The Detectability Lemma and Quantum Gap Amplification

Dorit Aharonov*, Itai Arad†, Zeph Landau‡ and Umesh Vazirani §

November 17, 2008

Abstract

The quantum analog of a constraint satisfaction problem is a sum of local Hamiltonians - each (term of the) Hamiltonian specifies a local constraint whose violation contributes to the energy of the given quantum state. Formalizing the intuitive connection between the ground (minimal) energy of the Hamiltonian and the *minimum* number of violated constraints is problematic, since the number of constraints being violated is not well defined when the terms in the Hamiltonian do not commute. The detectability lemma proved in this paper provides precisely such a quantitative connection. We apply the lemma to derive a quantum analogue of the classical gap amplification lemma of random walks on expander graphs. The quantum gap amplification lemma holds for local Hamiltonians with expander interaction graphs. Our proofs are based on a novel structure imposed on the Hilbert space, which we call the XY decomposition, which enables a reduction from the quantum non-commuting case to the commuting case (where many classical arguments go through).

The results may have several interesting implications. First, proving a quantum analogue to the PCP theorem is one of the most important challenges in quantum complexity theory. Our quantum gap amplification lemma may be viewed as the quantum analogue of the first of the three main steps in Dinur's PCP proof. Quantum gap amplification may also be related to fault tolerance of adiabatic computation, a model which has attracted much attention but for which no fault tolerance theory was derived yet. Thirdly, the detectability lemma, and the XY decomposition provide a handle on the structure of local Hamiltonians and their ground states. This may prove useful in the study of those important objects, in particular in the fast growing area of "quantum Hamiltonian complexity" connecting quantum complexity to condensed matter physics.

*School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel.
†Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, CA.
‡Department of Electrical Engineering and Computer Sciences, University of California at Berkeley,CA.
§Department of Electrical Engineering and Computer Sciences, University of California at Berkeley,CA.

1 Introduction

There is a close analogy between two fundamental notions from computational complexity theory and quantum physics: constraint satisfaction problems and the ground energy of local Hamiltonians. Each term in the local Hamiltonian specifies a local constraint whose violation contributes to the energy of the given quantum state. Hence the energy of the quantum state corresponds intuitively to the number of violated quantum constraints. A canonical example of this is the correspondence between the classical Cook-Levin theorem and its quantum analogue proved by Kitaev [1]. Kitaev showed that estimating the ground energy of a local Hamiltonian to within inverse polynomial accuracy (the quantum analogue of determining the minimal number of violated constraints) is complete for the quantum analog of NP, namely QMA.

But how accurate is this intuitive correspondence between the energy of a state and the number of violated quantum constraints? The main issue is that in the quantum case, the terms of the Hamiltonian do not commute in general. This means that it is not even a meaningful question to ask: how many constraints are violated by a given state? Or in keeping with the probabilistic nature of quantum physics: what is the probability that the given state violates at least k constraints?

The quantum detectability lemma, proved in this paper, provides a way of making sense of and answering these questions. It partitions the terms of the Hamiltonian into a constant number of sets of mutually commuting constraints, and shows that if the energy of a state is large, then so is the probability that a large number of constraints are violated in at least one of these sets. The partitioning is possible under the mild restriction that every particle (qubit) in the local Hamiltonian participates in a bounded number of constraints, and each term in the Hamiltonian is chosen from a finite set of possibilities. Clearly for any quantum state, there is some layer that accounts for a constant fraction of the energy of the state. The simplest version of the detectability lemma (the $\ell = 0$ case) asserts that the probability that at least one constraint in that layer is violated is proportional to the energy of the state. More precisely, this is the probability of getting a non-zero result if the energy of the state is measured with respect to that layer. The proof of even this simple case (even for two layers) is non-trivial, and relies on a certain decomposition of the Hilbert space described below, that we call the XY-decomposition. We remark that this case can be viewed as the converse of Kitaev's geometrical lemma [1] providing a lower bound on the ground energy of a sum of Hamiltonians in terms of the angle between their ground spaces; we explain this in Appendix D. The general form of the lemma lower bounds the probability that at least ℓ constraints are violated in this layer. The bound is non-trivial for ℓ smaller than some constant times the energy of the state. The proof of the general lemma is somewhat more complicated and is deferred to the appendix.

The XY-decomposition, which lies at the heart of the proof of the quantum detectability lemma, is interesting in its own right. It captures a structural relationship between the ground spaces of the different layers of the Hamiltonian. It partitions the Hilbert space into a tensor product of local spaces and then further decomposes each of these local spaces into commuting and non-commuting parts with respect to the Hamiltonian. Very roughly this allows for local analysis of the actions of the individual terms of the Hamiltonian the result of which is a parameter for which these actions can be shown to have exponential decay when restricted to the non-commuting part of the decomposition. In physics-speak, this shows that any local Hamiltonian system is a small parameter system.

Classical gap amplification, first proved in the context of saving random bits in RP and BPP amplification [2, 3], is a basic primitive in complexity theory. The idea is that if one is interested in amplifying the probability of hitting a given subset of the nodes (or edges) in a graph, then if the graph is an expander, a random walk would do almost as well as picking the nodes (or edges) independently. More generally a constraint satisfaction problem is represented by a hypergraph, with each hyperedge corresponding to a constraint. To amplify the gap between the acceptance and rejection probability, one considers the "t-step walk" on the hypergraph. Now, if the hypergraph is expanding, one can show that the gap gets amplified by a factor of $\Omega(t)$. This idea has since found many other important implications, for example, in Dinur's proof of the PCP theorem [4]. In this paper we prove a quantum analogue of the classical amplification lemma: the hyper-constraints are also generated from t terms in the original Hamiltonian which form a walk in the interaction graph. Each new hyper-constraint is the projection on the intersection of the t constraints on the walk. We show that if the original interaction graph was an expander, the average ground energy per term of the new Hamiltonian (consisting of the hyper-constraints) is $\Omega(t)$ times the average ground energy per term of the original Hamiltonian, thus establishing a bound similar to the classical lemma. The proof relies critically on the quantum detectability lemma, along with the classical analysis of walks on expander graphs.

Discussions and Possible Implications: The results in this paper are related to several important open problems in quantum computation complexity. First, the study of the computational complexity of local Hamiltonians has blossomed over the last few years, and touches upon efficient simulation of quantum systems and theoretical condensed matter physics. The techniques developed in this paper, the XY-decomposition and the quantum detectability lemma, can be expected to contribute to our understanding of this new area.

Second, the PCP theorem is arguably the most important development in computational complexity theory over the last two decades. Is there a quantum analogue? One natural formulation is the following: suppose we are given a local Hamiltonian on n qubits with the promise that the ground energy is either 0 or at least $1/p(n)$ for some polynomial $p(n)$. Is there a way to map this to a new local Hamiltonian such that the ground energy is either 0 or $\Omega(n)$? (see Section 7 for more precise definition). Proving such a quantum PCP theorem is a major challenge in quantum complexity theory; it would have implications for our understanding of inapproximability results of quantum complexity problems, quantum fault tolerance, the understanding of entanglement and notions such as no-cloning, as well as on the basic notion of energy gap amplification in condensed matter physics.

Our quantum gap amplification lemma can be viewed as a very weak form of the above statement. The problem is that checking the new t-walk constraints requires t queries, which is too large even if we wish to check a single constraint. Dinur's proof of the classical PCP theorem combines this kind of gap amplification with two other steps. In this sense quantum gap amplification is a possible first step towards emulating the outline of Dinur's proof in the quantum setting.

Gap amplification is tightly connected to adiabatic quantum computation, a model of quantum computation equivalent in power to the standard one, which has attracted considerable attention ([5, 6, 7, 8, 9, 10, 11, 12] and more). In adiabatic computation, the system evolves under a Hamiltonian with a non-negligible spectral gap between the ground state and the next excited state. Physical intuition suggests that such a model might be inherently robust to thermal noise [13]. Despite work on the subject [14, 15, 16, 17], including the development of quantum error correcting codes tailored for adiabatic evolution [18], an analogue to the threshold result of the standard model [19, 20, 21] is still missing. Can the spectral gap in adiabatic computation be amplified to a constant, to provide fault tolerance? This is probably impossible when the system of qubits is arranged on a line, since even though such a system can be adiabatically universal when the gap is inverse polynomial [22], Hastings has showed [23, 24] that adiabatic evolution in one dimension with constant spectral gap can be simulated efficiently classically. However, it may very well be true that amplification of the spectral gap in some well defined sense is possible, if the underlying geometry is that of an expander. It is likely that a proof of a quantum PCP theorem would pave the way to such a result, though we should caution that there is no proof showing such an implication.

Open Problems In this paper we handle the restricted case in which the local terms in the Hamiltonian are projections. We leave the general case for future work.

As a benchmark open problem, we pose the following question: prove an exponential size quantum PCP, in analogy with the first classical PCP results [25]. This already seems to require some non-trivial work in quantum information theory. We note that it is possible to prove a quantum PCP theorem where the proof is of doubly exponential size. Proving a quantum analogue of Dinur's degree reduction, which allows reducing the degree of the graph of interactions in the Hamiltonian, is a major open problem.

A related problem is to improve the parameters in current perturbation gadgets [26, 27, 28, 29] significantly.

2 Background - Local Hamiltonians and Local Projections

A k-local Hamiltonian H on n qubits is an operator $H : \mathbb{B}^{\otimes n} \to \mathbb{B}^{\otimes n}$ that can be written as a sum $H = \sum_{i=1}^{M} H_i$, where $M = poly(n)$ and every H_i is a Hermitian operator acting on at most k qubits. In this paper we restrict attention to the case where the H_i operators are projections. We will usually denote them by Q_i: $H = \sum_{i=1}^{M} Q_i$. Another assumption that we use is that every projection intersects with only a finite number of other projections. Together with the k-locality, this implies that the projections can be partitioned into a constant number, denoted g, of subsets (which we call *layers*) such that the projections in each layer are non-intersecting and thus commuting. We denote a system satisfying the above restriction by a k-QSAT system.

For a state $|\psi\rangle$, $\langle\psi|H|\psi\rangle = \sum_i \langle\psi|H_i|\psi\rangle$ is called the *energy* of the state. In our case, this can be any number between 0 and M. The minimal energy of the system is the lowest eigenvalue of the Hamiltonian, and is denoted by ϵ_0.

Deciding whether ϵ_0 is above some threshold a or below a threshold b with $a - b > 1/poly(n)$ is known as the k-local Hamiltonian problem (which is complete for Quantum NP). It can be viewed as the quantum analog of the $k-$SAT problem. We often refer to the projections Q_i as *constraints*, and when $\langle\psi|Q_i|\psi\rangle > 0$ we say that Q_i is violated (with respect to the state $|\psi\rangle$). Similarly, when $\langle\psi|Q_i|\psi\rangle = 0$, we say that Q_i is satisfied, or that $|\psi\rangle$ is in the accepting subspace of Q_i.

3 The XY decomposition

We consider a k-QSAT system over n qubits with M constraints that can be arranged in g layers. Let us start by describing at a high level the XY-decomposition and how it is applied:

We start with a decomposition of the Hilbert space into a tensor product of local spaces, and restrict our attention to only those terms of the Hamiltonian that act non-trivially on exactly one of these spaces. The way the actual decomposition is carried out depends upon an ordering of the layers, and is described in the pyramid construction below. Each of these local spaces can now be further decomposed into a direct sum of subspaces according to whether all the terms of the Hamiltonian acting on this subspace commute (X) or not (Y). This defines a natural XY-decomposition of any state. The main point of the XY-decomposition is that it allows us to capture some structural relationship between the ground spaces of the different layers of the Hamiltonian. The starting point for this is the observation that in each Y subspace there is some finite angle between the ground spaces of the different layers. Now stepping back, we can decompose the tensor product of all the local spaces into subspaces according to the number of Y components. The actions of the Hamiltonian on the local spaces collectively ensure that if we start from an arbitrary state and successively project it onto the ground spaces of the different layers, then the weight of the resulting state in each subspace decays exponentially in the number of Y components. This is a key property used in the proof of the detectability lemma.

3.1 Pyramids and Pyramid projections

We partition the Hilbert space of the k-QSAT system into a product of subspaces by defining the notion of a "pyramid", a special "connected" subset of the constraints, as follows. First, we arbitrarily order the layers from 1 to g. A pyramid is created by picking its apex - a constraint in the first layer

- and for each successive layer picking all constraints that intersect with the set of constraints picked in previous layers. We denote the Hilbert space of the qubits which participate in the pyramid by H_{pyr}. We now consider any maximal set of disjoint pyramids as illustrated in Fig. 1. Clearly the entire Hilbert space can be written as a tensor product of the pyramid spaces H_{pyr}, and constraints from different pyramids commute.

In the next step, we decompose the Hilbert space H_{pyr} of the first pyramid into a direct sum of subspaces $\{X_j\}$ and a a subspace Y: every space X_j is made of vectors which are simultaneous eigenvectors of *all* projections in the pyramid. Moreover, in every such X_j, each projection is allowed to take only one value - 0 or 1. Then Y is defined to be the residual subspace, i.e., the subspace that is orthogonal to all the X_j subspaces inside H_{pyr}. Clearly, all these spaces are orthogonal to each other. We refer to Y as the "non-commuting" part of the Hilbert space H_{pyr}; all other subspaces correspond to the "commuting parts". Of course, this decomposition can be done for every one of the pyramids.

We denote a sector of the XY decomposition by a string ν. ν specifies either an X_i space or a Y space at each location, and we define

$$|\nu| \stackrel{\text{def}}{=} \text{No. of } Y \text{ sites in } \nu \ . \tag{1}$$

We also define P_ν to be the projection into the tensor product of these spaces. Note that P_ν is by itself a product of all the corresponding P_{X_i}, P_Y projections. Every state in $|\psi\rangle \in H$ can therefore be written in the XY decomposition as:

$$|\psi\rangle = \sum_\nu P_\nu|\psi\rangle \stackrel{\text{def}}{=} \sum_\nu \lambda_\nu|\psi_\nu\rangle \ . \tag{2}$$

We will in fact use not one XY decomposition but a finite number of them. It is easy to see that there exists a constant $f(k,g)$ (independent of n) of XY decompositions such that every constraint in the top layer appears in one pyramid top in one of the XY decompositions, and so all top constraints are "covered" by one of the decompositions. But for now we fix one XY decomposition.

3.2 Commutation relations between projections inside the pyramids

For a fixed pyramid, we denote the operators which act on H_{pyr} and project of on the subspaces $\{X_j\}_j$ and Y by $\{P_{X_j}\}_j$, P_Y respectively. It is easy to verify that the projections form a valid decomposition of H_{pyr}: $[P_{X_i}, P_{X_j}] = [P_{X_i}, P_Y] = 0$ and $P_Y + \sum_j P_{X_j} = \mathbb{1}$. Moreover, the following holds:

- For any constraint Q in the pyramid: $[Q, P_{X_j}] = [Q, P_Y] = 0$.
- For every two constraints Q_1, Q_2 in pyramid and every X_j: $P_{X_j}[Q_1, Q_2]P_{X_j} = 0$.

3.3 The parameter θ

Definition 3.1 (The parameter θ) *Fix a pyramid. Consider the product $Q_0 \cdot Q_1 \cdot \ldots \cdot Q_N$ where every Q_i is either a projection from the pyramid or its complement, and every pyramid projection (or its complement) appears exactly once. Y does not contain any common eigenvector of* all *those projections. Hence there exists a constant $0 < \theta < 1$ such that for any possible pyramid in the system, and any order in which the constraints are chosen to appear in the product,*

$$\|P_Y \cdot Q_0 \cdot \ldots \cdot Q_N \cdot P_Y\| \le \theta \ . \tag{3}$$

$0 < \theta < 1$ *is a constant that depends only on the family of constraints and on the constant g (which determines the maximal number of constraints in a pyramid).*

3.4 The projections for the two layers case: Π_{red} and Π_{blue}

To introduce the behavior of the exponential decay, we start with the simpler case of two layers, which we call "blue" and "red". The structure in this case is illustrated in Fig. 2.

We define Π_{red} (Π_{blue}) to be the projection on the tensor product of the zero (accepting) subspaces of all the terms in the red (blue) layer. We call these subspaces the $\ell = 0$ subspaces because they contain $\ell = 0$ violations of the constraints in the layer. Later on, we will use the same notation to generalize the discussion to larger ℓ's.

A central observation is that we can present Π_{red}, Π_{blue} according to the pyramids structure. Take for example Π_{blue}. It is a projection into the zero energy space of the blue layer. It can therefore be written as the product $\Pi_{blue} = (\mathbb{1} - Q_1) \cdot (\mathbb{1} - Q_2) \cdots$ with Q_i being the blue constraints. We can therefore write it as $\Pi_{blue} = \Delta_{blue} R_{blue}$, where Δ_{blue} contains the constraints that are inside the pyramids (in this case these are simply the pyramids' tops), and R_{blue} contains the constraints outside the pyramids. Similarly, we define $\Pi_{red} = \Delta_{red} R_{red}$.

Because of the pyramids structure, the support of R_{red} and Δ_{blue} are non-intersecting, hence $R_{red}\Delta_{blue} = \Delta_{blue} R_{red}$ (note that this does not hold for R_{blue} and Δ_{red}), hence

$$\Pi_{red}\Pi_{blue} = \Delta_{red}\Delta_{blue} R_{red} R_{blue} \ . \tag{4}$$

3.5 The exponential decay for the 2 layers $\ell = 0$ case

We now prove the exponential decay behavior. Let us first coarse-grain the XY decomposition by gathering together all sectors with the same number of Y spaces. In other words, for every integer $0 \le s \le M$, define a projection

$$P_s \stackrel{\text{def}}{=} \sum_{|\nu|=s} P_\nu \ . \tag{5}$$

Then this is still a valid decomposition as the P_s are orthogonal to each other and $\sum_{s=0}^{m} P_s = \mathbb{1}$. The exponential decay lemma states that if we apply this decomposition to some state *after* applying the Π_{blue} and Π_{red} projections, then we can upper bound the weight of the s sector in terms of θ^s.

Lemma 3.2 (Exponential-decay lemma for $\ell = 0$) *Let $|\psi\rangle$ be an arbitrary (normalized) state, and consider the following normalized state $|\Omega\rangle \stackrel{\text{def}}{=} \frac{1}{x}\Pi_{red}\Pi_{blue}|\psi\rangle$ and its coarse grained XY decomposition $|\Omega\rangle = \sum_s P_s|\Omega\rangle \stackrel{\text{def}}{=} \sum_s \lambda_s|\Omega_s\rangle$. Then there exist weights $\{\eta_s\}$ such that $\sum_s \eta_s^2 \le 1$, and*

$$\lambda_s \le \frac{1}{x}\theta^s \eta_s. \tag{6}$$

Proof: To prove this claim, we take one step backwards, and write $|\Omega\rangle$ in terms of the fine-grained XY decomposition: $|\Omega\rangle = \sum_\nu \lambda_\nu |\Omega_\nu\rangle$. Then $\lambda_\nu^2 = \frac{1}{x^2}\langle\psi|\Pi_{blue}\Pi_{red}\ P_\nu\ \Pi_{red}\Pi_{blue}|\psi\rangle$. We use Eq. (4) and write $\Pi_{blue}\Pi_{red}\ P_\nu\ \Pi_{red}\Pi_{blue} = R_{blue}R_{red}\Delta_{blue}\Delta_{red}\ P_\nu\ \Delta_{red}\Delta_{blue}R_{red}R_{blue}$, and as P_ν commutes with $\Delta_{red}, \Delta_{blue}$, this is equal to $R_{blue}R_{red}P_\nu\ \Delta_{blue}\Delta_{red}\Delta_{blue}\ P_\nu R_{red}R_{blue}$. It follows that

$$\lambda_\nu^2 \le \frac{1}{x^2}\|P_\nu\Delta_{blue}\Delta_{red}\Delta_{blue}P_\nu\| \cdot \|P_\nu|\Phi\rangle\|^2 \tag{7}$$

with $|\Phi\rangle \stackrel{\text{def}}{=} R_{red}R_{blue}|\psi\rangle$. Let us estimate $\|P_\nu\Delta_{blue}\Delta_{red}\Delta_{blue}P_\nu\|$. Every operator in the product factors into a product of operators over every pyramid. Consider a pyramid site which is projected (by P_ν) into a Y subspace. For brevity, call it a Y pyramid. Let Q_{blue} be the blue constraint (the pyramid's top) and $Q_1, \ldots, Q_N$ the red constraints. We have $P_Y \cdot Q_{blue} \cdot Q_1 \cdots Q_N \cdot Q_{blue} \cdot P_Y = (P_Y \cdot Q_{blue} \cdot Q_1 \cdots Q_N \cdot P_Y) \cdot (P_Y \cdot Q_N \cdots Q_1 \cdot Q_{blue} \cdot P_Y)$, where we have used the fact that P_Y commutes with the constraints inside the pyramid. From Eq. (3), its norm is smaller or equal to θ^2, and since there are $|\nu|$ such Y sites, we deduce that $\|P_\nu \Delta_{blue}\Delta_{red}\Delta_{blue}P_\nu\| \le \theta^{|2\nu|}$. All together, this leads to $\lambda_\nu^2 \le \frac{1}{x^2}\theta^{2|\nu|}\langle\Phi|P_\nu|\Phi\rangle$, and summing over all ν with $|\nu| = s$, we obtain $\lambda_s^2 = \sum_{|\nu|=s}\lambda_\nu^2 \le \frac{1}{x^2}\theta^{2s}\eta_s^2$, where we have defined $\eta_s^2 \stackrel{\text{def}}{=} \langle\Phi|P_s|\Phi\rangle = \sum_{|\nu|=s}\langle\Phi|P_\nu|\Phi\rangle$. ■

One can also make a similar statement for more than two layers; the proof follows very similar lines. Here we do not state this lemma since the following result implies it as a special case.

3.6 Generalizing to many layers and $\ell > 0$

As was mentioned in the previous section, the Π_{red}, Π_{blue} projections can be regarded as $\ell = 0$ projections since they project into the subspaces of $\ell = 0$ violations in either the red layer or the blue layer. Generally, when we consider the i'th layer, we can define the projection $\Pi^{(i)}_{\le\ell}$ which projects into the subspace of ℓ or less violations in that layer. This allows us to derive an exponentially decaying bound on the similarly defined coefficients λ_s, except now the bound will contain some combinatorial factors depending on ℓ.

Lemma 3.3 (Exponential decay lemma for general ℓ) *Consider a k-QSAT system with g layers and M projections, drawn from a finite set that is characterized by a parameter $0 < \theta < 1$. Let $0 \le \ell \le M$ be an integer and let $|\psi\rangle$ be an arbitrary (normalized) state. Consider the normalized state*

$$|\Omega\rangle \stackrel{\text{def}}{=} \frac{1}{x}\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle \ , \tag{8}$$

and its coarse grained XY decomposition $|\Omega\rangle = \sum_s P_s|\Omega\rangle \stackrel{\text{def}}{=} \sum_s \lambda_s|\Omega_s\rangle$. Then there exist weights $\{\eta_0, \ldots, \eta_M\}$ such that $\sum_s \eta_s^2 \le 1$, and for every $s \ge \ell$,

$$\lambda_s \le \frac{1}{x}k^{g^2\ell}\left(\frac{\ell+1}{\ell!}\right)^g s^{g\ell}\theta^s\eta_s \ . \tag{9}$$

The proof of the above claim is more involved than the 2-layers, $\ell = 0$ case, and will therefore be given in Appendix E. The main difficulty here is the fact that when $\ell > 0$, the projections $\Pi^{(i)}_{\le\ell}$ are no longer a simple product of projections as in the $\ell = 0$ case. Instead, they can be thought of as a huge sum over similar products of projections, where each such product projects into a certain possible configuration of ℓ or less violations. This complicates the analysis, but other than that, the proof follows the same outline of the 2-layers, $\ell = 0$ case.

4 The detectability lemma for two layers and $\ell = 0$

In this section we prove the detectability lemma for the special case of two layers and $\ell = 0$. The proof is considerably simpler than the general proof, yet it demonstrates the ideas of the general case.

The general setup is similar to the one in Sec. 3.4. We consider a k-QSAT system with $\epsilon_0 > 0$ that can be arranged in two layers. The first layer is called the "blue layer" and the second layer is the "red layer" (see Fig. 2). We define Π_{red} as the projection that projects into the accepting (zero) space of the red, and similarly Π_{blue}. In addition, we assume that the constraints are drawn from a finite family of constraints with a parameter $0 < \theta < 1$ (see Sec. 3). Then the 2-layers, $\ell = 0$ detectability lemma is:

Lemma 4.1 (The detectability lemma for two layers and $\ell = 0$)
There exists a function $f(k)$ such that for every normalized state $|\psi\rangle$,

$$\max\left\{\|(\mathbb{1} - \Pi_{red})|\psi\rangle\|^2, \|(\mathbb{1} - \Pi_{blue})|\psi\rangle\|^2\right\} \ge \frac{1}{8}\Delta^2(0) \ , \tag{10}$$

where

$$\Delta^2(0) \stackrel{\text{def}}{=} 1 - \frac{1}{(\epsilon_0/f)\frac{(1-\theta^2)^3}{\theta^2} + 1} \ . \tag{11}$$

Appendix A proves, using simple algebra, that it suffices to prove the following lemma.

Lemma 4.2 *For every normalized state* $|\psi\rangle$*,* $\|\Pi_{red}\Pi_{blue}|\psi\rangle\|^2 \le 1 - \Delta^2(0)$ *.*

Proof of Lemma 4.2: Using the notation of Sec. 3.5, we define $|\Omega\rangle \stackrel{\text{def}}{=} \frac{1}{x}\Pi_{red}\Pi_{blue}|\psi\rangle$ and $x \stackrel{\text{def}}{=} \|\Pi_{red}\Pi_{blue}|\psi\rangle\|$. We wish to prove an upper bound for x. The idea of the proof is to estimate the total energy of $\langle\Omega|E|\Omega\rangle$. This energy has no contributions from the red layer since $|\Omega\rangle$ has been projected by Π_{red}, and so we may write:

$$\epsilon_0 \le \langle\Omega|\, E_{blue}\, |\Omega\rangle \ . \tag{12}$$

We will find an upper-bound for $\langle\Omega|\, E_{blue}\, |\Omega\rangle$ in terms of θ, k, and this would give us an inequality for x, θ, k, ϵ_0. Inverting that inequality will give us the desired result.

To estimate E_{blue} we consider first one possible XY decomposition. Let E^{top} be the energy of all the blue constraints from the pyramids in this decomposition - the "tops" of the pyramids. The main effort would be to find an upper-bound for $\langle\Omega|E^{top}|\Omega\rangle$. Once we do that, we can then repeat this process with other sets of pyramids (namely, other XY decompositions) until we cover all the blue constraints. All in all, there is a finite number $f(k)$ of XY decompositions that are needed for that. Therefore,

$$\epsilon_0 \le \langle\Omega|E_{blue}|\Omega\rangle \le f(k)\langle\Omega|E^{top}|\Omega\rangle \ . \tag{13}$$

Hence, it remains to bound $\langle\Omega|E^{top}|\Omega\rangle$. We start by applying the fine- and coarse-grained XY decompositions to $|\Omega\rangle$:

$$|\Omega\rangle = \sum_\nu \lambda_\nu|\Omega_\nu\rangle = \sum_s \lambda_s|\Omega_s\rangle \ . \tag{14}$$

Then as the XY projections commute with the projections in E^{top}, we get

$$\langle\Omega|E^{top}|\Omega\rangle = \sum_s \lambda_s^2\langle\Omega_s|E^{top}|\Omega_s\rangle \ . \tag{15}$$

Claim 4.3 $\langle\Omega_s|E^{top}|\Omega_s\rangle \le s$ *.*

The proof is short and can be found in Appendix C. Essentially, it follows from looking at the fine-grained XY sectors that contribute to s and noting that only the Y sites can contribute energy.

We can now use the above bound inside Eq. (15), together with the bound $\lambda_s^2 \le \frac{1}{x^2}\theta^{2s}\eta_s^2$, which follows from the exponential decay lemma 3.2. We get:

$$\langle\Omega|E^{top}|\Omega\rangle = \sum_s s\lambda_s^2 \le \sum_s \frac{1}{x^2}s\theta^{2s}\eta_s^2 \ . \tag{16}$$

In principle, inserting this into Eq. (13) we could simply bound every η_s^2 by 1, and, rearranging, get a bound on x^2. However, this bound would be bad for very small ϵ_0. Luckily, we can derive a stronger bound on η_s^2 for $s \ge 1$:

Claim 4.4 *For every* $s \ge 1$*,* $\eta_s^2 \le \frac{1-x^2}{1-\theta^2}$*.*

This follows by simple algebraic calculation, using the connection between x and the coefficients λ_s for $s \ge 1$. Claim 4.4 is proved in Appendix B. We can now finish the proof. Following Eq. (16) and Eq. (13) we have: $\epsilon_0 \le \langle\Omega|E_{blue}|\Omega\rangle \le f(k)\langle\Omega|E^{top}|\Omega\rangle \le f(k)\frac{1-x^2}{x^2}\cdot\frac{\theta^2}{(1-\theta^2)^3}$, which implies the result. ■

5 The (general) detectability lemma

The detectability lemma can be generalized for more than 2 layers and for $\ell > 0$. This generalization gives us a more detailed picture of the energy distribution. This is important when ϵ_0 is much bigger than 1 but is still smaller than its maximal value M. In such a case, the detectability lemma asserts that not only there exists a layer in which some violations are detectable - but that there must be a layer in which ℓ or more violations are detectable. In other words, it forbids a situation in which in all the layers the violations are of only few constraints and there is $1/poly$ weight on very high violations (so that the total energy $\geq \epsilon_0$). For the lemma to hold, we require that ℓ – the number of violations – does not exceed some normalized version of the minimal energy ϵ_0.

Lemma 5.1 (The general detectability lemma) *Consider a k-QSAT system with g layers and a ground energy $\epsilon_0 > 0$. Let $\Pi^{(i)}_{>\ell}$ denote a projection into the space of more than ℓ violations in the i'th layer. Then there exist integer functions $r(\theta, k, g), f(k, g) > 1$ such that for every $0 \leq \ell < \frac{1}{r}\left(\frac{\epsilon_0}{f} - \frac{1}{1-\theta}\right)$ and every normalized state $|\psi\rangle$ there is at least one layer i in which:*

$$\|\Pi^{(i)}_{>\ell}|\psi\rangle\|^2 \geq \frac{1}{(2g)^2}\Delta^2(\ell) \ . \tag{17}$$

$\Delta(\ell)$ is a function of $\ell, \epsilon_0, \theta, k, g$, and is given by

$$\Delta^2(\ell) = \begin{cases} 1 - \frac{1}{(\epsilon_0/f)\frac{(1-\theta^2)^3}{\theta^2}+1} & , \quad \ell = 0 \\ 1 - \frac{1}{1-\theta} \cdot \frac{1}{(\epsilon_0/f) - r\ell} & , \quad \ell > 0 \end{cases} \ . \tag{18}$$

The proof of the general detectability lemma is deduced from the exponential decay lemma for general ℓ, using similar reasoning to how the simpler detectability lemma is deduced from the $\ell = 0$ exponential decay lemma. However, the technical details are much more involved due to the same combinatorial factors that appear when moving from $\ell = 0$ to $\ell > 0$ in the exponential decay lemmas. The full proof in given in Appendix F.

6 The quantum gap amplification lemma

In this section we define a prove a quantum version of the well-known classical gap-amplification lemma. For completeness, we include a description and proof of the classical lemma in Appendix H.

The setting of the quantum amplification lemma is a natural generalization of the classical setting. We consider a d-regular expander graph $G = (V, E)$ with a second-largest eigenvalue $0 < \lambda(G) < 1$. On top of G we define a k-QSAT system as follows. We identify every vertex with a qudit of dimension q. Every edge $e \in E$ is identified with a projection Q_e on the two qudits that are associated with the vertices of the edge. This defines k-QSAT system with $k = 2\log(q)$ and a Hamiltonian $H = \sum_{e\in E} Q_e$. For any state $|\psi\rangle$, we define the quantum UNSAT of the system to be the average energy of the edges:

$$\mathrm{QUNSAT}_\psi(G) \stackrel{\mathrm{def}}{=} \frac{1}{|E|}\langle\psi|H|\psi\rangle = \frac{1}{|E|}\sum_{e\in E}\langle\psi|Q_e|\psi\rangle \ . \tag{19}$$

To define a new – "amplified" – constraint system, we consider all possible t-walks (t is fixed) $\mathbf{e} = (e_1, \ldots, e_t)$ and for each such walk, we define a $t\log(q)$-local projection $Q_\mathbf{e}$ by taking the intersection of all the accepting spaces along the path and defining it to be the accepting space of $Q_\mathbf{e}$. In other words, $Q_\mathbf{e}$ projects into the orthogonal complement of that space. We refer to the new system as G^t, and define

$$\mathrm{QUNSAT}_\psi(G^t) \stackrel{\mathrm{def}}{=} \frac{\sum_\mathbf{e}\langle\psi|Q_\mathbf{e}|\psi\rangle}{\#\text{ of }t\text{-walks}} \ , \quad \mathrm{QUNSAT}(G^t) \stackrel{\mathrm{def}}{=} \min_\psi \mathrm{QUNSAT}_\psi(G^t) \ . \tag{20}$$

As in the classical case, the quantum amplification lemma shows how QUNSAT(G^t) is amplified with respect to QUNSAT(G). The amplification is linear in t when QUNSAT(G) is far enough from 1, and then becomes saturated, just like in the classical case.

Lemma 6.1 (The quantum amplification lemma) *Consider a k-QSAT system on an expander graph $G = (V, E)$ as described above. Then there are positive functions $K(q, d, \theta)$, $c(\lambda)$, independent of the system size, such that*

$$\text{QUNSAT}(G^t) \geq c(\lambda) \cdot K(q, d, \theta) \cdot \min\left\{t \cdot \text{QUNSAT}(G), 1\right\} , \tag{21}$$

Proof: By definition, QUNSAT$(G) = \epsilon_0/|E|$ where ϵ_0 is the ground energy of G. Let $|\psi\rangle$ be a state for which QUNSAT$(G^t) = \text{QUNSAT}_\psi(G^t)$. We first notice that our k-QSAT system can be written with at most $g = 2d$ layers. We choose a layer i and expand $|\psi\rangle$ in terms of its violations in that layer:

$$|\psi\rangle = \sum_{j=0}^{|E|} \alpha_j |\psi_j\rangle \ . \tag{22}$$

Here $|\psi_j\rangle$ is the projection of $|\psi\rangle$ to the space with j violations in the i'th layer. Note that $|\psi_j\rangle$ has a well defined value of the number of violated constraints of the i'th level. We consider an auxiliary k-QSAT system G_i which has same underlying graph G and the same constraints of the i'th layer - but the rest of the constraints are null - i.e. they are always satisfied. It is clear that for every state $|\psi\rangle$, $\text{QUNSAT}_\psi(G^t) \geq \text{QUNSAT}_\psi(G_i^t)$. Moreover, all the projections in G_i^t commute within themselves and with the original projections of the i'th layer, therefore

$$\text{QUNSAT}_\psi(G_i^t) = \sum_j \alpha_j^2 \cdot \text{QUNSAT}_{\psi_j}(G_i^t) \ . \tag{23}$$

We will now show:

Claim 6.2

$$\text{QUNSAT}_{\psi_j}(G_i^t) \geq \begin{cases} t \cdot c(\lambda) \cdot \frac{j}{|E|} & , \ \textit{for } j \leq \frac{|E|}{t} \\ c(\lambda) & , \ \textit{for } j > \frac{|E|}{t} \end{cases} . \tag{24}$$

This claim follows almost immediately from the classical amplification lemma since the constraints in G_i are all commuting and can be considered classical. The full proof is given in Appendix I.

Let us now use it to estimate the amplification. Combining Eq. (24) with Eq. (23), we find

$$\text{QUNSAT}(G^t) = \text{QUNSAT}_\psi(G^t) \geq \text{QUNSAT}_\psi(G_i^t) \tag{25}$$

$$\geq t\frac{c(\lambda)}{|E|}\left(\alpha_1^2 + 2\alpha_2^2 + 3\alpha_3^2 + \ldots + \frac{|E|}{t}\alpha_{|E|/t}^2\right) + c(\lambda)\left(\alpha_{|E|/t+1}^2 + \ldots + \alpha_{|E|}^2\right) \ . \tag{26}$$

Therefore, as QUNSAT$(G) = \frac{\epsilon_0}{|E|}$, the amplification ratio we are looking for is

$$\frac{\text{QUNSAT}(G^t)}{\text{QUNSAT}(G)} \geq t\frac{c(\lambda)}{\epsilon_0}\left(\alpha_1^2 + 2\alpha_2^2 + 3\alpha_3^2 + \ldots + \frac{|E|}{t}\alpha_{|E|/t}^2\right) + c(\lambda) \cdot \frac{|E|}{\epsilon_0}\left(\alpha_{|E|/t+1}^2 + \ldots + \alpha_{|E|}^2\right) \tag{27}$$

The above equation is central and can be derived for *any* layer (namely, for any i). However, without additional information, it cannot be used to show amplification of QUNSAT(G^t). The reason is that the weights α_j^2 can theoretically conspire in such a way that no amplification would occur. For example, $1/poly(|E|)$ of the weight can be concentrated on $\alpha_{|E|}^2$ and the rest on α_0^2, and then there is no amplification since in these two sectors there is no amplification (one is completely satisfied and the other is completely saturated). Fortunately, we can use the detectability lemma to rule out the possibility that this sort of non-amplifying distribution appears *simultaneously* in all layers.

Specifically, we consider two cases: $\frac{\epsilon_0}{f} - \frac{4}{1-\theta} \leq 2r$ (the low-energy case) and $\frac{\epsilon_0}{f} - \frac{4}{1-\theta} > 2r$ (the high-energy case). In the former, we use the $\ell = 0$ detectability lemma, and in the latter we the $\ell > 0$ detectability lemma. The analysis of these two cases is quite straightforward, and is given in Appendix J. It shows that in the low-energy case, the amplification ratio in Eq. (27) must be bigger than $t \cdot c(\lambda) \cdot K_1(q, d, \theta)$, and in the high-energy case it is either larger than $t \cdot c(\lambda) \cdot K_2(q, d, \theta)$ or $\mathrm{QUNSAT}(G^t)$ is already larger than $c(\lambda) \cdot K_3(q, d, \theta)$. Here, $K_i(q, d, \theta), i = 1, 2, 3$ are some positive functions. Combining all these results prove the amplification lemma.

■

7 Discussion regarding quantum PCP

The classical PCP theorem (see Ref. [25]) states that the classical complexity class $\mathsf{PCP}[r, q]$ (with r being the number of random bits and q the number of queries) satisfies $\mathsf{NP} = \mathsf{PCP}[\log(n), 1]$. A natural generalization of the PCP complexity class to the quantum setting is:

Definition 7.1 (The class $\mathsf{QPCP}[q]$) *A language L is in $\mathsf{QPCP}[q]$ if there exists a BQP verifier V and a polynomial $p(\cdot)$ with the following properties. V receives as input a classical string x and a state $|\xi\rangle \in \mathbb{B}^{\otimes p(|x|)}$. V However, has access to only $\mathcal{O}(q)$ randomly chosen qubits from $|\psi\rangle$. We denote V's action on $(x, |\xi\rangle)$ by $V(x, |\xi\rangle)$. Then the condition for L to be in $\mathsf{QPCP}[q]$ is that*

- *If $x \in L$, there exists a witness $|\xi\rangle \in \mathbb{B}^{\otimes p(|x|)}$ such that $\Pr[V(x, |\xi\rangle)$ accepts$] \geq 2/3$.*
- *If $x \notin L$, then for every $|\xi\rangle \in \mathbb{B}^{\otimes p(|x|)}$ we have $\Pr[V(x, |\xi\rangle)$ accepts$] \leq 1/3$.*

Notice that we did not give the quantum verifier any random bits, since it is quantum and can generate randomness by itself. One naturally speculates:

Conjecture 7.2 (Quantum PCP) $\mathsf{QPCP}[1] = \mathsf{QMA}$.

An essentially equivalent way of formulating the quantum PCP theorem is in terms of local Hamiltonians: is it possible to efficiently transform any k-QSAT system with 1/poly promise gap into a k-QSAT system with constant promise gap. This corresponds to the inapproximability of max-3SAT.

Recently, Dinur gave a beautiful new proof of the classical PCP theorem [4], which works directly in this setting. She starts with a classical SAT system with 1/poly promise gap and successively amplifies the gap by repeated doubling. This doubling is accomplished by gap amplification followed by alphabet reduction and degree reduction to control the size and locality. Our quantum gap amplification lemma can be seen as a step towards emulating this approach in the quantum setting.

8 Acknowledgments

We are grateful to Matt Hastings and Tobias Osborne for exciting and inspiring discussions about the quantum gap amplification lemma, the possibility of a quantum PCP theorem, and possible avenues to prove (or disprove) it. We are also grateful to Michael Ben-Or, Avinatan Hassidim, and Barbara Terhal for useful discussions and comments.

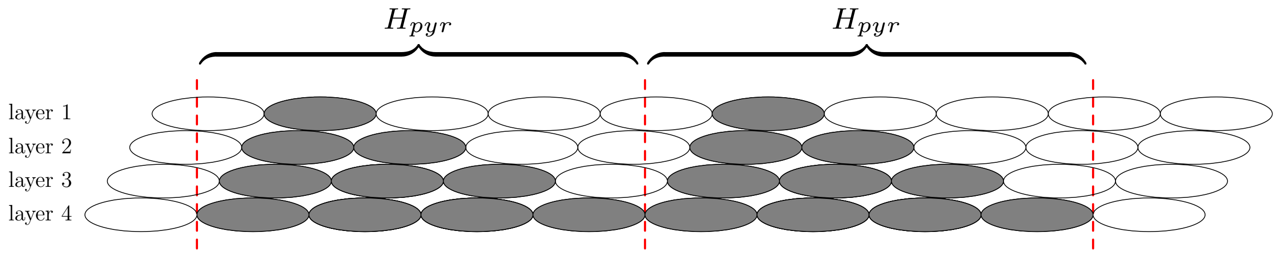

Figure 1: Illustration of the pyramids that define the XY decomposition.

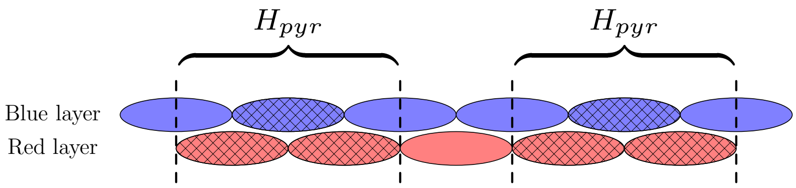

Figure 2: The setting of the XY decomposition in the case of two layers (colors).

References

[1] A. Kitaev, A. Shen, and M. Vyalyi, *Classical and Quantum Computation.* American Mathematical Society, 2002.

[2] M. Ajtai, J. Komlos, and E. Szemeredi, “Deterministic simulation in LOGSPACE,” in *Proceedings of the nineteenth annual ACM conference on Theory of computing*, pp. 132–140, ACM New York, NY, USA, 1987.

[3] R. Impagliazzo and D. Zuckerman, “How to recycle random bits,” in *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pp. 248–253.

[4] I. Dinur, “The pcp theorem by gap amplification,” *J. ACM*, vol. 54, no. 3, p. 12, 2007.

[5] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, “Quantum Computation by Adiabatic Evolution,” *Arxiv preprint quant-ph/0001106*, 2000.

[6] W. van Dam, M. Mosca, and U. Vazirani, “How Powerful is Adiabatic Quantum Computation?,” in *ANNUAL SYMPOSIUM ON FOUNDATIONS OF COMPUTER SCIENCE*, vol. 42, pp. 279–287, IEEE COMPUTER SOCIETY PRESS, 2001.

[7] W. van Dam and U. Vazirani, “More on the power of adiabatic computation,” *unpublished*, 2001.

[8] B. W. Reichardt, “The quantum adiabatic optimization algorithm and local minima,” in *STOC ’04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, (New York, NY, USA), pp. 502–510, ACM, 2004.

[9] J. Roland and N. Cerf, “Quantum search by local adiabatic evolution,” *Physical Review A*, vol. 65, no. 4, p. 42308, 2002.

[10] E. Farhi, J. Goldstone, and S. Gutmann, “Quantum Adiabatic Evolution Algorithms versus Simulated Annealing,” *Arxiv preprint quant-ph/0201031*, 2002.

[11] D. Aharonov and A. Ta-Shma, “Adiabatic quantum state generation and statistical zero knowledge,” in *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pp. 20–29, ACM Press New York, NY, USA, 2003.

[12] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, “Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation,” in *ANNUAL SYMPOSIUM ON FOUNDATIONS OF COMPUTER SCIENCE*, vol. 45, pp. 42–53, IEEE COMPUTER SOCIETY PRESS, 2004.

[13] A. Childs, E. Farhi, and J. Preskill, “Robustness of adiabatic quantum computation,” *Physical Review A*, vol. 65, no. 1, p. 12322, 2001.

[14] J. Ĺberg, D. Kult, and E. Sjöqvist, “Robustness of the adiabatic quantum search,” *Physical Review A*, vol. 71, no. 6, p. 60312, 2005.

[15] J. Roland and N. Cerf, “Noise resistance of adiabatic quantum computation using random matrix theory,” *Physical Review A*, vol. 71, no. 3, p. 32330, 2005.

[16] M. Sarandy and D. Lidar, “Adiabatic Quantum Computation in Open Systems,” *Physical Review Letters*, vol. 95, no. 25, p. 250503, 2005.

[17] D. A. Lidar, “Towards fault tolerant adiabatic quantum computation,” *Physical Review Letters*, vol. 100, no. 16, p. 160506, 2008.

[18] S. Jordan, E. Farhi, and P. Shor, “Error-correcting codes for adiabatic quantum computation,” *Physical Review A*, vol. 74, no. 5, p. 52322, 2006.

[19] D. Aharonov and M. Ben-Or, “Fault-tolerant quantum computation with constant error,” in *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pp. 176–188, ACM New York, NY, USA, 1997.

[20] E. Knill, R. Laflamme, and W. Zurek, “Resilient Quantum Computation,” *Science*, vol. 279, no. 5349, p. 342, 1998.

[21] A. Kitaev, “Fault-tolerant quantum computation by anyons,” *Annals of Physics*, vol. 303, no. 1, pp. 2–30, 2003.

[22] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, “The power of quantum systems on a line,” *arXiv preprint*, 2007.

[23] M. Hastings, “An Area Law for One Dimensional Quantum Systems,” *JSTAT, P*, vol. 8024, 2007.

[24] T. Osborne, “Simulating adiabatic evolution of gapped spin systems,” *Physical Review A*, vol. 75, no. 3, p. 32321, 2007.

[25] S. Arora and B. Barak, “Computational Complexity: A Modern Approach,” *to appear: http://www. cs. princeton.edu/theory/complexity*.

[26] J. Kempe, A. Kitaev, and O. Regev, “The Complexity of the Local Hamiltonian Problem,” *SIAM JOURNAL ON COMPUTING*, vol. 35, no. 5, p. 1070, 2006.

[27] R. Oliveira and B. Terhal, “The complexity of quantum spin systems on a two-dimensional square lattice,” *Quant. Inf. Comp.*, vol. 8, no. 10, pp. 0900–0924, 2008.

[28] S. Bravyi, D. DiVincenzo, D. Loss, and B. Terhal, “Quantum Simulation of Many-Body Hamiltonians Using Perturbation Theory with Bounded-Strength Interactions,” *Physical Review Letters*, vol. 101, no. 7, p. 70503, 2008.

[29] S. P. Jordan and E. Farhi, “Perturbative gadgets at arbitrary orders,” *Physical Review A (Atomic, Molecular, and Optical Physics)*, vol. 77, no. 6, p. 062329, 2008.

A Proof that Lemma 4.1 follows from Lemma 4.2

Proof of Lemma 4.1: To see how Lemma 4.1 follows from Lemma 4.2, assume that the latter holds, and assume that both $\|(1-\Pi_{red})|\psi\rangle\|^2 < \frac{1}{8}\Delta^2(0)$ and $\|(1-\Pi_{red})|\psi\rangle\|^2 < \frac{1}{8}\Delta^2(0)$. Then we write

$$\Pi_{blue}\Pi_{red}\Pi_{blue} = \mathbb{1} + (\Pi_{blue} - \mathbb{1}) + \Pi_{blue}(\Pi_{red} - \mathbb{1}) + \Pi_{blue}\Pi_{red}(\Pi_{blue} - \mathbb{1}) \ , \tag{28}$$

This way, except for the identity, every term on the RHS has a $\Pi_{blue} - \mathbb{1}$ or $\Pi_{red} - \mathbb{1}$ on its right side. We want also the left side to have such a term, so we continue in the same fashion, and write:

$$\Pi_{blue}(\Pi_{red} - \mathbb{1}) = (\Pi_{blue} - \mathbb{1})(\Pi_{red} - \mathbb{1}) + \Pi_{red} - \mathbb{1} \ , \tag{29}$$

$$\Pi_{blue}\Pi_{red}(\Pi_{blue} - \mathbb{1}) = (\Pi_{blue} - \mathbb{1})\Pi_{red}(\Pi_{blue} - \mathbb{1}) + (\Pi_{red} - \mathbb{1})(\Pi_{blue} - \mathbb{1}) + (\Pi_{blue} - \mathbb{1}) \ . \tag{30}$$

All together we have

$$\Pi_{blue}\Pi_{red}\Pi_{blue} = \mathbb{1} + \big[6 \text{ terms with } (\mathbb{1} - \Pi_{red}) \text{ or } (\mathbb{1} - \Pi_{red}) \text{ on both sides.}\big] \tag{31}$$

When we "sandwich" the above equation with $\langle\psi| \cdot |\psi\rangle$, the absolute value of each of the 6 terms will be smaller then $\Delta^2(0)/8$. This is due to the Cauchy-Schwartz inequality and the assumption that $\|(\mathbb{1} - \Pi_{red})|\psi\rangle\|$, $\|(\mathbb{1} - \Pi_{blue})|\psi\rangle\| < \sqrt{\Delta^2(0)/8}$. Therefore,

hence,

$$\|\Pi_{red}\Pi_{blue}|\psi\rangle\|^2 = \langle\psi|\Pi_{blue}\Pi_{red}\Pi_{blue}|\psi\rangle > 1 - \frac{6}{8}\Delta^2(0) > 1 - \Delta^2(0) \ , \tag{32}$$

which is a contradiction.

■

B Proof of bound on η_s

Here we prove Claim 4.4. Summing over Equation 6 squared, we get

$$1 \leq \frac{1}{x^2}\sum_{s=0}^{m} \theta^{2s}\eta_s^2 \ , \tag{33}$$

which is equivalent to

$$x^2 \leq \eta_0^2 + \sum_{s=1}^{m} \theta^{2s}\eta_s^2 \leq \eta_0^2 + \theta^2 \sum_{s=1}^{m} \eta_s^2 \ . \tag{34}$$

But $\sum_{s=0}^{m} \eta_s^2 \leq 1$, so $\eta_0^2 \leq 1 - \sum_{s=1}^{m} \eta_s^2$, and

$$x^2 \leq 1 - \sum_{s=1}^{m} \eta_s^2 + \theta^2 \sum_{s=1}^{m} \eta_s^2 = 1 - (1-\theta^2)\sum_{s=1}^{m} \eta_s^2 \ , \tag{35}$$

which leads to

$$\sum_{s=1}^{m} \eta_s^2 \leq \frac{1-x^2}{1-\theta^2} \ , \tag{36}$$

implying the desired bound.

C Proof of Claim 4.3

Proof: We will prove this claim on the fine-grained XY decomposition, by showing that $\langle\Omega_\nu|E^{top}|\Omega_\nu\rangle \le |\nu|$.

Essentially, the claim follows from the fact that only the Y sites can contribute energy. Indeed, consider an X pyramid, and let Q be its blue constraint. Then by definition, either $Q|\Omega_\nu\rangle = 0$ or $Q|\Omega_\nu\rangle = |\Omega_\nu\rangle$. If the site contributes non-zero energy, the latter must hold. But $|\Omega_\nu\rangle \propto \Pi_\nu|\Omega\rangle \propto P_\nu\Pi_{red}\Pi_{blue}|\psi\rangle$, and so we get

$$QP_\nu\Pi_{red}\Pi_{blue}|\psi\rangle = \Pi_\nu\Pi_{red}\Pi_{blue}|\psi\rangle \ . \tag{37}$$

We we show that the RHS of the above equation must vanish. Indeed, by Eq. (4), the LHS of the equation can be written as

$$QP_\nu\Delta_{red}\Delta_{blue}R_{red}R_{blue}|\psi\rangle \ . \tag{38}$$

But as the pyramids' projections commute with P_ν, we get

$$QP_\nu\Delta_{red}\Delta_{blue} = P_\nu Q\Delta_{red}\Delta_{blue}P_\nu \ , \tag{39}$$

and because in the X subspaces the blue and red constraints commute, this is equal to

$$P_\nu\Delta_{red}Q\Delta_{blue}P_\nu \ . \tag{40}$$

This expression must vanish since Δ_{blue} contains a $\mathbb{1} - Q$ term. It follows that the RHS of Eq. (37) must vanish and this proves the claim. ∎

D Relation of the simple detectability lemma and Kitaev's lemma

The $\ell = 0$ detectability lemma for the two layers can be seen as the converse of a special case of Kitaev's geometrical lemma, crucial in his proof of the quantum Cook-Levin theorem [1].

Lemma D.1 (Kitaev's lemma (see Ref. [1])) *Given finite-dimensional operators $P \ge 0, Q \ge 0$ with null eigenspaces, then*

$$P + Q \ge \min\Big\{\Delta(P), \Delta(Q)\Big\} \cdot (1 - \cos\alpha) \ , \tag{41}$$

where $\Delta(O) > 0$ is the smallest nonzero eigenvalue of O, and α the angle between the null spaces of P and Q.

Therefore if $\Delta(P), \Delta(Q)$ are fixed, by lower-bounding α, we can lower-bound the minimal energy of $P + Q$.

The 2-layers, $\ell = 0$ detectability lemma can be seen as the converse of this statement that holds in special case. In such case let the Q, P operators be the Π_{red} and Π_{blue} projections from Sec. 4. Then the detectability lemma can be used to lower-bound α. Indeed, for every state $|\psi\rangle$, Lemma 4.2 asserts that

$$\langle\psi|\Pi_{blue}\Pi_{red}\Pi_{blue}|\psi\rangle \le 1 - \Delta^2(0) \ . \tag{42}$$

Then the angle α is given by

$$\cos\alpha = \min_{\substack{|\psi\rangle\in H_P \\ \|\psi\|=1}} \langle\psi|Q|\psi\rangle = \min_{\|\psi\|=1} \langle\psi|\Pi_{blue}\Pi_{red}\Pi_{blue}|\psi\rangle \le 1 - \Delta^2(0) \ , \tag{43}$$

where H_P is the null space of P. Therefore, $1-\cos\alpha \ge \Delta^2(\ell=0,\epsilon_0)$, and combining it with Eq. (41), we get

$$\Delta^2(\ell=0,\epsilon_0) \le 1-\cos\alpha \le \epsilon_0 \ . \tag{44}$$

Moreover, looking at Eq. (18), we see that in the limit $\epsilon_0 \to 0$,

$$\left(\frac{1-\theta^2}{\theta}\right)^2 (1-\theta)\frac{\epsilon_0}{f} \le 1-\cos\alpha \le \epsilon_0 \ . \tag{45}$$

E Proving the exponential decay in the general case

In this section we prove the exponential decay in the general case, which is stated in Lemma 3.3 in Sec. 3.6. The proof follows essentially the path of the 2-layers, $\ell=0$ case. We start by proving the decay in the fine-grained XY decomposition. Consider then a given XY decomposition and some sector ν with $|\nu| \ge \ell$.

Claim E.1 *There exist $(\ell+1)^g$ states $|\Phi_j\rangle, j=1,\ldots,(\ell+1)^g$ with $\|\Phi_j\| \le 1$, such that the weight of every XY sector ν with $|\nu| \ge \ell$, is bounded by*

$$\lambda_\nu^2 \le \frac{1}{x^2}\frac{(\ell+1)^g}{(\ell!)^{2g}}\left(|\nu|^{g\ell}k^{g^2\ell}\theta^{|\nu|}\right)^2 \sum_{j=1}^{(\ell+1)^g} \|P_\nu|\Phi_j\rangle\|^2 \ . \tag{46}$$

Proof: By definition,

$$\lambda_\nu^2 = \langle\Omega|P_\nu|\Omega\rangle \tag{47}$$

$$= \frac{1}{x^2}\langle\psi|\Pi^{(1)}_{\le\ell}\cdots\Pi^{(g)}_{\le\ell}\cdot P_\nu\cdot\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle \tag{48}$$

$$= \frac{1}{x^2}\|P_\nu\cdot\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle\|^2 \ . \tag{49}$$

Let us estimate $\|P_\nu\cdot\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle\|$. Similarly to what was done in Sec. 3.4, we write the $\Pi^{(i)}_{\le\ell}$ projections in terms of projections inside and outside the pyramids. This time, however, there are $\ell+1$ possible terms:

$$\Pi^{(i)}_{\le\ell} = \sum_{j=0}^{\ell}\Delta^{(i)}_j\cdot R^{(i)}_{\le\ell-j} \ . \tag{50}$$

Here $\Delta^{(i)}_j$ denote the projection into the subspace in which all the constraints in the i'th layer *inside* the pyramid have exactly j violations, and $R^{(i)}_{\le j}$ denotes the projection into the subspace in which the constraints of the i'th layer *outside* the pyramids have j or less violations. Then As in the 2-layers case, because of the structure of the pyramids in which the support of every layer is included of the support of the layer we can "pull back" the $\Delta^{(i')}_{j'}$ operators across the $R^{(i)}_{\le j}$ operators as long as $i' < i$, and so

$$\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell} = \sum_{j_1,\ldots,j_g}(\Delta^{(g)}_{j_g}\cdots\Delta^{(1)}_{j_1})\cdot(R^{(g)}_{\le\ell-j_g}\cdots R^{(1)}_{\le\ell-j_1}) \ . \tag{51}$$

Combining the above equation with Eq. (3), we find

$$\|P_\nu\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle\| \le \sum_{j_1,\ldots,j_g}\|P_\nu(\Delta^{(g)}_{j_g}\cdots\Delta^{(1)}_{j_1})P_\nu\|\cdot\|P_\nu(R^{(g)}_{\le\ell-j_g}\cdots R^{(1)}_{\le\ell-j_1})|\psi\rangle\| \ . \tag{52}$$

We will upper-bound $\|P_\nu(\Delta^{(g)}_{j_g}\cdots\Delta^{(1)}_{j_g})P_\nu\|$. Every projection $\Delta^{(i)}_j$ can be written as a sum of products of the form $Q\cdot Q\cdot(\mathbb{1}-Q)\cdot\ldots$ that work on the projections of the i'th layer that are inside the pyramid, such that there are exactly j projections of the form Q and the rest is of the form $\mathbb{1}-Q$ - corresponding to exactly j violations.

The product $\Delta^{(g)}_{j_g}\cdots\Delta^{(1)}_{j_g}$, therefore, contains a huge number of such products. However, when we "sandwich" it between two P_ν projections, only few survive - those that are compatible with the X portion of P_ν. Let us estimate how many survive in a given layer. The X part is completely fixed, and therefore we have to choose from all the Y projections at most ℓ violations. There are $|\nu|$ Y sites and at each site there are at most k^g constraints, so overall, for $\ell\le|\nu|$, the number of surviving constraints in a single layer is bounded by

$$\binom{|\nu|k^g}{\ell}\le\frac{1}{\ell!}(|\nu|k^g)^\ell\ . \tag{53}$$

Considering all g layers, the total number of surviving terms is therefore bounded by $\left(\frac{1}{\ell!}|\nu|^\ell k^{g\ell}\right)^g$. The norm of each term is bounded by $\theta^{|\nu|}$ as there are $|\nu|$ Y sites. Therefore, the overall norm is bounded by

$$\|P_\nu(\Delta^{(g)}_{j_g}\cdots\Delta^{(1)}_{j_1})P_\nu\|\le\frac{1}{(\ell!)^g}|\nu|^{g\ell}k^{g^2\ell}\theta^{|\nu|}\ . \tag{54}$$

Thus far, we got

$$\|P_\nu\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle\|\le\frac{1}{(\ell!)^g}|\nu|^{g\ell}k^{g^2\ell}\theta^{|\nu|}\sum_{j_1,\ldots,j_g}\|P_\nu(R^{(g)}_{\le\ell-j_g}\cdots R^{(1)}_{\le\ell-j_g})|\psi\rangle\|\ . \tag{55}$$

There are $(\ell+1)^g$ terms in that sum, and so using standard Cauchy-Schwartz argument we get

$$\|P_\nu\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle\|^2\le\frac{(\ell+1)^g}{(\ell!)^{2g}}\left(|\nu|^{g\ell}k^{g^2\ell}\theta^{|\nu|}\right)^2\sum_{j_1,\ldots,j_g}\|P_\nu(R^{(g)}_{\le\ell-j_g}\cdots R^{(1)}_{\le\ell-j_g})|\psi\rangle\|^2\ . \tag{56}$$

Finally, grouping all the indices $(j_1,\ldots,j_g)$ into one big index j, and defining the un-normalized states

$$|\Phi_j\rangle\stackrel{\text{def}}{=}R^{(g)}_{\le\ell-j_g}\cdots R^{(1)}_{\le\ell-j_1}|\psi\rangle\ , \tag{57}$$

whose norm is smaller than or equal to 1, we get that for $|\nu|\ge\ell$,

$$\lambda^2_\nu=\langle\Omega|P_\nu|\Omega\rangle\le\frac{1}{x^2}\frac{(\ell+1)^g}{(\ell!)^{2g}}\left(|\nu|^{g\ell}k^{g^2\ell}\theta^{|\nu|}\right)^2\sum_{j=1}^{(\ell+1)^g}\|P_\nu|\Phi_j\rangle\|^2\ . \tag{58}$$

■

To prove Lemma 3.3, pass to the coarse grained XY decomposition by grouping together all the XY sectors with the same number of Y's. Then

$$\lambda^2_s=\sum_{|\nu|=s}\lambda^2_\nu\le\frac{1}{x^2}\frac{(\ell+1)^g}{(\ell!)^{2g}}\left(s^{g\ell}k^{g^2\ell}\theta^s\right)^2\sum_{j=1}^{(\ell+1)^g}\sum_{|\nu|=s}\|P_\nu|\Phi_j\rangle\|^2\ . \tag{59}$$

Defining

$$\eta^2_s\stackrel{\text{def}}{=}\frac{1}{(\ell+1)^g}\sum_{j=1}^{(\ell+1)^g}\sum_{|\nu|=s}\|P_\nu|\Phi_j\rangle\|^2=\frac{1}{(\ell+1)^g}\sum_{j=1}^{(\ell+1)^g}\|P_s|\Phi_j\rangle\|^2\ , \tag{60}$$

we find that $\sum_s\eta^2_s\le1$ (recall that $\|\Phi_j\|\le1$) and by Eq. (59), for every $s\ge\ell$,

$$\lambda_s\le\frac{1}{x}k^{g^2\ell}\left(\frac{\ell+1}{\ell!}\right)^g s^{g\ell}\theta^s\eta_s\ . \tag{61}$$

F Proving general detectability lemma, Lemma 5.1

To prove this lemma, we will prove the following auxiliary lemma

Lemma F.1 *Let $\Pi^{(i)}_{\le\ell} = \mathbb{1} - \Pi^{(i)}_{>\ell}$ denote the projection into ℓ or less violations in the i'th layer as in Sec. 3.6. Then*

$$\|\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle\|^2 \le 1 - \Delta^2(\ell) \ , \tag{62}$$

where $\Delta(\ell)$ is defined in Lemma 5.1, and in the $\ell > 0$ case we assume that $(\epsilon_0/f) - r\ell > \frac{1}{1-\theta}$.

The proof of Lemma F.1 would be given later in Sec. F.2. Based on it, we can prove Lemma 5.1 as follows

Proof of Lemma 5.1: Given the state $|\psi\rangle$ and an integer $\ell \ge 0$, assume that Eq. (62) holds and yet for every layer,

$$\|\Pi^{(i)}_{>\ell}|\psi\rangle\|^2 < \frac{1}{(2g)^2}\Delta^2(\ell) \ . \tag{63}$$

For brevity, we denote

$$x \stackrel{\text{def}}{=} \|\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle\| \ . \tag{64}$$

Then

$$x^2 = \langle\psi|\Pi^{(1)}_{\le\ell}\cdots\Pi^{(g-1)}_{\le\ell}\Pi^{(g)}_{\le\ell}\Pi^{(g-1)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle \ . \tag{65}$$

Every product of N operators can be written as:

$$O_1\cdots O_N = \mathbb{1} + (O_1 - \mathbb{1}) + O_1(O_2 - \mathbb{1}) + O_1O_2(O_3 - \mathbb{1}) \tag{66}$$

$$+ \ldots + (O_1\cdots O_{N-1})\cdot(O_N - \mathbb{1}) \ . \tag{67}$$

Expanding Eq. (65) this way, we get

$$x^2 = 1 + \langle\psi|(\Pi^{(1)}_{\le\ell} - \mathbb{1})|\psi\rangle + \langle\psi|\Pi^{(1)}_{\le\ell}\big(\Pi^{(2)}_{\le\ell} - \mathbb{1}\big)|\psi\rangle + \langle\psi|\Pi^{(1)}_{\le\ell}\Pi^{(2)}_{\le\ell}\big(\Pi^{(3)}_{\le\ell} - \mathbb{1}\big)|\psi\rangle + \ldots \tag{68}$$

The RHS of the above equation contains $2g-1$ terms of the form $\langle\psi|\Pi^{(1)}_{\le\ell}\cdots\Pi^{(i)}_{\le\ell}\big(\Pi^{(i+1)}_{\le\ell} - \mathbb{1}\big)|\psi\rangle$. Let us estimate their magnitude. By an expansion similar to Eq. (66), we write

$$\Pi^{(1)}_{\le\ell}\cdots\Pi^{(i)}_{\le\ell} = (\mathbb{1} - \Pi^{(i)}_{\le\ell}) + (\mathbb{1} - \Pi^{(i-1)}_{\le\ell})\Pi^{(i)}_{\le\ell} + \ldots \ . \tag{69}$$

Therefore $\langle\psi|\Pi^{(1)}_{\le\ell}\cdots\Pi^{(i)}_{\le\ell}\big(\Pi^{(i+1)}_{\le\ell} - \mathbb{1}\big)|\psi\rangle$ can be written as a sum of at most $2g$ terms, each of them is an inner product of $\langle\psi|(\mathbb{1} - \Pi^{(j)}_{\le\ell})$ times some projections, times $(\mathbb{1} - \Pi^{(i)}_{\le\ell})|\psi\rangle$. By our assumption, the norm of the ket and bra is smaller than $\Delta(\ell)/(2g)$ and as the norms of the projections are smaller than or equal to unity we find

$$|\langle\psi|\Pi^{(1)}_{\le\ell}\cdots\Pi^{(i)}_{\le\ell}\big(\Pi^{(i+1)}_{\le\ell} - \mathbb{1}\big)|\psi\rangle| \le 2g\frac{\Delta^2(\ell)}{(2g)^2} = \frac{\Delta^2(\ell)}{2g} \ . \tag{70}$$

Therefore, overall,

$$x^2 \le 1 + (2g-1)\frac{\Delta^2(\ell)}{2g} < 1 + \Delta^2(\ell) \ , \tag{71}$$

contradicting Eq. (62). ■

We now turn to the proof of Lemma F.1. The outline of the proof is very similar to the simple case of 2-layers, $\ell = 0$, and was discussed in Sec. 5. The main goal of the proof is to estimate the energy of the normalized state $\frac{1}{x}\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle$, which The has contributions from all layers. For every layer we will find a crude upper bound of its energy as a function of x (as well as of ℓ, k, g, θ). Summing all these bounds together, we will get an upper bound to the total energy. This energy is lower bounded by ϵ_0, the ground energy of the system, and this gives us an inequality. We then reverse it and extract an upper bound for x.

We start by using the XY decomposition to upper bound the energy of the first layer.

F.1 Estimating the energy of the first layer

Consider then an XY decomposition, and let E^{top} denote the energy of all the constraints of the first layer (the top layer in Fig. 1) that belong to the pyramids of the decomposition. We define $|\Omega\rangle$ to be the following normalized state:

$$|\Omega\rangle \stackrel{\text{def}}{=} \frac{1}{x}\Pi^{(g)}_{\le\ell}\cdots\Pi^{(1)}_{\le\ell}|\psi\rangle \ . \tag{72}$$

The entire section will be dedicated to proving the following lemma:

Lemma F.2 *For $\ell = 0$,*

$$\langle\Omega|E^{top}|\Omega\rangle \le \frac{1-x^2}{x^2}\frac{\theta^2}{(1-\theta^2)^3} \ , \tag{73}$$

and for $\ell > 0$ there is a positive function function $r(\theta, k, g)$ (independent of $|\Omega\rangle$) such that

$$\langle\Omega|E^{top}|\Omega\rangle \le r\ell + \frac{1}{x^2}\frac{1}{1-\theta} \ . \tag{74}$$

Proof:

The $\ell = 0$ case was essentially already proved in the 2-layers case in Sec. 4 (specifically, see Eq. (11)). The difference between the 2-layers case and the g-layers case are semantic and therefore we will only consider the $\ell > 0$ case.

Consider the coarse- and fine-grained XY decomposition of $|\Omega\rangle$,

$$|\Omega\rangle = \sum_\nu \lambda_\nu|\Omega_\nu\rangle = \sum_s \lambda_s|\Omega_\nu\rangle \ . \tag{75}$$

Since E^{top} a sum of the inverses of pyramid projections from the first layer, it must commute with the XY projections P_ν. Therefore,

$$\langle\Omega|E^{top}|\Omega\rangle = \sum_s \lambda_s^2\langle\Omega_s|E^{top}|\Omega_s\rangle \ . \tag{76}$$

Our first claim is

Claim F.3 *For every s with non-zero weight λ_s,*

$$\langle\Omega_s|E^{top}|\Omega_s\rangle \le \ell + s \ . \tag{77}$$

Proof: It is sufficient to prove that for every sector ν with non-zero weight, $\langle\Omega_\nu|E^{top}|\Omega_\nu\rangle \le \ell + |\nu|$.

E^{top} has one contribution from every pyramid top. Consider a sector ν of the fine-grained XY decomposition. It contains $|\nu|$ Y spaces and the rest are X spaces. The maximal energy contribution from the Y pyramids is therefore $|\nu|$. We will now show that the contribution from the X pyramids is at most ℓ. Essentially, the proof boils down to the fact that the projections commute on the X sectors and therefore if $|\Omega\rangle$ has more than ℓ violations on the X sectors then also $\Pi^{(1)}_{\leq \ell}|\psi\rangle$ has - which is impossible. The following argument shows this more formally.

Let Q_i be the projection in the first layer in the i'th pyramid, where ν has an X sector. Then either $Q_i|\Omega_\nu\rangle = 0$ or $Q_i|\Omega_\nu\rangle = |\Omega_\nu\rangle$. If the total contribution from all X sectors is larger than ℓ, there are $\ell+1$ pyramids in which $Q_i|\Omega_\nu\rangle = |\Omega_\nu\rangle$. For brevity, assume that these appear in the first $\ell+1$ pyramids. Then

$$(\prod_{i=1}^{\ell+1} Q_i)|\Omega_\nu\rangle = |\Omega_\nu\rangle \ . \tag{78}$$

Assuming that $\lambda_\nu \neq 0$, we get

$$(\prod_{i=1}^{\ell+1} Q_i)P_\nu|\Omega\rangle = P_\nu|\Omega\rangle \ . \tag{79}$$

Using the definition of $|\Omega\rangle$ in Eq. (72) and Eq. (51), the LHS of the above equation is equal to

$$\frac{1}{x}\sum_{j_1,\ldots,j_g}(\prod_{i=1}^{\ell+1} Q_i)P_\nu \cdot (\Delta^{(g)}_{j_g}\cdots\Delta^{(1)}_{j_g}) \cdot (R^{(g)}_{\leq \ell - j_g}\cdots R^{(1)}_{\leq \ell - j_g})|\psi\rangle \ . \tag{80}$$

The above expression vanishes. The reason is that it is a sum over terms which all contain

$$(\prod_{i=1}^{\ell+1} Q_i)P_\nu \cdot (\Delta^{(g)}_{j_g}\cdots\Delta^{(1)}_{j_g}) \ . \tag{81}$$

Using the fact that P_ν commutes with the constraints inside the pyramids, and that within an X sector the constraints of the pyramids commute within themselves, this is equal to

$$P_\nu(\prod_{i=1}^{\ell+1} Q_i)\cdot\Delta^{(1)}_{j_{g-1}}\cdot(\Delta^{(g)}_{j_g}\cdots\Delta^{(1)}_{j_{g-1}})P_\nu \tag{82}$$

But $\left(\prod_{i=1}^{\ell+1} Q_i\right)\cdot\Delta^{(1)}_{j_{g-1}}$ is identically zero as it must contain at least one term of the form $Q(\mathbb{1}-Q)$. It follows that $P_\nu|\Omega\rangle = 0$ which can only happen when $\lambda_\nu = 0$. ∎

Next, we bound the weights λ_s using the exponential decay of Sec. 3.6, which is proved in Sec. E. According to Lemma 3.3, there exists a set of weights η_s^2 such that $\sum_s \eta_s^2 \leq 1$ and for every $s \geq \ell$,

$$\lambda_s \leq \frac{1}{x}k^{g^2\ell}\left(\frac{\ell+1}{\ell!}\right)^g s^{g\ell}\theta^s\eta_s \ . \tag{83}$$

Bounding η_s by 1, we get that

$$\lambda_s \leq \frac{1}{x}k^{g^2\ell}\left(\frac{\ell+1}{\ell!}\right)^g s^{g\ell}\theta^s \ . \tag{84}$$

We are now in position to prove the main result of this section, Eq. (74). In Appendix G we use the above equation to show that it is possible to find a constant $r(\theta,k,g)$ such that for every $s > r\ell$,

$$\lambda_s^2 s \leq \frac{1}{x^2}\theta^s \ . \tag{85}$$

Consequently, from Claim F.3 it follows that

$$\langle\Omega|E^{top}|\Omega\rangle \le \ell + \sum_{s=0}^{r\ell}\lambda_s^2 s + \frac{1}{x^2}\sum_{s=r\ell+1}^{\infty}\theta^s \tag{86}$$

$$\le (r+1)\ell + \frac{1}{x^2}\frac{\theta^{(r\ell+1)}}{1-\theta} \ . \tag{87}$$

By redefining $r(\theta,k,g) \to r(\theta,k,g)+1$, and using the fact that $\theta^{(r\ell+1)} < 1$, we recover Eq. (74).

Finally, we can now prove Lemma F.1.

F.2 Proof of Lemma F.1

To prove Lemma F.1, we use Lemma F.2 to estimate the total energy of the system. To estimate the energy of the first layer, we apply Lemma F.2 several times using different XY decomposition. The XY decompositions are chosen such that every constraint in the first layer appears in exactly one XY decomposition. One can easily verify that the total number of such decompositions that is needed for this task is upper bounded by some constant $f_1(k,g)$. Therefore, the total energy of the first layer is bounded by

$$\langle\Omega|E_1|\Omega\rangle \le \begin{cases} f_1 \frac{1-x^2}{x^2}\frac{\theta^2}{(1-\theta^2)^3} & , \ \ell = 0 \\ f_1\left[r\ell + \frac{1}{x^2}\frac{1}{1-\theta}\right] & , \ \ell > 0 \end{cases} \tag{88}$$

To bound the energy of the second layer, we can apply the derivation of the first layer, with some trivial modifications:

$$g \to g-1 \ , \tag{89}$$

$$|\psi\rangle \to \Pi^{(1)}_{\le\ell}|\psi\rangle \ . \tag{90}$$

In addition, we need to update the functions $r(\theta,k,g)$ and $f_1(k,g)$. It is easy to see that both of them can be decreased. Therefore, it is not surprising to see that the upper bound of the $\langle\Omega|E_2|\Omega\rangle$ is smaller than the upper bound of $\langle\Omega|E_1|\Omega\rangle$, and this true for all the other layers. Consequently, by setting $f(k,g) \stackrel{\text{def}}{=} g f_1(k,g)$, we get:

$$\epsilon_0 \le \langle\Omega|E|\Omega\rangle \le \begin{cases} f \frac{1-x^2}{x^2}\frac{\theta^2}{(1-\theta^2)^3} & , \ \ell = 0 \\ f\left[r\ell + \frac{1}{x^2}\frac{1}{1-\theta}\right] & , \ \ell > 0 \end{cases} \ . \tag{91}$$

Here ϵ_0 is the ground energy of the system. Lemma F.1 is now proved by inverting this inequality:

$$\text{for } \ell = 0\text{:} \quad x^2 \le \frac{1}{(\epsilon_0/f)\frac{(1-\theta^2)^3}{\theta^2}+1} = 1-\Delta^2(0) \ , \tag{92}$$

$$\text{for } \ell > 0\text{:} \quad x^2 \le \frac{1}{1-\theta}\cdot\frac{1}{(\epsilon_0/f)-r\ell} = 1-\Delta^2(\ell) \ . \tag{93}$$

Note, of course, that the $\ell > 0$ inequality is only valid for $(\epsilon_0/f) - r\ell > 0$.

G Finding $r(\theta,k,g)$

In this section we prove that it is possible to find a constant $r(\theta,k,g)$ such that for every $s > r\ell$,

$$\frac{1}{x^2}\left(\frac{\ell+1}{\ell!}\right)^{2g} k^{2g^2\ell} s^{2g\ell+1}\theta^{2s} \le \frac{1}{x^2}\theta^s \ . \tag{94}$$

Eliminating a factor $\frac{\theta^s}{x^2}$ and taking a log of the equation, we find the following sufficient condition

$$2g[\log(\ell+1) - \log(\ell!)] + 2g^2\ell\log(k) + (2g\ell+1)\log(s) + s\log(\theta) < 0 \tag{95}$$

which is equivalent to

$$\frac{2g}{s}[\log(\ell+1) - \log(\ell!)] + 2g^2\frac{\ell}{s}\log(k) + (2g\ell+1)\frac{\log(s)}{s} + \log(\theta) < 0 \tag{96}$$

Re-arranging it gives

$$\frac{2g}{s}[\ell\log(s) - \log(\ell!)] + \frac{2g}{s}\log(\ell+1) + 2g^2\frac{\ell}{s}\log(k) + \frac{\log(s)}{s} < \log(1/\theta)\ . \tag{97}$$

On the LHS we have the sum of 4 terms. For $s > 3$, $\log(s)/s$ is monotonically decreasing ($\log(\cdot)$ is the natural logarithm). So if the above condition holds for $s = r\ell$ with $r > 3$, it would hold for any $s > r\ell$. Therefore, a sufficient condition is

$$\frac{2g}{s}[\ell\log(s) - \log(\ell!)] + \frac{2g}{r\ell}\log(\ell+1) + \frac{2g^2\log(k)}{r} + \frac{\log(r)}{r} < \log(1/\theta)\ . \tag{98}$$

Next, for every $\ell \geq 1$, the term $\frac{\log(\ell+1)}{\ell}$, which appears in the second element is smaller than 1, therefore a sufficient condition is

$$\frac{2g}{s}[\ell\log(s) - \log(\ell!)] + \frac{2g}{r} + \frac{2g^2\log(k)}{r} + \frac{\log(r)}{r} < \log(1/\theta)\ . \tag{99}$$

Let us now analyze the first term. Using Sterling's approximation, we get

$$\frac{2g}{s}[\ell\log(s) - \log(\ell!)] \leq \frac{2g}{s}[\ell\log(s) - \ell\log(\ell) + \ell] \tag{100}$$

$$= \frac{2g}{s/\ell}\log(s/\ell) + \frac{2g}{s/\ell}\ . \tag{101}$$

Again, using the assumption that $s/\ell > r > 3$, then $\frac{\log(s/\ell)}{s/\ell} < \log(r)/r$, and $\frac{2g}{s/\ell} < 2g/r$. So overall, we find that as long as $r > 3$, a sufficient condition for Eq. (94) is

$$(2g+1)\frac{\log(r)}{r} + \frac{4g + 2g^2\log(k)}{r} < \log(1/\theta)\ . \tag{102}$$

The LHS of the above inequality approaches zero as $r \to +\infty$, hence we can find an $r(\theta, k, g) > 3$ that satisfies it.

H Proof of Classical Amplification lemma

We consider a d-regular expander graph $G = (V, E)$ with $n = |V|$ vertices and second largest eigenvalue $0 < \lambda(G) < 1$. With every node of G we associate a variable that takes values in a finite alphabet Σ with $|\Sigma| = q$. Every edge is associated with a local constraint on the two values of the nodes in the edge. We refer to the set of constraints as a constraint system $\mathcal{C}$. Let σ denote an assignment of the variables. Then $\mathrm{UNSAT}_\sigma(G)$ is the *fraction of unsatisfied edges under the assignment* σ.

In the amplification lemma, we define a new constraint system on G using the notion of a t-walk. A t-walk on a graph G is a sequence of $t+1$ adjacent vertices, corresponding to a path of t steps on G. We denote it by $\mathbf{e} = (e_1, e_2, \ldots, e_t)$, with $e_1, \ldots, e_t$ being the edges along the path.

The new constraint system is defined as follows. Consider all possible t-walks on G, and for each t-walk $\mathbf{e} = (e_1, \ldots, e_t)$ we define a constraint that is satisfied if and only if all the constraints along the path are satisfied. With some abuse of notation, we will call the new constraint system G^t and define its UNSAT with respect to σ by

$$\text{UNSAT}_\sigma(G^t) = \frac{\#\text{ of unsatisfied } t\text{-walks}}{\text{total }\#\text{ of } t\text{-walks}} \ . \tag{103}$$

The amplification lemma shows that that by moving from G to G^t, the UNSAT is "amplified" by a factor of t, provided that $\text{UNSAT}_\sigma(G)$ is not too close to 1.

Lemma H.1 (The classical amplification lemma) *Let $G = (V, E)$ be an expander graph with second largest eigenvalue $0 < \lambda < 1$, and let $\mathcal{C}$ be a constraint system on it using an alphabet Σ. Let G^t denote the t-walk constraint system that was defined above. Define $c(\lambda) \stackrel{\text{def}}{=} \left(2 + \frac{2}{1-\lambda}\right)^{-1}$. Then for every assignment σ,*

$$\text{UNSAT}_\sigma(G^t) \geq \begin{cases} t \cdot c(\lambda) \cdot \text{UNSAT}_\sigma(G) & , \quad \text{UNSAT}_\sigma(G) \leq \frac{1}{t} \\ c(\lambda) & , \quad \text{UNSAT}_\sigma(G) \geq \frac{1}{t} \end{cases} . \tag{104}$$

Proof: Given the assignment σ, we let $F \subseteq E$ denote the set of unsatisfied edges in G. Obviously, $\text{UNSAT}_\sigma(G) = \frac{|F|}{|E|}$. Consider the homogeneous probability distribution over all t-walks. We define a random variable $Z(\mathbf{e})$ that counts the number of unsatisfied edges in the t-walk $e = (e_1, \ldots, e_t)$. Then

$$\text{UNSAT}_\sigma(G^t) = \Pr[Z(\mathbf{e}) > 0] \ . \tag{105}$$

Moreover, since $Z(\mathbf{e})$ is a non-negative random variable that is not identically 0,

$$\Pr[Z(\mathbf{e}) > 0] \geq \frac{\mathbb{E}^2[Z(\mathbf{e})]}{\mathbb{E}[Z^2(\mathbf{e})]} \ . \tag{106}$$

In what follows, we will lower-bound $\mathbb{E}[Z(\mathbf{e})]$ and upper-bound $\mathbb{E}[Z^2(\mathbf{e})]$. To do that, we write $Z(\mathbf{e}) = \sum_{i=1}^t Z_i(\mathbf{e})$, where $Z_i(\mathbf{e})$ is the random variable that is equal to 1 if the i'th edge of $\mathbf{e}$ is unsatisfied and to 0 otherwise. It is easy to see that for every i, $\mathbb{E}[Z_i(\mathbf{e})] = |F|/|E|$, and therefore

$$\mathbb{E}[Z(\mathbf{e})] = t\frac{|E|}{|F|} \ . \tag{107}$$

To bound $\mathbb{E}[Z^2(\mathbf{e})]$, we write

$$\mathbb{E}[Z^2(\mathbf{e})] = \sum_{i,j} \mathbb{E}[Z_i(\mathbf{e})Z_j(\mathbf{e})] = \sum_{i=1}^t \mathbb{E}[Z_i^2(\mathbf{e})] + 2\sum_{i<j} \mathbb{E}[Z_i(\mathbf{e})Z_j(\mathbf{e})] \ . \tag{108}$$

Note that $Z_i^2(\mathbf{e}) = Z_i(\mathbf{e})$ and so

$$\mathbb{E}[Z^2(\mathbf{e})] = t\frac{|F|}{|E|} + 2\sum_{i<j} \mathbb{E}[Z_i(\mathbf{e})Z_j(\mathbf{e})] \ . \tag{109}$$

To estimate $\mathbb{E}[Z_i(\mathbf{e})Z_j(\mathbf{e})]$, we can use the expansion properties of G, which imply that as i, j grow apart, $Z_i(\mathbf{e})$ and $Z_j(\mathbf{e})$ become more and more independent, and so $\mathbb{E}[Z_i(\mathbf{e})Z_j(\mathbf{e})] \to \mathbb{E}[Z_i(\mathbf{e})] \cdot \mathbb{E}[Z_j(\mathbf{e})] = \left(\frac{|F|}{|E|}\right)^2$. The exact statement is that for $i > j$,

$$\mathbb{E}[Z_i(\mathbf{e})Z_j(\mathbf{e})] \leq \frac{|F|}{|E|}\left(\frac{|F|}{|E|} + |\lambda|^{i-j-1}\right) \ . \tag{110}$$

The proof of this fact is standard, and is given in Ref. [4], and will therefore be omitted.

Inserting this into Eq. (109), we arrive to

$$\mathbb{E}[Z^2(\mathbf{e})] = t\frac{|F|}{|E|} + 2\frac{|F|}{|E|}\sum_{i=1}^{t}\sum_{j=i+1}^{t}\left(\frac{|F|}{|E|} + |\lambda|^{i-j-1}\right) \tag{111}$$

$$\leq t\frac{|F|}{|E|} + t(t-1)\left(\frac{|F|}{|E|}\right)^2 + \frac{2t}{1-\lambda}\frac{|F|}{|E|} \tag{112}$$

$$= t\frac{|F|}{|E|}\left(1 + \frac{|F|}{|E|}(t-1) + \frac{2}{1-\lambda}\right) . \tag{113}$$

Using Eq. (106), we get

$$\mathrm{UNSAT}_\sigma(G^t) \geq \frac{t\mathrm{UNSAT}_\sigma(G)}{1 + t\mathrm{UNSAT}_\sigma(G) + \frac{2}{1-\lambda}} \stackrel{\mathrm{def}}{=} F\big(\mathrm{UNSAT}_\sigma(G)\big) , \tag{114}$$

where $F(x) = \frac{tx}{1+tx+\frac{2}{1-\lambda}}$. If $x \leq 1/t$, $F(x) \geq \frac{tx}{2+\frac{2}{1-\lambda}}$. On the other hand, as $F(x)$ is monotonically increasing for $x > 0$, then for $x \geq 1/t$, $F(x) \geq F(1/t) = \frac{1}{2+\frac{2}{1-\lambda}}$. Setting $x = \mathrm{UNSAT}_\sigma(G)$ and recalling that $c(\lambda) \stackrel{\mathrm{def}}{=} \left(2 + \frac{2}{1-\lambda}\right)^{-1}$. completes the proof. ∎

I Proof of Claim 6.2

Proof: This follows from the classical amplification lemma. We expand $|\psi_j\rangle$ as a superposition $|\psi_j\rangle = \sum_\nu \beta_\nu|\psi_\nu\rangle$, where $|\psi_\nu\rangle$ has a well-defined value (1 or 0, namely violating or not) at each edge of G_i, with the total number of violations being exactly j. Moreover, it is easy to see that as the projection into the state $|\psi_\nu\rangle$ commutes with the projections of G_i, then

$$\mathrm{QUNSAT}_{\psi_j}(G_i) = \sum_\nu \beta_\nu^2 \cdot \mathrm{QUNSAT}_{\psi_\nu}(G_i^t) , \tag{115}$$

$$\mathrm{QUNSAT}_{\psi_j}(G_i^t) = \sum_\nu \beta_\nu^2 \cdot \mathrm{QUNSAT}_{\psi_\nu}(G_i^t) , \tag{116}$$

hence it is sufficient to prove Eq. (24) for $\mathrm{QUNSAT}_{\psi_\nu}(G_i^t)$. This, however, follows directly from the classical amplification lemma since under the state $|\psi_\nu\rangle$ the constraints of G_i have a well-defined, classical values. We can therefore treat the situation as a classical system G_c with some assignment σ and $\mathrm{UNSAT}_\sigma(G_c) = j/|E|$. According to the classical amplification lemma, if $j/|E| \leq 1/t \Leftrightarrow j \leq |E|/t$ then $\mathrm{UNSAT}_\sigma(G_c^t) \geq t \cdot c(\lambda) \cdot \frac{j}{|E|}$, otherwise, $\mathrm{UNSAT}_\sigma(G_c^t) \geq c(\lambda)$. But as everything is classical for G_i and G_i^t in the ν sector then,

$$\mathrm{UNSAT}_\sigma(G_c^t) = \mathrm{QUNSAT}_{\psi_\nu}(G_i^t) \tag{117}$$

and this proves the claim. ∎

J The last step in the quantum gap amplification proof

Here we give the last step in proving the quantum amplification by using Eqs. (25,27). The idea is to consider two possible cases: $\frac{\epsilon_0}{f} - \frac{4}{1-\theta} \leq 2r$ (the low-energy case) and $\frac{\epsilon_0}{f} - \frac{4}{1-\theta} > 2r$ (the high-energy case). For the former we use the $\ell = 0$. In the former, we use the $\ell > 0$ detectability lemma. Let us start with the low energy case.

J.1 The low energy case: $\frac{\epsilon_0}{f} - \frac{4}{1-\theta} \le 2r$

Here we estimate the amplification using the $\ell = 0$ detectability. Specifically, Lemma 5.1 ensures us that there a layer i in which,

$$\alpha_1^2 + \alpha_2^2 + \ldots + \alpha_{|E|}^2 \ge \frac{1}{(2g)^2}\Delta^2(0) \ . \tag{118}$$

On the other hand, it is easy to see that Eq. (27) implies

$$\frac{\text{QUNSAT}(G^t)}{\text{QUNSAT}(G)} \ge t\frac{c(\lambda)}{\epsilon_0}\left(\alpha_1^2 + \alpha_2^2 + \alpha_3^2 + \ldots + \alpha_{|E|}^2\right) \ . \tag{119}$$

Therefore,

$$\frac{\text{QUNSAT}(G^t)}{\text{QUNSAT}(G)} \ge t \cdot c(\lambda) \cdot (2g)^{-2} \cdot \frac{\Delta^2(0)}{\epsilon_0} \ . \tag{120}$$

Let us now lower bound the expression $\frac{\Delta^2(0)}{\epsilon_0}$. $\Delta^2(0)$ is a continuous function of ϵ_0 that is bounded between 0 and 1 for $\epsilon_0 \ge 0$. We have to worry about to things: (i) if ϵ_0 becomes too large, the ratio might become small, and (ii) as $\epsilon_0 \to 0$, also $\Delta^2(0) \to 0$. The first worry is taken cared by fact that in the low-energy case ϵ_0 is upper bounded by $\frac{\epsilon_0}{f} - \frac{4}{1-\theta} \le 2r$. The second one is taken cared by noticing the approach of $\Delta^2(0)$ to 0 as ϵ_0 is linear in ϵ_0 (see Eq. (18)). Therefore as $\epsilon_0 \to 0$, the ratio approaches some positive constant. All in all, we conclude that in the low-energy case,

$$\frac{\text{QUNSAT}(G^t)}{\text{QUNSAT}(G)} \ge t \cdot c(\lambda) \cdot K_1(q, d, \theta) \ . \tag{121}$$

J.2 The high energy case: $\frac{\epsilon_0}{f} - \frac{4}{1-\theta} \ge 2r$

In the high-energy case, we use the detectability lemma with a particular ℓ to show the amplification. We choose ℓ as large as possible so that $(\ell+1)/\epsilon_0$ will be lower-bounded by a positive function of q, d, θ. Specifically, the high energy condition implies $\frac{\epsilon_0}{f} - \frac{2}{1-\theta} \ge 2r$, and so we choose[1]

$$\ell = \left\lfloor \frac{1}{r}\left(\frac{\epsilon_0}{f} - \frac{2}{1-\theta}\right)\right\rfloor \ge 2 \ . \tag{122}$$

Then on one hand,

$$\ell < \frac{1}{r}\left(\frac{\epsilon_0}{f} - \frac{2}{1-\theta}\right) \ , \tag{123}$$

and so $(1-\theta)\left(\frac{\epsilon_0}{f} - r\ell\right) > 2$, yielding a finite detectability in Lemma 5.1:

$$\Delta^2(\ell) > 1 - \frac{1}{2} = \frac{1}{2} \tag{124}$$

$$\Downarrow \tag{125}$$

$$\alpha_{\ell+1}^2 + \alpha_{\ell+2}^2 + \ldots + \alpha_{|E|}^2 \ge \frac{1}{(2g)^2}\Delta^2(\ell) \ge \frac{1}{8g^2} \ . \tag{126}$$

[1]Note that by assumption ϵ_0 is larger than $2r$, which can only happen when $|E|$ – the total number of constraints in the system – satisfies $|E| > 2r$, therefore the ℓ we choose makes sense.

On the other hand, Eq. (122) also implies

$$\ell + 1 \geq \frac{1}{r}\left(\frac{\epsilon_0}{f} - \frac{2}{1-\theta}\right) \ , \tag{127}$$

and so

$$\frac{\ell+1}{\epsilon_0} \geq \frac{1}{r}\left(\frac{1}{f} - \frac{2}{\epsilon_0(1-\theta)}\right) \ . \tag{128}$$

But $\frac{\epsilon_0}{f} - \frac{4}{1-\theta} > 0$, therefore

$$\frac{\ell+1}{\epsilon_0} \geq \frac{1}{2fr} \ . \tag{129}$$

Let us now return to Eq. (27). By omitting all the α_i^2 terms with $i \leq \ell$, we obtain

$$\frac{\mathrm{QUNSAT}(G^t)}{\mathrm{QUNSAT}(G)} \geq t\frac{c(\lambda)}{\epsilon_0}(\ell+1)\left(\alpha_{\ell+1}^2 + \ldots \alpha_{|E|/t}^2\right) + c(\lambda)\cdot\frac{|E|}{\epsilon_0}\left(\alpha_{|E|/t+1}^2 + \ldots + \alpha_{|E|}^2\right) \tag{130}$$

Define

$$A = \alpha_{\ell+1}^2 + \ldots + \alpha_{|E|/t}^2 \ , \tag{131}$$
$$B = \alpha_{|E|/t+1}^2 + \ldots + \alpha_{|E|}^2 \ . \tag{132}$$

Then $A + B \geq \frac{1}{8g^2}$ and

$$\frac{\mathrm{QUNSAT}(G^t)}{\mathrm{QUNSAT}(G)} \geq t\frac{c(\lambda)}{\epsilon_0}(\ell+1)A + c(\lambda)\cdot\frac{|E|}{\epsilon_0}B \ . \tag{133}$$

If $A \geq \frac{1}{16g^2}$ then from the above equation and by Eq. (129),

$$\frac{\mathrm{QUNSAT}(G^t)}{\mathrm{QUNSAT}(G)} \geq t\cdot c(\lambda)\cdot\frac{\ell+1}{\epsilon_0}A \geq t\cdot c(\lambda)\cdot\frac{1}{2fr}\cdot\frac{1}{16g^2} \stackrel{\mathrm{def}}{=} t\cdot c(\lambda)\cdot K_2(q,d,\theta) \ . \tag{134}$$

If, on the other hand, $B \geq \frac{1}{16g^2}$ then we can use Eq. (25) to conclude that

$$\mathrm{QUNSAT}(G^t) \geq c(\lambda)\cdot\frac{1}{16g^2} \stackrel{\mathrm{def}}{=} c(\lambda)\cdot K_3(d) \ . \tag{135}$$

Combining all these 3 results, it is straightforward to define a function $K(q,d,\theta)$ such that

$$\mathrm{QUNSAT}(G^t) \geq c(\lambda)\cdot K(q,d,\theta)\cdot\min\left\{t\cdot\mathrm{QUNSAT}(G), 1\right\} \ . \tag{136}$$