy with measurement-based quantum computation here. We've taken a difficult statement about computation and complexity (the QPCP theorem) and recast it into a statement that is just about entanglement and Hamiltonians/states - much like how in MBQC we shift from the study of computation to the study of computationally useful states.

3 Step 1 - Quantum Codes and Local Indistinguishability

We start by taking inspiration from a property of error correcting codes - namely, local indistinguishability.

Consider a quantum error correcting code of distance-d; such an error correcting code is able to correct an erasure error on a subset S of qubits as long as $|S| \leq d$, as we can consider a completely depolarizing channel:

$$\mathcal{E}_S(\rho) = \rho_{-S} \otimes \nu_S = \rho_{-S} \otimes \frac{1}{4^{|S|}} \sum_{a,b \in \{0,1\}^S} (X^a Z^b) \rho_S (Z^b X^a) \tag{3.1}$$

which replaces the qubits in S with a maximally mixed state, equivalent to an erasure error.

Now, suppose we are able to correct this erasure error via quantum error correction, i.e. obtain all of the lost information of ρ_S from the remaining portion ρ_{-S}, applying a recovery operation such that $\mathcal{R}(\rho_{-S}) = \rho$. We could then consider a thought experiment where instead of throwing away ρ_S and putting it to a mixed state in the depolarizing procedure, we kept it. But then if we apply the error correction procedure, we would recover what would seem to be another copy of ρ_S - thus violating the no cloning theorem!

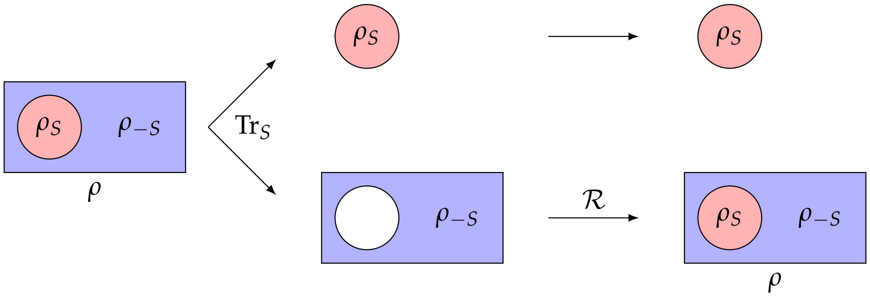

Figure 3.1: Thought experiment to demonstrate the local indistinguishability of quantum codes. We take a code state ρ and trace out a subset S of qubits such that $|S| < d$ (i.e. the S qubits are hit with an erasure error). We then apply the recovery operation to recover the full quantum state ρ. We would then seem to have two copies of ρ_S, which seemingly violates the no-cloning theorem, unless ρ_S contains no information about ρ.

How do we resolve this apparent contradiction? Well, if we are able to recover the entirety of the quantum state ρ by applying an error correction procedure to ρ_{-S} (the information content between the two states is the same - one can interconvert between them via Tr_S and $\mathcal{R}$), it follows that no information about ρ is containing in ρ_S - hence every correctable region contains no information about the encoded state, and we therefore obtain the following Proposition:

Proposition: Local Indistinguishability of QEC Codes

Let C be a $[n, k, d]$ error correcting code and S a subset of the qubits with $|S| < d$. Then the reduced density matrix ρ_S of any code state ρ on the set S is an invariant of the chode, and for all code states ρ, φ holds that $\rho_S = \varphi_S$.

It is worth pointing out that this is a very distinct feature from classical error correction, which is built on repetition - here there is very much the opposite going on!

This property of quantum codes will turn out to be one of the keys in constructing NLTS Hamiltonains out of them - let us explore why by considering this property in more detail.

4 Step 2 - Local Indistinguishability and Circuit Lower Bounds

4.1 An Easy Bound from Indistinguishability

Let us formalize the notion of local indistinguishability:

Definition: Local Indistinguishability

Two states ρ, φ are d-locally indistinguishable if for every subset $S \subset [n]$, $|S| \leq d$, the density matrices ρ_S, φ_S are equal.

The canonical example of local indistinguishable states are the GHZ states:

$$|\text{GHZ}_\pm\rangle = \frac{|0\rangle^{\otimes n} + |1\rangle^{\otimes 1}}{\sqrt{2}} \tag{4.1}$$

which is a locally indistinguishable state as the reduced density matrix on any subset S of qubits with $|S| < n$ is:

$$\rho_{\text{GHZ}_\pm, S} = \frac{|0\rangle\langle 0|^{|S|} + |1\rangle\langle 1|^{|S|}}{2} \tag{4.2}$$

i.e. the maximally mixed state. The states have highly global properties - the ± relative phase can only be observed by looking at all n qubits, and any subset yields no information - they are $n-1$ locally indistinguishable.

This property of states yields lower bounds (which sounds like a very good feature, recalling the conditions on NLTS) on circuits required to produce them:

Lemma: Simple Local Indistinguishabilility Lower Bounds

If $|\rho\rangle, |\varphi\rangle$ are distinct, d-locally indistinguishable states, then neither state can be generated by a circuit of depth $\Omega(\log d)$.

Proof. suppose there is a depth-t circuit of U acting on n qubits with $|\rho\rangle = U|0\rangle^{\otimes n}$, and suppose $2^t < d$. Now consider the Hamiltonian:

$$H_U = U H_0 U^\dagger = U \sum_{i=1}^{n} |1\rangle\langle 1|_i U^\dagger \tag{4.3}$$

The unique ground state of H_0 is $|0\rangle^{\otimes n}$, and the unique ground state of $U H_0 U^\dagger$ is therefore seen to be $|\rho\rangle = U|0\rangle^{\otimes n}$. If $h_i = U|1\rangle\langle 1|_i U^\dagger$ are the local terms of H_U, it acts non-trivially only on the lightcone L_i, so:

$$\langle\varphi|H_U|\varphi\rangle = \sum_{i=1}^{n} \langle\varphi|h_i|\varphi\rangle = \sum_{i=1}^{n} \text{Tr}(h_i \varphi_{L_i}) = \sum_{i=1}^{n} \text{Tr}(h_i \rho_{L_i}) = \langle\rho|H_U|\rho\rangle = 0 \tag{4.4}$$

where in the second equality we can reduce the computation to within the lightcone, and in the third equality we use that $2^t < d$ and so within the lightcone the two states are indistinguishable. So if $|\rho\rangle$ is the ground state of H_U, it is equal to $|\varphi\rangle$; contradiction. This implies that $2^t \geq d$, or $t \geq \log(d)$. We thus get our lower bound. □

The proof yields the useful fact:

Proposition: Hamiltonian Distinguishability

If $|\rho\rangle, |\varphi\rangle$ are d-locally indistinguishable states, and H is a d-local Hamiltonian, $\langle\rho|H|\rho\rangle = \langle\varphi|H|\varphi\rangle$, so H cannot distinguish ρ and φ.

The above Proposition is useful, but incomplete - namely, it is not robust. We will want lower bounds for all states within a small δ radius of the state ρ, and not just for the state ρ (notice that in the previous example, the spectral gap is 1, so it is only robust to $O(\frac{1}{n})$ perturbations, but we want robustness to constant perturbations). To this end we want an improvement, and it comes in the form of the Lemma below:

4.2 A Robust Bound and Well-Spread Distributions

Lemma: Circuit Depth Lower bounds for Distributions

Let D be a probability distribution on $\{0,1\}^n$ generated by measuring hte output of a quantum circuit in the standard basis. If $S_1, S_2 \subset \{0,1\}^n$ satisfy $D(S_1), D(S_2) \geq \mu$, then the depth of the circuit is at least $\frac{1}{3}\log(\frac{\text{dist}(S_1,S_2)^2}{400n\ln\frac{1}{\mu}})$

[2] gives two proofs of this fact; one algebraic and one through the lens of local indistinguishability. We sketch the proof of the second. We have all the mathematical machinery necessary, save for two pieces -

However, we briefly mention some of the ideas that go into the local indistinguishability lens proof - We want to lower bound the circuit depth of a state $|\rho\rangle$ s.t. measuring $|\rho\rangle$ in the standard basis yields the distribution D, and to apply the Simple Local Indistinguishability lower bound, we have to create a state $|\varphi\rangle$ which has the same reduced density matrices - instead, we create a state with approximately the same reduced density matrices. This is done by considering a series of Hamming Balls around S_1, and a sequence of states:

$$|\varphi_i\rangle = \left(\sum_{x\in\{0,1\}^n} (-1)^{x\in B_i}|x\rangle\langle x|\right)|\rho\rangle \tag{4.5}$$

which flip the sign of basis strings inside the ball B_i. With the additional machinery of *approximate ground state projectors* K (which for ground states $|\psi_0\rangle$ satisfy $\left\||\psi_0\rangle\langle\psi_0| - K\right\| \leq \delta$ - in particular "maximal quality" AGSPs can be built from the optimal polynomial approximation to the AND function), one is able to show that one of the $|\varphi_i\rangle$s are approximately locally indistinguishable, and obtains the desired bound on the circuits required to produce the distribution.

This Lemma generalizes (and makes robust) the GHZ example, where we had 50/50 of the mass onto two specific points of $0000\ldots$ and $1111\ldots$. We can actually be more free in the distribution. We just need the masses to be distributed in a way such that both are above some small μ bounded away from zero, and the masses can be in some regions S_1, S_2 (sufficiently far in Hamming distance) rather than single well-separated points. Specifically, we are interested in well-spread distributions:

Definition: Well-spread Distributions

Let D be a probability distribution on $\{0,1\}^n$, and suppose $S_1, S_2 \subset \{0,1\}^n$ satisfy $D(S_1), D(S_2) \geq \mu = O(1)$ and $\mathrm{dist}(S_1, S_2) \geq \omega(\sqrt{n})$. Then D is *well-spread*.

for which the above Lemma then tells us that the circuits associated with such states are above a constant in depth! Therefore, it remains to find Hamiltonians whose low-energy states have this property. For this, we look back to where we found our inspiration for local indistinguishability - to the land of quantum error correction.

5 Step 3 - Finding Well-Spread Distributions in Good Quantum Codes

5.1 Classical Tanner Codes and Approximate Codeword Clustering

Before diving into quantum error correcting codes with the desired property, let's take a look at a specific class of classical codes to guide us. In particular, we consider a classical LDPC (low-density parity check[1]) code known as a Tanner code. It is constructed by taking a regular graph $G = (V, E)$ with degree d and $|V| = n$ vertices and a classical linear code $C \subset \{0,1\}^d$ (with the resulting code denoted as $T = T(C, G)$). The physical bits correspond to the edges of the graph (so the code is on $m = |E|$ bits), and for each vertex $v \in V$, the corresponding local check term verifies that the bits on the edges adjacent to v are members of C.

Tanner codes have the clustering of approximate codewords property:

Property: Clustering of approximate codewords (Classical)

If $y \in \{0,1\}^m$ is a word satisfying most checks, $|\mathbb{H}y| \leq \delta n$ for small δ where $\mathbb{H} \in \mathbb{F}_2^{m \times n}$ is the linear parity check matrix of T[a], then there exists constants c_1, c_2 such that either:

$$|y| \leq c_1 \delta n \text{ or } |y| \geq c_2 n. \tag{5.1}$$

[a] A parity check matrix is one method to specify a classical code - in this formulation, a $[n, k]$ classical code (the set of codewords) consists of the kernel of a $n - k \times n$ parity check matrix $\mathbb{H}$, i.e. all length-n binary vectors that satisfy $\mathbb{H}x = 0$. A code encoding k bits has 2^k possible codewords, so $\dim(\ker \mathbb{H}) = k$ and $\mathbb{H}$ must therefore have linearly independent rows. The parity check formulation elucidates the error detection and recovery procedure, as since $\mathbb{H}y = 0$ for all codewords, if we induce an error $y' = y + e$ then $\mathbb{H}y' = \mathbb{H}e$ which provides the error syndrome from which we can diagnose the error and correct.

Since approximate codewords cluster together, if we look at the low-energy space of a Tanner code, we see a clustered space (specifically, the support is focused on $G^\delta = \{y : |\mathbb{H}y| \leq \delta m\}$ which can be partitioned into clusters).

where the clusters are $O(\delta n)$ in diameter (from $|y| \leq c_1 \delta n$) and they are separated by distance $O(n)$ (from $|y| \geq c_2 n$). This is not quite sufficient to have a well-spread distribution (specifically - we have no way to show that the distribution is not completely skewed on a single cluster; in order to have a well-spread distribution we need two clusters with $O(1)$ mass on them), but it gets the support right, and when we promote it to the quantum code, we get exactly the desired property!

The proof of the above fact involves some slightly technical, but not hard graph theory proofs concerning expansion. We record it here, starting with a definition.

[1] which are defined classically as codes with parity check matrix such that each column has a fixed small $j \geq 3$ number of 1s, and each row contains a small fixed number of 1s - the quantum analog is to be a code which is the simultaneous +1 eigenspace of a collection of projectors such that each projector acts non-trivially on at most l physical qubits, and each physical qubit is acted on by at most l qubits, for $l = O(1)$

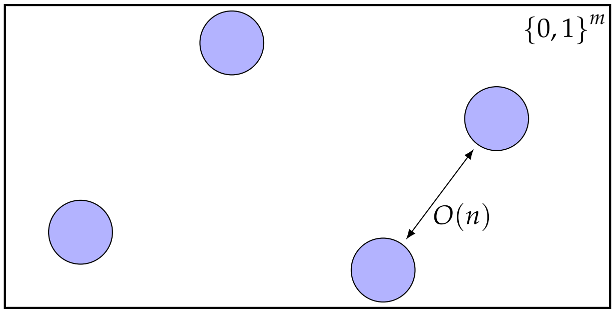

Figure 5.1: Cartoon of low-energy subspace of the Tanner code - because of the clustering of approximate codewords property, the subspace is split into clusters - which has the correct support for a well-spread distribution (but is not quite a well-spread distribution).

Definition: γ-expansion

Let G be a d-left regular bipartite graph between vertex sets L and R. A subset $A \subset L$ is said to be γ-expanding if $|\Gamma(A)| \geq (1-\gamma)d|A|$ where $\Gamma(A) \subset R$ is the set of neighbours of A. We say that G is (γ, α)-small set expanding if every set A of size $\leq \alpha|L|$ is γ-expanding.

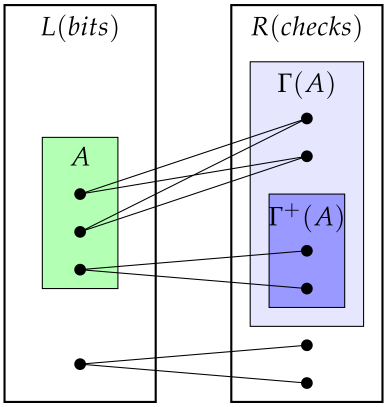

Figure 5.2: Cartoon of small set expansion. Pictured is a 2-left regular bipartite graph, and an A which is γ-expanding for $\gamma \geq \frac{1}{3}$. $\Gamma^+(A)$ is the subset of $\Gamma(A)$ with a unique neighbour in A (which is used in the proof of the following Lemma). We denote L as bits and R as checks as in the following Lemma we consider forming a bipartite graph between these two.

Lemma: Criterion for Clustering of approximate codewords

For a classical error correcting code with check matrix $H \in \mathbb{F}_2^{m \times n}$, draw the interaction graph G between the set of vertices V and the set of checks C, with an edge $v \sim c$ if v participates in the check. If G is (γ, α)-small set expanding for $\gamma < \frac{1}{2}$, then the code has the clustering of approximate codewords property.

Proof. Let $y \in \{0,1\}^n$. If $|y| < \alpha n$ then y is the indicator vector for a small subset $A \subset V$, and $|\Gamma(A)| \geq (1-\gamma)d|A|$. Let $\Gamma^+(A)$ be the subset of $\Gamma(A)$ with unique neighbour in A. Since the number of edges between A and $\Gamma(A)$ is $d|A|$, then

$$d|A| \geq |\Gamma^+(A)| + 2(|\Gamma(A)| - |\Gamma^+(A)|) = -|\Gamma^+(A)| + 2(1-\gamma)d|A| \tag{5.2}$$

Therefore, $|\Gamma^+(A)| \geq (1-2\gamma)d|A|$. Since every check in $\Gamma(A)$ is adjacent to a unique vertex in A, $\Gamma^+(A)$ is a subset of the checks that will be violated by y. Set $c_2 = \alpha$ and $c_1 = \frac{m}{(1-2\gamma)dn}$. Then if $|y| < \alpha n$, then:

$$\delta m \geq |\mathbb{H}y| \geq |\Gamma^+(A)| \geq (1-2\gamma)d|A| = (1-2\gamma)d|y| \tag{5.3}$$

so $|y| < c_1 \delta n$. □

So, if the expansion of the graph G chosen is good by the above criterion, Tanner codes have clustered approximate codewords!

5.2 CSS and Stabilizer Codes

We now go to the quantum realm. Two notions in quantum error correction that will be useful are CSS codes and stabilizer codes. CSS codes are constructed by two classical linear codes C_x, C_z that are kernels of parity check matrices $\mathbb{H}_x, \mathbb{H}_z$. They are combined in such a way to make a quantum code (see [5] 10.4.2, for example) that makes use of the error-correcting properties of the classical code; with C_x correcting the bit flip errors and C_z the phase flip errors.

CSS codes are a class of all stabilizer codes, for which the codespace is the unique $+1$ eigenspace of a set of commuting Pauli operators; i.e. stabilizers of the code. We recall that for a given stabilizer code with stabilizers $\{C_i\}_{i=1}^n$, we can construct a local Hamiltonian as:

$$H = \sum_{i=1}^m h_i := \sum_{i=1}^m \frac{\mathbb{I} - C_i}{2}. \tag{5.4}$$

so if we find a CSS code with the desired low-energy subspace well-spread property, we can construct the Hamiltonian as:

$$H = H_x + H_z = \sum_{w_x} \frac{\mathbb{I} - C_{w_x}}{2} + \sum_{w_z} \frac{\mathbb{I} - C_{w_z}}{2} \tag{5.5}$$

where we sum over the rows w_x/w_z of the parity check matrices $\mathbb{H}_x, \mathbb{H}_z$ (i.e. the stabilizers of C_x, C_z).

5.3 Quantum Tanner Codes and Approximate Codeword Clustering

The desired CSS/stabilizer code we need comes from the quantum promotion of the Tanner code. This has a technical group-theoretic construction[2] - we omit this but refer to [2] for the formal definition. However, hopefully the discussion of the classical code is convincing enough to suggest that the quantum analog of the Tanner code would satisfy the quantum analog of approximate codeword clustering (which is another proof we omit).

Property: Clustering of approximate codewords (Quantum)

Consider a $[n, k, d]$ CSS code defined by classical codes (C_x, C_z). Define G_z^δ as the set of vectors which violate at most a δ-fraction of checks from C_z, i.e. $G_z^\delta = \{y : |\mathbb{H}_z y| \leq \delta m_z\}$ and similarly define G_x^δ. The CSS clusters approximate codewords if there exist constants c_1, c_2, δ_0 such that for sufficiently small $0 \leq \delta < \delta_0$ and for every vector $y \in \{0,1\}^n$:

1. If $y \in G_z^\delta$, then either $|y|_{C_x^\perp} \leq c_1 \delta n$ or else $|y|_{C_x^\perp} \geq c_2 n$
2. If $y \in G_x^\delta$, then either $|y|_{C_z^\perp} \leq c_1 \delta n$ or else $|y|_{C_z^\perp} \geq c_2 n$

[2] One defines two classical tanner codes from obtaining two graphs from the balanced product complex of two Cayley graphs...

5.4 Finshing the Proof

As a final step, we show that the clustering of code-words property of the Quantum tanner code gives us well-spread distributions - the Circuit Depth Lower bounds Lemma then tells us that the Hamiltonian constructed from Eq. (5.5) is NLTS.

Consider a state ψ in the low-energy subspace of H such that $\mathrm{Tr}(H\psi) \leq \epsilon n$.

We then consider three steps.

1. **The supports are focused on $G^{O(\epsilon)}$.** Consider the distributions D_x, D_z generated by measuring ψ in the X/Z bases. The clustering of approximate code-words property then tells us D_x is almost completely supported on $G_x^{O(\epsilon)}$ and equivalently for D_z (and this amount of support can be made tighter by considering states of sufficiently small energy - the proof in the thesis takes states of energy $\epsilon/200$, for example) - this is as in the case for the classical Tanner code.

2. **The supports are well-clustered.** Similarly, the clustering of approximate code-words means that $G_x^{\epsilon_1}$ (And analogously $G_z^{\epsilon_2}$) can be partititoned into clusters. Formally, this is done by taking $x, y \in G_x^{\epsilon_1}$ and defining an equivalence realtion by $x \sim y$ iff $|x \oplus y|_{C_x^{\perp}} \leq 2c_1\epsilon_1 n$. The clusters in this partition (by the definition of the equivalence relation) are separated by $\geq c_2 n$, so they are also well-separated. So, both in the X and Z bases we have clustered distributions! - This is again, as in the classical tanner code!

3. **The distributions are not concentrated on any one cluster.** We are close to having a well-spread distribution, but we need to show that the distributions are not completely focused on a single cluster.

   We do not know how to show this property for D_z, D_x individually. However, we can show that it is impossible for both to be concentrated on any one cluster! This follows from the following uncertainty principle:

   Proposition: Uncertainty of Distributions

   Given a state ψ and measurement distributions D_x, D_z, for all subsets $S, T \subset \{0,1\}^n$, $D_x(T) \leq 2\sqrt{1 - D_z(S)} + \sqrt{|S| \cdot |T|/2^n}$.

   which tells us that if there is a cluster B_z that D_z is almost completely concentrated on (i.e) that $D_x(B_x)$ is bounded above, i.e. D_x is not completely concentrated on B_x. The quantum picture has an edge over the classical picture in that there are not just one clustered universe, but two which control the spread of mass of the other.

   From this, we obtain that one of D_z, D_x is well-spread, and hence that D_z, D_x requires a circuit depth of $\Omega(\log n)$ and hence ψ requires a circuit depth of $\Omega(\log n)$ - i.e. all low-energy ψ of the Hamiltonian corresponding to the quantum Tanner code cannot be generated by a constant depth circuit, and so it is an NLTS Hamiltonian. □

6 Future Directions

The proof of NLTS has no bearing on whether QPCP is true, but has taken away one of the most obvious lines of attack towards refuting it - low depth circuits are no longer a generic enough ansatz to classically approximate all local Hamiltonians! But perhaps a different description is possible, e.g. if all low-energy states of local Hamiltonians could be described by a Clifford circuit, we can use the Gottesman-Knill theorem to efficiently describe the state and calculate its energy and the QPCP conjecture could be disproven.

However, the techniques used in the result are unlikely to directly build towards a proof of the QPCP theorem. Proofs of the classical PCP theorem have already been shown to be incompatible with quantization - a birds-eye reason being the incompatibility of quantum error correction with the (classical) error correction techniques used - essentially, violations of the no-cloning theorem. In fact, the property of local indistinguishability used to prove NLTS is actually a roadblock in a QPCP proofs - as classical PCP theorems can be framed as an elegant locally testable (!!) code plus a satisfying assignment for the formula.

In studying this result, I have two general take-home messages (beyond the actual very fascinating result!). The first is that we cannot always let our intuition lead us astray - although the NLTS Hamiltonians presented here are highly engineered objects, they nevertheless exist and fly in the face of our intuition considering the robustness of entanglement. The second is the usefulness in reframing problems and combining sources of inspiration - NLTS was successfully able to be proven by reducing a problem about computation to one about Hamiltonians and entanglement, and by taking inspiration from developments in the seemingly distant field of quantum error correction.

References

[1] Anurag Anshu, Nikolas P. Breuckmann, and Chinmay Nirkhe. Nlts hamiltonians from good quantum codes, 2022.

[2] Chinmay Nirkhe. *Lower bounds on the complexity of quantum proofs*. PhD thesis, EECS Department, University of California, Berkeley, Aug 2022.

[3] Chinmay Nirkhe. Nlts hamiltonians from codes. Presented at Simons Institute Quantum Colloquium, 2022.

[4] Dorit Aharonov, Itai Arad, and Thomas Vidick. The quantum pcp conjecture. 2013.

[5] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.