

1 QMA

The class NP (non-deterministic polynomial time) contains many thousand of the most important computational problems. Of these problems, the vast majority are NP-complete. This means that these are the hardest problems in NP. By this we mean that, if anyone of them can be solved by a polynomial time algorithm, then every problem in NP can be solved by a polynomial time algorithm. The cornerstone of this theory of NP-completeness is the Cook-Levin theorem, which states that 3-SAT is NP-complete.

A language L is in NP if there is a polynomial time proof checker C and a polynomial $poly$, with the following property: if $x \in L$ then there is a string y with $|y| \leq poly|x|$, such that $C(x,y) = 1$. If $x \notin L$, then for every y such that $|y| \leq poly(|x|)$, $C(x,y) = 0$.

Kitaev gave the quantum analogue of the Cook-Levin theorem by showing that QSAT the quantum analogue of 3-SAT is complete for the quantum analogue of NP, called BQNP or QMA.

QMA is the quantum generalization of MA — the probabilistic analogue of NP. To define MA, we simply replace the deterministic polynomial time proof checker with a probabilistic polynomial time proof checker C . Now if $x \in L$, then there is a string y with $|y| \leq poly|x|$, such that $C(x,y) = 1$ with probability at least $2/3$. If $x \notin L$, then for every y such that $|y| \leq poly(|x|)$, $C(x,y) = 0$ with probability at least $2/3$.

To define BQNP, the quantum analogue of MA, we replace the probabilistic polynomial time proof checker by a quantum polynomial time proof checker. Equally important, the witness string y is now allowed to be a quantum witness, i.e., it can be a superposition over strings of length at most $poly(|x|)$.

BQP is trivially contained in BQNP since it can be simulated by the verifier alone. MA is also contained in BQNP since quantum machines can perform the classical computations of their classical counterparts. Kitaev's proof that QSAT is BQNP-complete implies a non-trivial upper bound, showing that $BQNP \subseteq P^{#P}$.

A QMA-Complete Problem

Recall that a Hamiltonian acting on n qubits is a 2^n dimensional Hermitian matrix. Say that a Hamiltonian is c -local if it acts as the identity on all except c of the qubits. Consider the following problem: **Local Hamiltonians or Q5SAT:** Let H_j (for $j = 1, \dots, r$) be 5-local Hamiltonians on n qubits (each specified by complex $2^5 \times 2^5$ matrices.). Assume that each H_j is scaled so that all eigenvalues λ of H_j satisfy $0 \leq \lambda \leq 1$. Let $H = \sum_{j=1}^r H_j$. There is a promise about H that either all eigenvalues of H are $\geq b$ or there is an eigenvalue of H that is $\leq a$, where $0 \leq a < b \leq 1$ and the difference $b - a$ is at least inverse polynomial in n , i.e., $b - a \geq \frac{1}{poly(n)}$. The problem asks whether H has an eigenvalue $\leq a$.

The Connection with 3-SAT

In 3-SAT, we are given a formula f on n variables in 3-CNF (conjunctive normal form.) That is, f is a conjunction of many clauses c_i :

$$f(x_1, x_2, \dots, x_n) = c_1 \wedge c_2 \wedge \dots \wedge c_m,$$

where each clause c_j is a disjunction of three variables or their negations. For example, c_j may be $(x_a \vee \bar{x}_b \vee x_c)$.

We would like to make a corresponding Hamiltonian H_i for each clause c_i . H_i should penalize an assignment which does not satisfy the clause c_i . In the example where $c_j = (x_a \vee \bar{x}_b \vee x_c)$, we want to penalize the assignment state $|010\rangle$. If our notion of *penalize* is to have a positive eigenvalue, then we can let $H_j = |010\rangle\langle 010|$, and define the other H_i 's similarly, i.e., each H_i has a 1 eigenvalue with a corresponding eigenvector that causes clause c_i to be false.

Finally, we let

$$H = \sum_{i=1}^m H_i,$$

so that H is a sum of 3-local Hamiltonians. It is not hard to see that the smallest eigenvalue of H is the minimum (over all assignments) number of unsatisfied clauses. In particular, H has a 0 eigenvalue exactly when there is a satisfying assignment for f .

For general QSAT instances, the Hamiltonians H_j cannot be simultaneously diagonalized in general, and the problem appears much harder.

Membership in QMA

We can assume without loss of generality that each H_j is just a projection matrix $|\phi_j\rangle\langle\phi_j| \otimes I$. The prover would like to provide convincing and easily verifiable evidence that $H = \sum H_j = \sum (A_j \otimes I)$ has a small eigenvalue $\lambda \leq a$. The proof consists of (a tensor product of) polynomial in n copies of the corresponding eigenvector $|\eta\rangle$.

$\lambda = \sum_j \langle \eta | H_j | \eta \rangle$. Given a single copy of $|\eta\rangle$, the verifier can flip a coin with bias $\frac{\lambda}{r}$ as follows:

1. Pick $H_j = |\phi_j\rangle\langle\phi_j|$ at random
2. Measure $|\eta\rangle$ by projecting onto $|\phi_j\rangle$.

This succeeds with probability $\frac{\lambda}{r}$. Given the promise that $\lambda \leq a$ or $\lambda \geq b$, it suffices for the verifier to repeat this test $\frac{r^2}{(b-a)^2}$ times to conclude with high confidence that $\lambda \leq a$. Thus polynomial in n copies of $|\eta\rangle$ are sufficient. Note that since the verifier is performing each test randomly and independently, the prover gains no advantage by sending an entangled state to the verifier. Exercise: prove this.

QMA-Completeness

To show that QSAT is complete in QMA, we need to show that the universal BQNP problem reduces to it. That is, given a quantum circuit $U = U_L U_{L-1} \dots U_1$ and a promise that exactly one of the following holds:

1. $\exists |\eta\rangle, U$ accepts on input $|\eta\rangle$ with probability $\geq p_1 = 1 - \epsilon$

2. $\forall |\eta\rangle, U$ accepts on input $|\eta\rangle$ with probability $\leq p_0 = \epsilon$,

The challenge is to design an instance of QSAT which allows us to distinguish the above two cases. i.e. we wish to specify a sum of local Hamiltonians that has an eigenvector with small eigenvalue if and only if $\exists |\eta\rangle$ that causes U to accept with high ($\geq p_1$) probability.

The construction of the local Hamiltonian is analogous to Cook's theorem. The quantum analogue of the accepting tableau in Cook's theorem will be the computational history of the quantum circuit:

$$|T\rangle = \sum_{t=0}^L |\phi_t\rangle \otimes |t\rangle$$

where $|\phi_0\rangle$ is a valid initial state and $|\phi_i\rangle = U_i |\phi_{i-1}\rangle$. Thus the computation history $|T\rangle$ is an element of $(\mathcal{C}^2)^{\otimes n} \otimes \mathcal{C}^{L+1}$. It is a superposition over time steps of the state of the quantum bits as the quantum circuit operates on them.

Now the idea of the QMA-completeness proof is to design the hamiltonian H such that:

1. if there exists $|\eta\rangle$ where U accepts with probability at least $1 - \epsilon$, then the computational history $|T\rangle$ of the quantum circuit U on input η is an eigenvector with eigenvalue at most $\frac{\epsilon}{L+1}$
2. if U rejects every input with probability at least $1 - \epsilon$, then all the eigenvalues of H are at least $\frac{c(1-\epsilon)}{L^3}$

H will be the sum $H_{initial} + H_{final} + H_{propagate}$. The first two terms are simple and express the condition that the computational history starts with a valid input state, and ends in an accepting state.

We consider the first m bits of U 's state the input bits and the remaining $n - m$ bits to be the clean work bits. The design of the $H_{initial}$ component should then reflect that at time 0, all of the work bits are clear:

$$H_{initial} = \sum_{s=m+1}^n \Pi_s^{(1)} \otimes |0\rangle \langle 0|$$

where $\Pi_s^{(1)}$ denotes projection onto the s -th qubit with value 1.

Assume that the state of the first qubit at the output determines whether or not the input is accepted. Then H_{final} needs to indicate that at time L the first qubit is a 1:

$$H_{final} = \Pi_1^{(0)} \otimes |L\rangle \langle L|.$$

The most complicated component of H is $H_{propagate}$, which captures transitions between time steps. $H_{propagate} = \sum_{j=1}^L H_j$, where

$$\begin{aligned} H_j = & -\frac{1}{2} U_j \otimes |j\rangle \langle j-1| \\ & -\frac{1}{2} U_j^\dagger \otimes |j-1\rangle \langle j| \\ & + \frac{1}{2} I \otimes (|j\rangle \langle j| + |j-1\rangle \langle j-1|) \end{aligned}$$

The fact that the computational history is a superposition over time steps is quite crucial here. To check that the correct operation has been applied in step j , it suffices to restrict attention to the $j - 1$ -st and j -th bit of the clock (assuming that the clock is represented in unary). Now the quantum register is in a superposition over its state at time $j - 1$ and at time j . Locally checking this superposition is sufficient to determine whether its clock j component is the result of applying the quantum gate U_j to the clock $j - 1$ component. This is precisely what the Hamiltonian H_j above is designed to do.

Next we show that an accepting history of computation is an eigenvector of H with eigenvalue 0.

Let $|T\rangle = \sum_{t=0}^L |\phi_t\rangle \otimes |t\rangle$. We analyze the contribution from each component of H . If $|T\rangle$ starts with qubits $m + 1$ through n clear, $H_{initial}$ does not contribute to $H|T\rangle$. If $|T\rangle$ is a computation of U , that is, $|\phi_t\rangle = U_t |\phi_{t-1}\rangle$ for all t , then from $H_{propagate}$ we get:

$$\begin{aligned} H_j T &= -\frac{1}{2} U_j |\phi_{j-1}\rangle |j\rangle - \frac{1}{2} U_j^\dagger |\phi_j\rangle |j-1\rangle \\ &\quad + \frac{1}{2} |\phi_j\rangle |j\rangle + \frac{1}{2} |\phi_{j-1}\rangle |j-1\rangle \\ &= -\frac{1}{2} |\phi_j\rangle |j\rangle - \frac{1}{2} |\phi_{j-1}\rangle |j-1\rangle \\ &\quad + \frac{1}{2} |\phi_j\rangle |j\rangle + \frac{1}{2} |\phi_{j-1}\rangle |j-1\rangle \\ &= 0, \end{aligned}$$

for no contribution from $H_{propagate}$.

Finally, if U accepts with probability at least $1 - \epsilon$, only H_{final} contributes a penalty to the sum, for an eigenvalue of at most $\frac{\epsilon}{L+1}$.

The hard part of the proof lies in showing the converse. That if there is no $|\eta\rangle$ which U accepts with high probability, then all eigenvalues of H are large. We refer the interested reader to [?] for the proof of this.

Upper bound on QMA

One consequence of this proof of QMA-completeness is the following:

Theorem: $QMA \subseteq P^{\#P}$.

Replace H with $I - H$, so it either has an eigenvalue greater than or equal to $a' = 1 - a$ or all eigenvalues are smaller than $b' = 1 - b$. Consider the trace of H^k . This is either at least a'^k or at most Nb'^k . We can make sure that $a'^k \gg Nb'^k$, by choosing $k \gg n^d \log N$. So we just need to estimate $Tr(H^k)$ in $P^{\#P}$.

To see this, write $Tr(H^k) = Tr((\sum_j H_j)^k) = Tr(\sum_{j_1, \dots, j_k} H_{j_1} \cdots H_{j_k}) = \sum_{j_1, \dots, j_k} Tr(H_{j_1} \cdots H_{j_k})$. Each trace in this sum is itself just a sum of exponentially many easy to compute contributions, and thus the entire sum is easily seen to be estimated in $P^{\#P}$.