The Detectability Lemma and Quantum Gap Amplification

Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh Vazirani

March 26, 2009

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: the quantum object corresponding to a max-k-SAT formula $f = c_1 \ldots c_m$ is a local Hamiltonian $H = H_1 + \ldots + H_m$, where each term H_j acts non-trivially on at most k qubits. We may assume $H_j = I - P_j$ for some projection operator P_j, so that H_j may be thought of as imposing the constraint that the quantum state should lie in the null space of H_j (the range of P_j). Very roughly, then, there is an analogy between the total energy of a given n qubit state acted upon by H and the number of violated constraints in a k-SAT formula. In a nutshell, this is the content of Kitaev's quantum analog of the Cook-Levin theorem. 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 first issue is that a quantum state can make an arbitrary angle with the null space of H_j, and therefore we can only speak of the degree to which it violates the constraint H_j. This is easily dealt with by speaking about the probability that each H_j is violated. A more serious problem is that the terms of the Hamiltonian may not commute in general. This means that it is not even meaningful to ask: what is the probability that a given state violates at least k constraints. Our main result in this paper is the *quantum detectability lemma*, which provides a way of making sense of and answering these questions.

The detectability lemma helps carry over some of our classical intuition about constraint satisfaction problems to the quantum context. The hope is that this will prove to be a useful tool in the newly evolving area of quantum Hamiltonian complexity, which provides a bridge between classical complexity theory and condensed matter physics. In this paper, we apply the detectability lemma to gap amplification, a technique first introduced in the context of randomness efficient amplification of the success probability of RP and BPP algorithms, and more recently applied by Dinur in the context of the PCP theorem. Indeed, our research was originally motivated by the question of whether there is a quantum analog of the PCP theorem, and quantum gap amplification translates one of several steps in Dinur's proof to a quantum setting.

Before we can state the detectability lemma, we must identify some further structure of local Hamiltonians. We will assume that the terms of the Hamiltonian H can be partitioned into a constant number(g) of sets, called layers, such that all terms within a layer commute (and therefore can be simultaneously measured). Such a partitioning is possible under the mild restriction that the Hamiltonian is of *constant degree*: every particle (qubit) in the local Hamiltonian participates in a bounded number of the local constraints. Let us denote by Π_k the product of all the projections participating in the kth layer, which is simply the projection on the common ground space of all projections in the layer. The probability that a given state $|\psi\rangle$ lies in this common ground space is simply $||\Pi_k|\psi\rangle||^2$, and therefore the probability that at least one constraint in layer k is violated is $1 - ||\Pi_k|\psi\rangle||^2$. Now suppose that the ground state energy $\langle\psi|H|\psi\rangle$ of state $|\psi\rangle$ is ϵ_0. The detectability lemma states:

$||\Pi_1 \cdots \Pi_g|\psi\rangle||^2 \leq 1 - c(g)\epsilon_0$, for some constant $c(g)$. This means that if we sequentially test whether any constraint in layer k is violated, starting with $k = g$ down to $k = 1$, then we will detect a violation with probability at least $c(g)\epsilon_0$. A more general version of the lemma lowerbounds the probability that at least l constraints are violated.

A corollary to the detectability lemma is that for any state $|\psi\rangle$, there is some layer j, such that the probability that at least one contraint in layer j is violated by $|\psi\rangle$ is at least $c'(g)\epsilon_0$.

2 Detectability Lemma

Many of the main components of the proof exist in the following special example consisting of two layers: the Hamiltonian $H = \sum_{i+1}^{n-1} H_i$ will be such that H_i will be two-local and acting on the i-th and $i+1$-st qubit. The first layer will consist of the constraints (all of which commute with each other because they act on different qubits) $\{H_1, H_3, \dots\}$ and the second layer consists of the constraints $\{H_2, H_4, \dots\}$; call these the *odd* and *even* layers respectively. We let $P_i = 1 - H_i$ and let $\Pi_{\text{odd}} = P_1 P_3 \dots$, $\Pi_{\text{even}} = P_2 P_4 \dots$ be the projection onto the ground space for each layer. Within this example, our main theorem says that $||\Pi_{\text{odd}}\Pi_{\text{even}}|| < 1 - c$ for a constant that does not depend on n, the number of qubits.

Now, think of the n qubit space as a tensor product of consecutive 4-qubit spaces, consisting of the first four, the second four, etc. Notice that many of the P_i act only on one of these 4-qubit spaces; the only P_i that act between spaces are $P_4, P_8 \dots$. Consider the first 4-qubit spaces along with P_1, P_2, P_3 all of which act there. We split that space into two orthogonal subspaces, one where the three operators all commute which we call X and the rest, which we call Y. We do the analogous thing on all the other 4-qubit spaces and if we then expand according to the 4 qubit tensor product structure, we end up with a decomposition of the space as a direct sum of subspaces, each of which are a product of local X or Y subspaces. Each term is an invariant subspace for the action of all the P_i's except $\{P_4, P_8 \dots\}$; and thus analysis of sums and products of combinations of these P_i's can be done locally. We'll call this an XY decomposition of the space.

The essential analysis boils down to the following observation: the operator $P_1 P_3 P_2$ has norm strictly less than one (call it θ) on the Y component for the first 4 qubits, since an eigenvector with eigenvalue 1 would be an invariant 1-dimensional space for all three operators and hence would have been in the X part of the decomposition. This observation, when used on each of the 4-qubit components and with a little more work, essentially shows that $\Pi_{\text{odd}}\Pi_{\text{even}}$ produces a multiplicative decay factor of θ for each Y component of a given part of an XY decomposition.

This shows that for any vector ψ, the vector $\Pi_{\text{odd}}\Pi_{\text{even}}|\psi\rangle$ has to have exponentially small mass in the components of the XY decomposition with many Y terms. It also shows that for $||\Pi_{\text{odd}}\Pi_{\text{even}}|\psi\rangle||$ to be close to 1, it must have less than constant mass in those components with at least one Y term. Together, these two properties imply that the energy contribution (which can be simply analyzed using the XY decomposition) of $H - (H_4 + H_8 + \dots$ on $\Pi_{\text{odd}}\Pi_{\text{even}}|\psi\rangle$ is less than a constant. Combining this with the same analysis with different groupings of the qubits (i.e. different XY decompositions) allows the same conclusion for the energy contributions of the other H_i thus contradicting the energy assumption and thus no such ψ can exist.

The general proof has the same essence as the proof in this example. In the general setting of more layers and not nearest neighbor local constraints, we use the term *pyramid* to describe the spaces and operators corresponding to the 4-qubit spaces and the operator $P_1 P_3 P_2$ in the above example. The motivation for this terminology comes from describing $P_1 P_3 P_2$ as a single operator P_2 in the even layer and the two projections "underneath" (i.e. that act on common qubits) in the odd layer. The XY decomposition then corresponds to grouping these pyramids into non-overlapping groups (which we term *Ponzi's*) that produces the comparable tensor product structure. In our example above, it would only take one other XY decomposition to cover all the constraints; in the general case it will take a finite number (but more than two) Ponzis to cover all the constraints.