

Quantum NP - A Survey

Dorit Aharonov*and Tomer Naveh†

Abstract

We describe Kitaev's result from 1999, in which he defines the complexity class QMA, the quantum analog of the class NP, and shows that a natural extension of 3-SAT, namely local Hamiltonians, is QMA complete. The result builds upon the classical Cook-Levin proof of the NP completeness of *SAT*, but differs from it in several fundamental ways, which we highlight. This result raises a rich array of open problems related to quantum complexity, algorithms and entanglement, which we state at the end of this survey. This survey is the extension of lecture notes taken by Naveh for Aharonov's quantum computation course, held in Tel Aviv University, 2001.

1 Introduction

The field of complexity theory has witnessed several fundamental results in the recent decade or two; It is now a rich field involving deep questions and leading to the discovery of beautiful and unexpected structures, with important contributions to the understanding of the notion of classical probabilistic and deterministic computation. With the stormy entrance of quantum computation into the life of theoretical computer scientists, it seems only natural to ask whether such a rich theory of complexity can also be developed for the quantum model; it is probably true that such interesting structures and results await for us down the road of quantum complexity theory, with perhaps insights to be drawn from them regarding the quantum computational power. Several important results have already been discovered[16, 11], and there are surely more to come. It is not unreasonable to also hope that quantum complexity can significantly contribute to the understanding of classical complexity in unexpected ways; A puzzling example in which quantum arguments are used in order to prove an entirely classical result in the area of locally decodable codes was recently found[8]. We thus view the development of the field of quantum complexity as an important direction that holds the promise of a rich area of study with possible implications to the understanding of quantum algorithmic theory, as well as to classical complexity theory and to the foundations of quantum physics.

Perhaps the most basic and fundamental result in classical complexity theory, is the Cook-Levin theorem[14], which states that *SAT*, the problem of whether a Boolean formula is satisfiable or not, is *NP* complete. This result opened the door to the study of the extremely

*School of Engineering and Computer Science, The Hebrew University, Jerusalem, Israel, and the Mathematical Sciences Research Institute, Berkeley, California

†Department for Computer Science, Tel Aviv University, Tel Aviv, Israel

expressive complexity class NP , and the rich theory of NP -completeness, and was an important building block in many later results in theoretical computer science and complexity theory, such as the PCP theorems, hardness of approximation results and the proof that $IP = PSpace$. In the heart of this result stands the very basic understanding that computation is local.

We devote this manuscript to the survey of a result by Kitaev[9, 10], which is the quantum analog of the Cook-Levin Theorem. Kitaev first defines the quantum analog of NP , and then defines a complete problem which can be viewed as a generalization of SAT to the quantum world. The proof follows the lines of the Cook-Levin proof, but defers from it in some fundamental points; We highlight those as we go along. The classical proof is quite simple; The quantum counterpart is rather complicated and long. However, there are several reasons to study this theorem, apart from the elegance of the proof, and from the naturalness of the question. First, there is a lot to be learned from the comparison of the classical proof and its (much more involved) quantum counterpart; Understanding the exact places where those differ is insightful. Secondly, the result raises a rich array of natural and interesting open problems related to this subject; We list those at the end of the survey, after the proof. Our proof follows closely the proof given by Kitaev[9, 10], with minor deviations; Our main contribution here is adding explanations and clarifications, hopefully providing some intuition behind the proof, and highlighting some open problems. We hope that this survey will provide an easy access to Kitaev's fundamental result and to the rich array of open questions it raises.

2 Definition of QMA

We would like to define a complexity class which will be the quantum analog of NP:

Definition 1 NP: $L \in NP$ if there exists a deterministic polynomial time verifier V such that:

- $\forall x \in L \exists y \ |y| = poly(|x|), V(x, y) = 1.$
- $\forall x \notin L \forall y \ |y| = poly(|x|), V(x, y) = 0.$

By $|x|$ we mean the number of bits in the binary string x . However, when trying to define the quantum analog, we immediately encounter an obstacle. We cannot require the verifier to answer 0 or 1 deterministically, because we will not be able to distinguish between this case and the case in which the verifier outputs these values with extremely high probability. Since the fact that states are continuous is inherent to quantum computation, we resort to defining QMA, the quantum analog of MA, which is the probabilistic version of NP.

Informally, MA can be thought of as a probabilistic analog of NP, allowing for two-sided errors.

Definition 2 MA: $L \in MA$ if there exists a probabilistic polynomial time verifier V such that:

- $\forall x \in L \exists y \ |y| = poly(|x|), Pr(V(x, y) = 1) \geq \frac{2}{3}$

- $\forall x \notin L \quad \forall y \quad |y| = \text{poly}(|x|), \Pr(V(x, y) = 1) \leq \frac{1}{3}$

MA is naturally viewed as a game or interaction between 2 parties - Merlin, which has infinite computational power, and Arthur, which is limited to a polynomial time machine (the above V). Merlin should answer queries such as “is $x \in L$?”, and accompany the answer with a polynomial witness y which Arthur can verify in polynomial time. Note that when showing that a problem is in MA, we should also show that Merlin cannot fool Arthur - i.e. that when $x \notin L$ there is no witness y that can persuade the verifier to believe that $x \in L$ with probability $\geq \frac{1}{3}$.

We will define QMA analogously, where the verifier V is a quantum machine, and the witness y is a state of a polynomial number of qubits. We denote by \mathcal{B} the Hilbert space of one qubit.

Definition 3 QMA: $L \in \text{QMA}$ if there exists a quantum polynomial time verifier V and a polynomial p such that:

- $\forall x \in L \quad \exists |\xi\rangle \in \mathcal{B}^{p(|x|)}, \Pr(V(|x\rangle|\xi)) = 1) \geq 2/3$
- $\forall x \notin L \quad \forall |\xi\rangle \in \mathcal{B}^{p(|x|)}, \Pr(V(|x\rangle|\xi)) = 1) \leq 1/3$

Another possible definition would be to take $|\alpha\rangle$ as a classical witness, i.e. a basis state, but leave V to be a quantum machine. We call this class Quantum Classical MA (QCMA).

Definition 4 QCMA: $L \in \text{QCMA}$ if there exists a quantum polynomial time verifier V and a polynomial p such that:

- $\forall x \in L \quad \exists y \quad |y| = \text{poly}(|x|), \Pr(V(|x\rangle|y)) = 1) \geq 2/3$
- $\forall x \notin L \quad \forall y, \quad |y| = \text{poly}(|x|), \Pr(V(|x\rangle|\alpha)) = 1) \leq 1/3$

Claim 1 $MA \subseteq QCMA \subseteq QMA$.

Proof: The left inclusion is trivial. The right inclusion follows from the fact that the quantum verifier can force Merlin to send him a classical witness by measuring the witness before applying on it the quantum algorithm. \square

It is unclear whether the two classes, $QCMA$ and QMA are the same; See open question 5 for further discussion. In any case, for the purposes of this paper, we will limit ourselves to the class QMA, where the witnesses are quantum.

2.1 Amplification

In all the above definitions of MA , $QCMA$ and QMA , we have used as our completeness parameter (i.e. one minus the error probability in case $x \in L$) the value $2/3$ and as our soundness parameter (the bound on error probability in case $x \notin L$) the value $1/3$. We can denote this choice by $MA(2/3, 1/3)$ or $QMA(2/3, 1/3)$. In general, we can define the classes $MA(c, s)$ or

$\text{QMA}(c, s)$ with general completeness and soundness parameters, which are functions of the input's length. It turns out that we have a lot of freedom in the choice of these parameters, and we can make them either polynomially close to each other, or exponentially close to 1 or 0, without changing the complexity classes we are dealing with. In other words, amplification of the completeness and the soundness from polynomial separation to exponentially small error can be done in polynomial overhead. This is done using parallel repetition and taking the appropriate majority. More formally, for the case of the classical class MA :

Theorem 1 $MA(c, c - 1/n^g) \subseteq MA(2/3, 1/3) = MA(1 - e^{-n^g}, e^{-n^g})$ where we require g to be a constant and $0 < c, c - 1/n^g < 1$.

Proof: If c and s are separated by some $1/\text{poly}(n)$, we run the verifier polynomially many times, say m , using independent random coins at each time. In case $x \in L$ the expected number of acceptances is at least cm , whereas in case $x \notin L$ it is at most sm ; The Chernoff bound¹[13] guarantees that we can distinguish between the cases with only polynomially number of independent experiments with exponentially small error. This proves the inclusions $\mathcal{MA}(c, c - 1/n^g) \subseteq \mathcal{MA}(2/3, 1/3) \subseteq \mathcal{MA}(1 - e^{-n^g}, e^{-n^g})$ where the other direction is trivial. \square

Hence, we can conveniently move between the definition of MA with either one of these three possible choices of parameters.

Remark 1 *The class MA as we defined it has two sided errors; In fact, this class is equivalent to MA with only one sided error, i.e. with completeness 1 and soundness bounded away from $1/19, 7/8$. It is unclear whether the same holds in the quantum case; See open question 4.*

This nice freedom in the choice of parameters, due to the parallel repetition, holds also in the quantum case, however with a slightly more complicated proof.

Theorem 2 $QMA(c, c - 1/n^g) \subseteq QMA(2/3, 1/3) = QMA(1 - e^{-n^g}, e^{-n^g})$ where we require g to be a constant, $0 < c, c - 1/n^g < 1$.

Proof: The proof of this theorem is slightly more subtle than the simple proof in the classical case. We will first prove that $QMA(2/3, 1/3)$ is contained in $QMA(1 - e^{-n^g}, e^{-n^g})$, i.e. that we can amplify soundness and completeness exponentially. The idea of remains the same as in the classical case: the verifier should perform polynomially many independent experiments and output the majority vote. However, unlike in the classical case, the verifier cannot perform many independent experiments on the same witness provided by the prover since after measuring it the witness will have changed; Neither can the verifier copy the quantum witness state before verifying it, due to the no cloning theorem[18] which states that an unknown quantum state cannot be copied. The verifier thus needs to ask the prover to provide him with polynomially many copies of the witness. This is problematic, since the prover might try

¹The Chernoff bound guarantees that the average of polynomially many repetitions of independent experiments will converge exponentially fast to the expected value

to cheat by entangling the witnesses he provides. We will have to show that such a strategy cannot help the prover in case x is not in the language.

We construct a new verifier which runs in parallel polynomially many copies of the verifier V , then outputs the majority. The existence of a witness for the new verifier in case $x \in L$ is trivial since it is simply duplicate copies of the original witness. To prove soundness, one might suspect that entanglement between the provers can be used to bypass the fact that the error goes exponentially to 0. To show this cannot happen, we treat the verifiers as if they are applied one after the other, and not in parallel. This is correct since the verifiers operate on different qubits and so they commute. We know that the probability that the first copy of V outputs 1 is less than $1/3$. After the first verifier was applied, we can apply the second verifier. The second verifier gets as an input some state, which can be conditioned on the result of the measurement of the first verifier. However, regardless of what this output was, it is still correct that the probability for an output 1 is less than $1/3$. And so on for the remaining of the verifiers. Hence, the probability for the majority of the verifiers being 1 can be bounded from above by the probability that polynomially many independent Bernoulli trials with bias $1/3$ will be 1, which decays exponentially by Chernoff.

The idea of the inclusion $\text{QMA}(c, c - 1/n^g) \subseteq \text{QMA}(2/3, 1/3)$ is exactly the same, except that instead of majority vote among the polynomially many verifiers, we need to count the number of accepting verifiers, and accept only if this number is above $(c + s)/2$ times the number of experiments. The other inclusions are trivial. \square

In the rest of the survey, we will interchange between the different choices of parameters according to our convenience.

2.2 Complexity

Before we continue, let us summarize what is known about these classes in terms of complexity. The most important class in quantum complexity theory is the class BQP , which consists of those problems which can be solved by a quantum machine with error probability bounded below half; This is considered as the class of tractable problems on a quantum computer. It is to be compared with the class BPP which is the same class for classical computers. Of course, we have that $BPP \subseteq BQP$, and that $BQP \subseteq QMA$. But how powerful is the class QMA ? Can we upperbound it? Adleman *et. al.* proved that BQP is contained in a large class, called PP . A language L is in PP if there exists a Turing machine that runs in polynomial time on an input x , and such that if $x \in L$ it outputs 1 with probability larger than $1/2$, and if $x \notin L$ it outputs 0 with probability larger than $1/2$. Note that the difference between the output probability and $1/2$ can be exponentially small. This makes the class possibly much stronger than the class BPP ; In particular, PP contains NP (see [5] lecture 7). It turns out that the above upper bound on BQP can be generalized to prove the same inclusions for the class QMA , i.e. $QMA \subseteq PP$. This fact was first noted by Kitaev and Watrous[17] who build on a simplification of [1] by Fortnow and Rogers[6] to prove it. To summarize we have that:

Theorem 3 $BPP \subseteq BQP \subseteq QCMA \subseteq QMA \subseteq PP$.

This is almost all that is known regarding the relation of BQP and QMA to classical complexity classes. To give intuition about what this upper bound means regarding the quantum complexity power, we note that the class PP is known to be contained in perhaps a more natural class, $PSPACE$, which is the class of languages that can be recognized by a Turing machine that uses polynomial *space* (but can take exponential amount of time.)

We now proceed to define the complete problem for QMA : Local Hamiltonian.

3 The Local Hamiltonian Problem

In this section we will define what can be thought of as the quantum analog of 3-SAT, called the “local Hamiltonian problem”.

Definition 1 5-Local Hamiltonian problem

- **Input:** H_1, \dots, H_r , A set of r Hermitian positive semi definite matrices operating on the space of five qubits, $\mathcal{B}^{\otimes 5}$, with bounded norm $\|H_i\| \leq 1$. Each matrix comes with a specification of the 5 qubits (out of the total n qubits) on which it operates. Each matrix entry is given with $\text{poly}(n)$ many bits. Apart from H_i we are also given two real numbers, a and b (again, with polynomially many bits) such that $b - a > 1/\text{poly}(n)$.
- **Output:** Is the smallest eigenvalue of $H = H_1 + H_2 + \dots + H_r$ smaller than a or are all eigenvalues larger than b ?

We slightly abuse notation here by writing $H = H_1 + H_2 + \dots + H_r$; H_i are matrices operating on different qubits, and the summation is over their extension to the entire set of qubits (tensor product with identity). This abuse of notation will be used throughout the paper, and it will be clear that we mean the summation of the operators as operators on n qubits.

Note that the defined problem is a promise problem: we are promised that one of the two possible outputs occurs. In other words, we don’t care what the output is for Hamiltonians with minimal eigenvalue between a and b .

In the same way, one can naturally define the k -local Hamiltonian problem for any k . We will see that 5-local Hamiltonian is QMA complete, and the reason for the number five will only be apparent towards the very end of the proof. However it is unclear whether it is necessary to consider 5 local Hamiltonian or whether a smaller number suffices; See open question number 1 for further discussion.

3.1 Connection to 3-SAT

We now show that the local Hamiltonian problem is a natural generalization of 3-SAT to the quantum world. For this, we explain how 3-SAT can be viewed as a 3-local Hamiltonian problem. We will work with qubits, but all operations are now classical operations in disguise. Let $\phi = C_1 \wedge C_2 \wedge \dots \wedge C_r$ be a 3-SAT formula on n variables, where each C_i is a clause, i.e. an OR over three variables or their negations. For every clause C_i we define a 8×8 matrix

H_i , operating on three qubits. H_i is a projection on the unsatisfying assignment of C_i . For example, for the clause $C_i = X_1 \vee X_2 \vee \neg X_3$ we get the matrix:

$$H_i = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} = |001\rangle\langle 001|$$

since 001 is the only unsatisfying assignment for C_i . H_i defined this way is a projection matrix. Moreover, it is Hermitian. If we look at H_i , its eigenvectors are all basis vectors of three qubits, with the vectors corresponding to satisfying assignments having eigenvalues 0, and the vector of the unsatisfying assignment corresponding to the eigenvalue 1. We then consider the operation of H_i on all the qubits, by taking tensor product of H_i with identity on the rest of the qubits. We denote the new matrix by H_i too, again by slight abuse of notation; It will be clear from context which of these we are talking about. If z is an assignment to the n variables which satisfies a clause C_i , then $H_i|z\rangle = 0$. Otherwise, $H_i|z\rangle = |z\rangle$. We can view this as if the matrix H_i “penalizes” assignments that do not satisfy C_i by giving them one unit of “energy”. We denote $H = \sum_{i=1}^r H_i$, and observe that $H|z\rangle = q|z\rangle$ where q is the number of clauses unsatisfied by z . All eigenvalues of H are non negative integers, and zero is an eigenvalue of H if and only if H corresponds to a satisfiable formula. Otherwise, the smallest eigenvalue of H is at least 1. Thus, 3-SAT is equivalent to the following problem: “Is the smallest eigenvalue of H 0 or is it at least 1?”, which is an instance of the 3-local Hamiltonian problem.

4 Local Hamiltonians is in QMA

Theorem 1 *The k -Local Hamiltonian problem is in QMA for any $k = O(\log(n))$.*

Proof: We first want to show that if the Hamiltonian H has an eigenvalue smaller than a , i.e. if we are in a “yes” instance, then there exists a witness that Marlin can use to convince Arthur for this fact. The obvious witness to use, is simply an eigenstate with eigenvalue smaller than a . Let us denote this ground state by $|\eta\rangle$ and its corresponding eigenvalue by λ . We will construct a procedure which outputs 1 with probability which is related to this eigenvalue. To illustrate the idea, consider first the simpler case in which all the Hamiltonians H_i are merely projections, $H_i = |\alpha_i\rangle\langle\alpha_i|$. In this case, we note that

$$\lambda = \langle\eta|H|\eta\rangle = \sum_{i=1}^r \langle\eta|H_i|\eta\rangle = \sum_{i=1}^r \langle\eta|\alpha_i\rangle\langle\alpha_i|\eta\rangle = \sum_{i=1}^r |\langle\eta|\alpha_i\rangle|^2 \quad (1)$$

or,

$$\lambda/r = (1/r) \sum_{i=1}^r |\langle\eta|\alpha_i\rangle|^2 \quad (2)$$

We note that $|\langle \eta | \alpha_i \rangle|^2$ is exactly the probability to get a positive answer when measuring the state $|\eta\rangle$ in the basis $|\alpha_i\rangle$ and the subspace orthogonal to it. Thus, equation 2 gives the following interpretation of λ/r : It is simply the probability to get the answer 1 when we pick i randomly between 1 and r and measure $|\eta\rangle$ in the basis $|\alpha_i\rangle$ and the subspace orthogonal to it. This implies an easy way to design an experiment, or a quantum verification procedure on the input state η , which outputs 1 with probability $1 - \lambda/r$ and 0 otherwise. Pick an $i \in \{1, \dots, r\}$ uniformly at random, measure $|\eta\rangle$ in the basis $|\alpha_i\rangle$ and the orthogonal subspace, and output 0 if the measurement resulted in a projection on $|\alpha_i\rangle$; output 1 otherwise. The probability for 1 is exactly $1 - \lambda/r \geq 1 - a/r$. On the other hand, if H is a “no” instance, i.e. all eigenvalues are larger than b , then for any vector $|\eta\rangle$,

$$\langle \eta | H | \eta \rangle = \sum_{i=1}^r \langle \eta | H_i | \eta \rangle \geq b; \quad (3)$$

The probability for 1 in the experiment is in this case

$$1 - \langle \eta | H | \eta \rangle / r \leq 1 - b/r; \quad (4)$$

Since we know that $b - a \geq 1/n^g$, we also have that the probabilities for 1 for the “yes” and “no” instances are polynomially different: $1 - a/r > 1 - b/r + 1/n^g$. We can amplify this difference using the amplification theorem and this proves that the problem is indeed in QMA, if the Hamiltonians are simple projections.

We remark that since the projections are local, i.e. involve at most $\log(n)$ qubits, such a measurement can be performed by a polynomial quantum verifier.

To deal with the more general case, where H_i are general Hermitian positive semidefinite matrices with norm at most 1, we note that any such matrix can be written in its spectral decomposition,

$$H_i = \sum_{j=1}^{\dim(H_i)} w_j^i |\alpha_j^i\rangle \langle \alpha_j^i| \quad (5)$$

We now impose the following trick which enables us to toss a coin with probability $1 - \langle \eta | H_i | \eta \rangle$. We first add one qubit to the system, in the state $|0\rangle$. We then apply the following unitary transformation on the qubits of H_i and on the extra qubit:

$$T |\alpha_j^i\rangle |0\rangle = |\alpha_j^i\rangle (\sqrt{w_j^i} |0\rangle + \sqrt{1 - w_j^i} |1\rangle) \quad (6)$$

We now prove that the measurement of the extra qubit outputs 1 with probability $1 - \langle \eta | H_i | \eta \rangle$. To see this, write

$$|\eta\rangle = \sum_j y_j |\alpha_j^i\rangle |\beta_j^i\rangle \quad (7)$$

using the Schmidt decomposition. After the transformation T , this state evolves to

$$T |\eta\rangle = \sum_j y_j |\alpha_j^i\rangle |\beta_j^i\rangle (\sqrt{w_j^i} |0\rangle + \sqrt{1 - w_j^i} |1\rangle) \quad (8)$$

The probability to measure 1 is then the squared norm of the following vector:

$$\sum_j \sqrt{1 - w_j^i} y_j |\alpha_j^i\rangle |\beta_j^i\rangle. \quad (9)$$

This squared norm is just

$$\sum_j (1 - w_j^i) |y_j|^2 = 1 - \sum_j w_j^i |y_j|^2 \quad (10)$$

but we know that

$$\langle \eta | H_i | \eta \rangle = \sum_j w_j^i |y_j|^2. \quad (11)$$

We can now describe the exact verification procedure: Pick a random index i , and perform the above test for H_i : Add one qubit, apply T and measure the extra qubit. The outcome will be 1 with probability

$$\sum_i (1/r) (1 - \langle \eta | H_i | \eta \rangle) = 1 - \langle \eta | H | \eta \rangle / r \quad (12)$$

If we are in a “yes” instance, this number will be larger than $1 - a/r$; If we are in a “no” instance, it will be smaller than $1 - b/r$, and the proof is completed just as in the simple projections case. \square

5 QMA Completeness

In this section we will show that the 5-local Hamiltonian problem is QMA-hard. The proof is complicated, and we will start with an overview.

5.1 Reminder of the Cook-Levin Proof

The proof that local Hamiltonian is QMA complete bears a lot of resemblance to Cook-Levin’s proof that 3SAT is NP-Complete. Let us briefly sketch the idea underlying the Cook-Levin’s proof, so that we can refer to it later on. Consider an NP problem, L . There is a Turing machine which operates on x, y where x is a supposedly member of L and y is a supposed witness for this fact, and M checks that x is in the language using y . We now want to construct a reduction to 3-SAT, i.e. to design a Boolean formula which is satisfiable if and only if x is in the language, i.e. if the Turing machine performed a successful computation which started with the input and ended with “accept” or 1, in the first site on the tape. To construct such a formula, we consider the variables $x_{i,t}$ where i runs over all reachable locations on the tape in the polynomial time limit and t runs over the time steps. $x_{i,t}$ are variables which can get any of some constant number of possible values; These values correspond to a finite description of the state of the Turing machine related to the location i on the tape at time t . They include what is written on the tape at that time and that location, the state of the

Turing machine at that time, and whether the head of the Turing machine is at that location or not. An assignment to these variables can be viewed as a *history* of some computation; a description of how the Turing machine evolved in time. The 3-SAT formula we construct is essentially checking that this evolution is a valid evolution of the Turing machine. Each clause in the formula will look at three subsequent cells at some time t , say $x_{i-1,t}$, $x_{i,t}$ and $x_{i+1,t}$ plus the cell $x_{i,t+1}$. Given the values of $x_{i-1,t}$, $x_{i,t}$ and $x_{i+1,t}$, it is possible to know whether the value of $x_{i,t+1}$ is valid or not; Thus, the clause is satisfied if and only if $x_{i,t+1}$ evolves from $x_{i-1,t}$, $x_{i,t}$ and $x_{i+1,t}$ by a valid computation. We also add clauses that check that the input is really x , i.e. clauses of the form $x_{i,0}$ if the i 'th input bit was 1, and $\neg x_{i,0}$ if the i 'th input bit was 0. Finally, we check that the output is accept by adding the clause $x_{1,T}$, which is satisfied if the first site is 1 at the end of the computation at time T . Each of these verifications is local, since the evolution of the Turing machine is local, and thus each corresponds to a clause. Note that our variables have a constant but possibly large set of possible values; It is easy to see that such formulas can be converted to formulas over Boolean variables, and that each clause can be converted to many clauses each operating only on three variables. All these details are none of our concern; The main issue, which we will try to mimic in the quantum case, is that the history of a Turing machine can be verified locally.

5.2 The Quantum analog- Sketch

The idea of the quantum proof is very similar. We know that L is in QMA ; Thus, there exists a quantum circuit, using two-qubit gates, which accepts an input x with some witness $|\xi\rangle$ with high probability (we will assume it is exponentially close to 1) if x is in L and rejects with exponentially close to one probability if $x \notin L$ given any witness. We want to reduce this problem to the local Hamiltonian problem, i.e. to construct a Hamiltonian which will have small eigenvalue in the $x \in L$ case and only large eigenvalues otherwise.

How to construct the analog? Drawing from the Cook-Levin proof, we want the history of the computation to be our witness, which we hope to be able to verify locally. Our first guess for the quantum witness would thus be the sequence of states which constitute the history of the computation:

$$|x\rangle|\xi\rangle, U_1|x\rangle|\xi\rangle, U_2U_1|x\rangle|\xi\rangle, \dots, U_T \cdots U_2U_1|x\rangle|\xi\rangle. \quad (13)$$

However, there is a serious problem with this suggestion. Let us assume for a moment that U_1 is simply the identity gate, and all we want to check is whether the first and second states given to us by the prover are the same, and we want to do this via a local Hamiltonian. In general, we want to design a local Hamiltonian which when applied on $|\alpha\rangle|\alpha\rangle$ it behaves differently than when applied on $|\alpha\rangle|\beta\rangle$, if $|\beta\rangle$ is quite different from $|\alpha\rangle$. The problem is that a local Hamiltonian has access only to the reduced density matrix of the state $|\alpha\rangle|\beta\rangle$ to five qubits at a time. In other words, $\langle \eta | H | \eta \rangle$ for any state $|\eta\rangle$ will be exactly the same if we move to $|\eta'\rangle$ as long as it has the same reduced density matrices as $|\eta\rangle$ on all sets of five qubits. It is very easy to construct two states which agree on all density matrices of five qubits, but are completely different due to their overall correlations or entanglement. Hence, using only local Hamiltonians we cannot hope to be able to verify the correctness of the time evolution if the states are given to us sequentially. However, entanglement which was the source of this

problem, can also help us solve it. Consider the following superposition

$$\frac{1}{\sqrt{2}}(|0\rangle|\alpha\rangle + |1\rangle|\beta\rangle) \tag{14}$$

From the reduced density matrix of just the *first* qubit, we can learn a lot about whether the states $|\alpha\rangle$ and $|\beta\rangle$ are the same or different; in fact, the reduced density matrix of the first qubit tells us the angle between these two states, as one can easily verify. This means that if the histories are given to us in *superposition*, there is hope that local measurements or observables like our local Hamiltonian will be able to verify the correctness of the time evolution.

The idea is therefore to ask the prover for the history of the computation, not in the form of sequential states but rather in a superposition over all *time leafs*:

$$|\eta\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T U_t \dots U_1 |x, \xi\rangle |t\rangle \tag{15}$$

we will see later how this state can actually be verified for correctness. In the book[9] this idea of moving from time evolution to a time-independent local Hamiltonian is attributed to Feynman[4].

Except for this main difference of using superposition over time instead of sequential time, there is another essential difference in the proof. In the classical case the eigenvalues are integers, and so to show soundness one only has to show that the resulting formula is not satisfiable if x is not in the language. The corresponding statement would be that the smallest eigenvalue of H is not 0; In the classical case, this automatically means that it is at least 1. In the quantum case, due to the continuous nature of the model, the fact that the smallest eigenvalue is larger than 0 is not enough; One actually has to show that it is at least polynomially bounded away from zero, because the accuracy achieved by the verification process is only polynomial, i.e. we can only amplify a polynomial separation and not an exponentially small separation. To bound the lowest eigenvalue from below Kitaev uses a geometrical argument, augmented with some nice ideas of how to perform the analysis involving the known theory of random walks on the line, represented here by the time axis.

5.3 The reduction

Let L be a problem in QMA. Then there exists a quantum circuit Q with two-qubit gates U_1, \dots, U_T such that for an input $|x\rangle$ and a witness $|\xi\rangle$ the output qubit has more than $1 - e^{-n}$ probability to collapse on $|1\rangle$ if $x \in L$ and less than e^{-n} probability to collapse the first qubit on $|0\rangle$ otherwise. Given this sequence of gates, we will construct an input to the local Hamiltonian problem, i.e. a sequence of local matrices. For now, our matrices will not be completely local, but instead will operate on two qubits among the n computer qubits plus an extra $T + 1$ dimensional Hilbert space, which will serve as a clock, and which is augmented to the right of all the other qubits. We will modify the Hamiltonian later on so it is truly 5-local, but for the sake of simplicity we present the main part of the proof using this extra $T + 1$ dimensional Hilbert space. We denote by the subscript C the subspace of the clock. The subscript i at the

foot of operators means that they operate on qubit number i ; The projection operator $\Pi^{|\alpha\rangle}$ means project on the subspace spanned by $|\alpha\rangle$. Our Hamiltonian will be a sum of three main terms, $H = H_{in} + H_{out} + H_{prop}$.

- H_{in} is a matrix that checks that the input for the first n qubits is indeed x , where we do not care about the witness; It can be anything. This check need to verify that the i th bit is indeed x_i , at time 0, for all i between 1 to n . This is done by projecting the state to time 0 (by projecting the clock state to time 0), and then projecting the remaining state to the space orthogonal to $|x_i\rangle$:

$$H_{in} = \sum_{i=1}^n \Pi_i^{|\neg x_i\rangle} \otimes |0\rangle\langle 0|_C \quad (16)$$

- H_{out} is a matrix that checks that the output is 1 at time T , again, by first projecting the clock to time T and then projecting the state to the subspace orthogonal to $|1\rangle$ on the first qubit which carries the answer of the quantum circuit:

$$H_{out} = \Pi_1^{(0)} \otimes |T\rangle\langle T|_C \quad (17)$$

- H_{prop} checks that the propagation of the computational process is done according to the given circuit. It is a sum of T terms,

$$H_{prop} = \sum_{t=1}^T H_{prop}(t) \quad (18)$$

where each term checks that the propagation from time $t - 1$ to t is correct:

$$H_{prop}(t) = \frac{1}{2}(I \otimes |t\rangle\langle t| + I \otimes |t-1\rangle\langle t-1| - U_t \otimes |t\rangle\langle t-1| - U_t^\dagger \otimes |t-1\rangle\langle t|) \quad (19)$$

During the proof of the completeness part, it will become clear why each of these terms really verifies what we claim it does. Note that each term in the above Hamiltonian indeed satisfies our constraints of being Hermitian, positive semi-definite and of norm at most 1; There is one problem, which is that it is only *log*-local and not local, since it operates on two qubits among the main qubits plus the clock space (which can be represented by logarithmically many qubits, which is the reason why we call it *log*-local.) We will fix this problem only much later and for now we work with the Hamiltonian H as defined. To complete the reduction we also need to specify a and b ; We let $a = 1/T^{10}$, and $b = 1/4(T + 1)^3$.

The claim is that the constructed Hamiltonian has an eigenvalue less than a if x is in the language that the quantum circuit accepts, and otherwise all the eigenvalues of the Hamiltonian are larger than b . Once we prove both claims (completeness and soundness) we will be done; The two together imply that solving the local Hamiltonian problem for the Hamiltonian that is associated with a certain circuit, is a way to decide the answer of the circuit, (i.e. solving the Hamiltonian problem is QMA hard: any QMA problem can be solved using a machine that solves the local Hamiltonian problem.) It will remain only to deal with the locality problem which we will do at the very end.

5.4 Completeness

To prove completeness, we want to show that a “yes” instance of the QMA problem transforms to a “yes” instance in the Local Hamiltonian problem. If $x \in L$, the H we constructed has an eigenvalue smaller than a . For this, it suffices to prove the following claim:

Claim 2 *If x is accepted by the circuit Q , for some quantum witness $|\xi\rangle$, with probability which is larger than $1 - \epsilon$, then the Hamiltonian $H = H_{in} + H_{prop} + H_{out}$ constructed above given x and the circuit Q has an eigenvector with eigenvalue $\leq \epsilon$.*

Proof: To see why the claim is true, in analogy with the classical case, the state we will use is the history of the computation

$$|\eta\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T U_t U_{t-1} \dots U_1 |\gamma_0\rangle \otimes |t\rangle \quad (20)$$

where $|\gamma_0\rangle$ is the state at the beginning of the computation (a tensor product of the input and the witness to the machine) and $|t\rangle$ is a clock state.

The intuition is that this state is “almost” a zero eigenstate of the Hamiltonian H , since it “almost” satisfies all the tests this local Hamiltonian checks. More formally, we claim that

$$\langle \eta | H | \eta \rangle \leq \epsilon. \quad (21)$$

which suffices to prove the claim.

To calculate $\langle \eta | H | \eta \rangle$ we first note that

$$H_{in} |\eta\rangle = 0. \quad (22)$$

It is less obvious but can be easily checked that for each $t = 1, \dots, T$

$$H_{prop}(t) |\eta\rangle = 0. \quad (23)$$

The reader is recommended to verify this step, since it explains the definition of the propagation Hamiltonian, which is one of the main conceptual steps in the proof. The intuition is that the propagation Hamiltonian is composed of four parts, all confined to the projections on the span of the two time leaves $|t-1\rangle$ and $|t\rangle$. Two terms in the Hamiltonian $H_{prop}(t)$, $I \otimes |t\rangle\langle t| + I \otimes |t-1\rangle\langle t-1|$ correspond simply to picking out the state at those times. In addition, there are two extra terms: the term, $U_t \otimes |t\rangle\langle t-1|$ which corresponds to a forward propagation in time, and a term $U_t^\dagger \otimes |t-1\rangle\langle t|$ which corresponds to backwards propagation in time; When applied on the projection of the state to the two time steps $|t-1\rangle$ and $|t\rangle$, the forward propagation in time term picks just the $t-1$ time step and propagates it forward by applying U_t to it, and then the resulting state gets canceled with the t time step; The same happens with the backwards propagation term which picks up the time step t , propagates it one step backwards by applying U_t^\dagger and then this term gets canceled with it. with the $t-1$ time step.

It is left to check what happens to $|\eta\rangle$ when we apply H_{out} . When we apply H_{out} on $|\eta\rangle$ we get a projection on the part of $|\eta\rangle$ which rejects. Since the probability for rejection is $\leq \epsilon$, we get that the norm squared of $H_{out}|\eta\rangle$ is at most ϵ , and hence $\langle \eta | H_{out} | \eta \rangle = \|H_{out}|\eta\rangle\|^2 \leq \epsilon$. Hence, the minimal eigenvalue of H is less than ϵ . \square

5.5 Soundness

To complete the reduction, we need to show that if $x \notin L$, the minimal eigenvalue of H is larger than the chosen b .

Theorem 2 *If $x \notin L$ then the minimal eigenvalue of H is $\geq \frac{1}{4(T+1)^3}$.*

Proof: To prove this theorem we will put together several lemmas. The idea is to write H as a sum of two Hamiltonians, $H_1 = H_{in} + H_{out}$, $H_2 = H_{prop}$, and to use the following geometrical lemma, which gives a lower bound on the lowest eigenvalue of a sum of two Hamiltonians, given some conditions on the eigenvalues and eigenspaces of the two Hamiltonians.

Lemma 1 *Let H_1 and H_2 be two Hermitian positive semi-definite matrices, and let N_1 and N_2 be the eigenspaces of the eigenvalue 0, respectively. If the angle between N_1 and N_2 is some $\theta > 0$, and the second eigenvalue of both H_1 and H_2 is $\geq \lambda$ then the minimal eigenvalue of $H_1 + H_2 \geq \lambda \sin^2(\theta/2)$.*

Proof: Consider an eigenvector of $H_1 + H_2$, $|\delta\rangle$ such that $\| |\delta\rangle \| = 1$. For at least one of the subspaces N_1 or N_2 The angle between $|\delta\rangle$ and this subspace is at least $\frac{\theta}{2}$. W.L.O.G let this subspace be N_1 . We have

$$\langle \delta | (H_1 + H_2) | \delta \rangle = \langle \delta | H_1 | \delta \rangle + \langle \delta | H_2 | \delta \rangle \geq \langle \delta | H_1 | \delta \rangle.$$

We write

$$|\delta\rangle = |\mu\rangle + |\mu^\perp\rangle$$

where $|\mu\rangle, |\mu^\perp\rangle$ are the projections of $|\delta\rangle$ onto N_1 and the orthogonal subspace to N_1 respectively. Then

$$\langle \delta | H_1 | \delta \rangle = \langle \mu^\perp | H_1 | \mu^\perp \rangle \geq \| |\mu^\perp\rangle \|^2 \lambda$$

where the first equality follows from the fact that N_1 and its complement are invariant to the application of H_1 and the second follows from the definition of H_1 and λ . We also know that $\| |\mu^\perp\rangle \|^2 \geq \sin^2(\theta/2)$ because the angle between N_1 and $|\delta\rangle$ is at least $\theta/2$, and this completes the proof. \square

To use the geometrical lemma, we will assume $x \notin L$ and give lower bounds on the second eigenvalues of H_1 and H_2 , as well as a lower bound on θ . We will first bound the second eigenvalues of H_1 and H_2 .

Lemma 2 *The second eigenvalue of H_1 is at least 1.*

Proof: The second eigenvalue of H_1 is ≥ 1 since H_{in} and H_{out} are projections and hence their eigenvalues are 0 and 1. Since the eigenspaces of the eigenvalue 1 of H_{in} and H_{out} are orthogonal (because they operate on different times), they commute, and so their second eigenvalue is simply the minimal second eigenvalue of the two. \square

Lemma 3 *The second eigenvalue of $H_2 = H_{prop}$ is at least $\frac{1}{2(T+1)^2}$.*

Proof: It turns out that for this and further arguments it is simpler to look at H_{prop} in a rotated basis. The eigenvalues of a matrix are not changed when looked at in a different basis. Hence we define the unitary matrix R as follows:

$$R = \sum_{t=0}^T U_t \dots U_1 \otimes |t\rangle\langle t|. \quad (24)$$

R is unitary since it is a block diagonal matrix with each of its blocks unitary. What R does is basically rotate the basis in each time leaf to the basis which one gets if one applies the first t computation steps on the computational basis. Hence, in the new rotated basis, the computation is simply the identity. Now, it is easy to check that

$$R^\dagger H_{prop}(t)R = \frac{1}{2}(I \otimes |t\rangle\langle t| + I \otimes |t-1\rangle\langle t-1| - I \otimes |t-1\rangle\langle t| - I \otimes |t-1\rangle\langle t|) \quad (25)$$

We can write $H_{prop} = I \otimes A$ where A is a $(T+1) \times (T+1)$ of the form:

$$A = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\ -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & 0 & 0 & 0 & 0 \\ 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & 0 & 0 \\ 0 & 0 & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

$$= I - \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix} = I - B$$

The eigenvalues of $R^\dagger H_{prop}R$ or equivalently of H_{prop} are simply the eigenvalues of A (with multiple appearances), or 1 minus those of B ; It suffices then to find the eigenvalues of B .

Interestingly, the matrix B is a familiar matrix from the theory of random walks and we will use this fact in the analysis of its eigenvalues. For a direct proof see Kiteav[10]. Here we refer to the theory of random walks due to its intriguing connection with the subject at hand. For a nice exposition of random walks, see Lovasz's survey[12]. Returning to our matrix B , it turns out that it is the stochastic matrix corresponding to a simple random walk on the time axis, from 0 to T . with a loop at both ends. The largest eigenvalue of this matrix is 1, corresponding to the eigenvector which is the uniform limiting distribution. This eigenvalue gives the 0 eigenvalue of A and hence of H_{prop} . In random walk theory, one is very interested

in the second eigenvalue of the stochastic matrices corresponding to random walks since the second eigenvalue is directly related to the rate at which the random walk mixes to its limiting distribution. B 's second largest eigenvalue λ_2 is bounded from below by the conductance ϕ of the graph on which the random walk is applied, using Jerrum and Sinclair's bound[15]:

$$1 - \lambda_2 \geq \phi^2/2 \quad (26)$$

The conductance of the random walk is $\frac{1}{T+1}$ which gives $1 - \lambda_2 \geq \frac{1}{2(T+1)^2}$. Since 1 minus the second largest eigenvalue of B is exactly the second smallest eigenvalue of A , this implies the desired result. \square

It is left to give a lower bound on the angle between the two null spaces.

Lemma 4 *The angle between N_1 and N_2 satisfies $\sin^2(\theta/2) \geq \frac{1}{2(T+1)}$.*

Proof: $H_1 = H_{in} + H_{out}$ is a projection, and hence the null space is simply the subspace orthogonal to the space on which H_1 projects. Hence, N_1 is equal to the direct sum of three subspaces:

$$N_1 = (|x\rangle\langle x| \otimes W \otimes |0\rangle\langle 0|) \oplus (|1\rangle\langle 1| \otimes W \otimes |T\rangle\langle T|) \oplus_{t=1}^{T-1} (W \otimes |t\rangle\langle t|) \quad (27)$$

where W is the entire Hilbert space for the remaining of the qubits. N_2 , the null space of H_{prop} , is exactly the space spanned by all valid computations starting with an arbitrary state $|\alpha\rangle$ on the qubits of the input and witness together. These are all states of the form:

$$|\eta\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T U_t \dots U_1 |\alpha\rangle \otimes |t\rangle. \quad (28)$$

The fact that such states are in the null space of H_{prop} was shown before; The fact that all states in the null space of H_{prop} are of this form follows from looking at the rotated $R^\dagger H_{prop} R$, as before. The null space of the rotated H_{prop} is simply the entire space on the computer register times the null space of the clock matrix A ; The null space of the matrix A is exactly all constant vectors. This is a standard claim, following from the fact that the random walk B defines on the line is aperiodic, ergodic, and converges to the uniform vector. (One can readily prove this fact also from scratch, by considering the effect of A on the eigenvector corresponding to eigenvalue 1, and looking at the maximal coordinate.)

Now, to find N_2 , the null space of H_{prop} , we have to rotate the null space of $RH_{prop}R^\dagger$ (which is all the Hilbert space on the computer qubits times constant vectors on the clock space) back to the original basis, by applying R and R^\dagger from both sides. It is easy to see that for any state

$$|\alpha\rangle \otimes \frac{1}{\sqrt{T+1}} \sum_{t=0}^T |t\rangle \quad (29)$$

rotating it back gives a state of the form of equation 28.

We now want to bound the angle between N_1 and N_2 , which is the minimal angle between two vectors from both spaces. Any vector in N_2 is of the form of equation 28, i.e. a history of a certain computation, and the angle ϕ between such a history $|\eta\rangle$ and N_1 is given by

$$\cos^2(\phi) = \|\Pi_{N_1}|\eta\rangle\|^2 \quad (30)$$

where Π_{N_1} denotes the projection onto N_1 . This is true since $|\eta\rangle$ is of norm 1. Thus we have that the angle θ between N_1 and N_2 is the minimal angle ϕ between a history vector and the space N_1 , or equivalently:

$$\cos^2(\theta) = \max_{|\eta\rangle \in N_2} \{\|\Pi_{N_1}|\eta\rangle\|^2\} \quad (31)$$

We now claim that for any $|\eta\rangle \in N_2$ we have:

$$\|\Pi_{N_1}|\eta\rangle\|^2 \leq 1 - \frac{1}{2(T+1)}. \quad (32)$$

The proof of this will complete the proof of the lemma, using equation 31. To prove the upper bound of equation 32, we observe that the norm squared of the projection onto N_1 is simply the sum of the norms squared on the projections on the different parts of N_1 , as a direct sum of subspaces. If we write N_1 as a direct sum of the different spaces spanned by times $t = 1, \dots, T-1$, then since $|\eta\rangle$ is the uniform superposition over time, the projection of $|\eta\rangle$ on each of the middle time step gives $\frac{1}{T+1}$, and so the total contribution of the middle time leafs is $\frac{T-1}{T+1}$.

We now claim that the contribution of the first and last leafs together is far from the maximal possible contribution $\frac{2}{T+1}$. Intuitively, this is due to the fact that the projection on N_1 sums up the projection on x as input in the beginning of the computation plus the projection on “accept” at the end of the computation. However, since $|\eta\rangle$ represents a valid computation by a circuit that does not accept x , it cannot be the case that both projections are maximal. To quantify this statement, we observe that we can write $|\eta\rangle$ as a sum of two states:

$$|\eta\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T |\gamma_t\rangle \otimes |t\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T (a|\gamma_t^1\rangle + b|\gamma_t^2\rangle) \otimes |t\rangle = a|\eta_1\rangle + b|\eta_2\rangle \quad (33)$$

where $|\gamma_0^1\rangle$ is the normalized projection of $|\gamma_0\rangle$ on the input being x , and $|\gamma_0^2\rangle$ is the normalized projection on the orthogonal subspace, and $|\gamma_t^1\rangle, |\gamma_t^2\rangle$ are simply the states obtained from the initial states by applying the computation. The norm squared of the projection of the first time leaf of $|\eta\rangle$, $\frac{1}{\sqrt{T+1}}|\gamma_0\rangle \otimes |0\rangle$, onto N_1 is $\frac{a^2}{T+1}$. The norm squared of the projection of the last time leaf of $|\eta\rangle$, $\frac{1}{\sqrt{T+1}}|\gamma_0\rangle \otimes |T\rangle$, onto N_1 is the norm squared of

$$\frac{1}{\sqrt{T+1}}(a|\delta_T^1\rangle + b|\delta_T^2\rangle) \quad (34)$$

where $|\delta_T^1\rangle, |\delta_T^2\rangle$ are the projections of $|\gamma_T^1\rangle, |\gamma_T^2\rangle$ on “accept”, respectively (we are using the linearity of projection.) Now we have that

$$\| |\delta_T^1\rangle \|^2 \leq e^{-n} \quad (35)$$

since the circuit accepts with probability less than e^{-n} if x is not in the language. Hence,

$$\|a|\delta_T^1\rangle + b|\delta_T^2\rangle\| \leq 2e^{-n} + b^2 \quad (36)$$

and so the total norm squared of the projection on N_1 is at most

$$\|\Pi_{N_1}|\eta\rangle\|^2 \leq \frac{1}{T+1}(a^2 + 2e^{-n} + b^2) + \frac{T-1}{T+1} \leq (1 - \frac{1}{2(T+1)}) \quad (37)$$

using the fact that $a^2 + b^2 = 1$. This completes the proof. \square

To complete the proof of the theorem 2, we simply apply the geometrical lemma 1 using the bounds we have shown for the eigenvalues and for θ . \square

This completes the proof of the hardness of Local Hamiltonian for QMA, if we are allowed to use Hamiltonians which operate on spaces of polynomial dimension; In the next section we make the last step that is needed to convert the Hamiltonian to a 5-local Hamiltonian consisting of terms operating on five qubits only.

6 Improving from *log* local to 5-local

To move from operators on the entire clock to local operators, represent the time in unary representation on T qubits which will serve as the clock qubits. For example, time $t = 4$ is represented by the T qubit state $|111100\dots 00\rangle$. To modify the Hamiltonian accordingly, we replace all operators on the clock space by operators that operate on three qubits at most. We apply the following modifications:

$$\begin{aligned} |t\rangle\langle t-1| &\longmapsto |110\rangle\langle 100| \otimes I \\ |t-1\rangle\langle t| &\longmapsto |100\rangle\langle 110| \otimes I \\ |t\rangle\langle t| &\longmapsto |110\rangle\langle 110| \otimes I \\ |t-1\rangle\langle t-1| &\longmapsto |100\rangle\langle 100| \otimes I \end{aligned} \quad (38)$$

where in all these cases $|110\rangle\langle 100|$ or the similar terms operate on qubits $t-1, t, t+1$ of the clock qubits and the identity I operates on the remaining $T-2$ clock qubits. This will hold in all terms of H_{prop} , except for two exceptions to the above - when $t = 1$ and $t = T$, in order not to refer to bits 0 and $T+1$ of the clock which do not exist. For $t = 1$ we will drop the first bit of the 3-bit operator, so the operator $|1\rangle\langle 0|$ on the original clock becomes $|10\rangle\langle 00|$ on the first two bits; and similarly $|0\rangle\langle 1|$ on the original clock becomes $|00\rangle\langle 10|$ on the first two bits in the unary clock. For the case $t = T$ we drop the 3^{rd} bit of the operators in the same manner. The final H_{prop} which we get is

$$H'_{prop}(t) = \frac{1}{2}(I \otimes |110\rangle\langle 110| + I \otimes |100\rangle\langle 100| - U_t \otimes |110\rangle\langle 100| - U_t^\dagger \otimes |100\rangle\langle 110|) \quad (39)$$

where the three qubit operators operate on qubits $t-1, t, t+1$. For $H'_{prop}(1), H'_{prop}(T)$ we get a slightly different expression with the clock operators operating only on two qubits as explained above.

As for H_{in} and H_{out} , we again change $|t\rangle\langle t|$ to be an operator on three qubits for the middle time leafs and two qubits for the beginning and end leafs.

We first claim that restricting ourselves to the subspace spanned by states with the clock qubits being in valid unary representations, all previous claims hold. Explicitly, as can be easily checked:

Claim 3 *For any state $|\eta\rangle$ which represents a valid history of a computation, (in unary representation) we still get $H'_{prop}|\eta\rangle = 0$.*

From this, using exactly the same arguments as used before, we have that

Claim 4 *If $|\eta\rangle$ is a history of an accepting computation, then $\langle\eta|H'|\eta\rangle \leq \epsilon$.*

Hence, completeness will go through with these modifications. However, soundness will not go through because of the following reason. The T qubits that we have introduced have many more possible states except for valid unary representations of some time step. The H' we have defined operates on such states as well; To prove soundness, we need to show that among such states there are no states of small eigenvalue in case of x not in the language. This might be complicated, and we resort to a different solution.

In addition to the modifications of the existing terms in the Hamiltonian, we introduce a new term which penalizes the state of the clock qubits if they are not unary representation of some time step. We call this term H_{clock} ; It locally checks that the clock bits are a valid unary representation. All we need to check is that two consequent bits cannot be in the state $|01\rangle$; This can be done by a sum of local projections, as follows:

$$H'_{clock} = \sum_{t=1}^T |01\rangle\langle 01|_{t-1,t} \otimes I. \quad (40)$$

Our (truly!) final Hamiltonian is defined to be

$$H' = H'_{in} + H'_{out} + H'_{prop} + H'_{clock} \quad (41)$$

Clearly, $|\eta\rangle$ (with a unary clock) is an eigenvector of eigenvalue $\leq \epsilon$ of the new Hamiltonian H' , since it is a zero eigenvector of H'_{clock} , and so completeness is preserved.

For the proof of soundness, we observe that H' keeps the subspace that is spanned by all states in which the clock qubits are valid unary representations invariant; Let us call this subspace \mathcal{D} . The orthogonal subspace, \mathcal{D}^\perp , is also invariant under the operation of H' . H' operates on \mathcal{D} just as the previous H did, and hence on this subspace the lower bound on the eigenvalues holds as before; On the orthogonal subspace \mathcal{D}^\perp the eigenvalue of H' is at least 1 since H'_{clock} detects at least one violation. Hence, overall, the lower bound from theorem 2 holds here too. This completes the proof of completeness of 5-local Hamiltonian.

Remark 2 Why 5? *We remark here regarding the necessity of three qubits Hamiltonians instead of one qubit Hamiltonians to control the propagation in time. One can naively suggest to use the one qubit Hamiltonian $|1\rangle\langle 0|$ operating on the t^{th} clock qubit to represent the propagation*

from $t - 1$ to t , instead of the three qubit operator we use. This suggestion does not work for the following reason: the history states will no longer be eigenstates of H'_{prop} , since the one qubit time propagation terms might cause valid time leafs to propagate to invalid ones; E.g., $|11100000\rangle$ will propagate by $U_6 \otimes |1\rangle\langle 0|_6$ to $|11100100\rangle$. This does not happen when three qubit operators are used. It is an open question whether this obstacle can be overcome to show that 3-local or 4-local Hamiltonian is QMA complete; See open question 3.

7 Discussion and Open Questions

We have presented here a beautiful result by Kitaev which we believe is a fundamental stepping stone for the field of quantum complexity. We collect here a list of open questions it raises.

The first set of problems is related to the question of the expressiveness of the class QMA. There are hundreds of NP complete problems, from an enormous variety of fields; So far, the only interesting quantum MA complete problem we know of is the Local Hamiltonian problem.

Open Question 1 *Find more quantum MA complete problems.*

In particular,

Open Question 2 *Is there a natural QMA complete problem which is not quantum related?*

We have proved that 5-local Hamiltonian is QMA complete. What is the importance of the number five? It is unknown whether 5 qubits are necessary. Perhaps even 2-locality suffices to achieve QMA completeness. This is very different from the classical situation, where it is known that 3-SAT is NP complete but 2-SAT can be solved in polynomial time.

Open Question 3 *What is the complexity of k -local Hamiltonian with $k = 2, 3, 4$?*

The next question is related to the definition of the class QMA:

Open Question 4 *Is QMA with two sided errors the same as QMA with one sided error?*

This holds in the classical case, and the question is whether it holds quantumly. Kitaev and Watrous[11] show the equivalence of one and two sided errors in certain cases (quantum interactive proofs with more rounds) but the proof does not carry over to this case.

Another open question which is related to the definition of QMA is

Open Question 5 *Is $QCMA = QMA$?*

Due to the results presented in this survey, it seems reasonable to assume that the answer is yes. Our intuition behind this conjecture is that the quantum verifier limits itself in its tests of the quantum state to the reduced density matrices of five qubits; It therefore does not care about longer range entanglement. Perhaps such states that are specified by short range entanglement can be efficiently generated. In other words:

Open Question 6 Consider a (possibly very complicated) n qubit state $|\xi\rangle$. Is there an efficient circuit that generates a state $|\xi'\rangle$ which has (almost) the same reduced density matrices to any subset of five qubits?

If this can be done, this will prove the equality $QCMA = QMA$ since the classical witness can be the description of the quantum circuit, and the verifier can generate the state on its own. A proof that shows this equivalence is likely to be very insightful regarding quantum correlations. This question touches upon the interesting question of whether it is possible to develop a quantum analog of the beautiful theory of pseudo-random generators; In pseudo-random generation, one generates probability distributions that are very different from the uniform distribution but such that any circuit of some restricted set (say of bounded size) cannot tell the difference. The question we are asking is of a similar type, and can be viewed as a *pseudo quantum generator* type question, since we need the state to pass the test of the very restricted verifier who only looks at sets of five qubits.

The results presented here highlight a very interesting connection between Hamiltonians and unitary gates or quantum circuits, which Kitaev attributes to Feynman[4]. We view this connection as fascinating and potentially very powerful. In particular, in exactly the same way QMA circuits are translated to a local Hamiltonian, one can also translate BQP circuits to a local Hamiltonian; It is very interesting to ask whether the other direction of moving from groundstates of local Hamiltonians to efficient circuits that generate them also holds. We cannot hope for a constructive version of this direction since it is QMA hard, but an existence proof of such short circuits for ground states of local Hamiltonians will be very interesting, leading to a positive answer to open question 5 and to implications to the very important task of quantum state generation (see [2].)

Open Question 7 Given a local Hamiltonian, does there exist a polynomial size quantum circuit that generates (a state with non negligible projection on) its ground state?

We remark that Feynman's point of view[4] of moving from circuits to time independent Hamiltonians, enables one to translate short computation time into large spectral gap of Hamiltonians. The spectral gap of $\Omega(1/T^3)$ achieved in Kitaev's proof is indeed expressed in terms of the computational time T . It is interesting to compare this to the framework of adiabatic quantum computation[3] in which the opposite direction of moving from large spectral gaps to efficient state generation is taken, when one aims at generating a groundstate of a final Hamiltonian H_f by designing a sequence of local Hamiltonians H_0, \dots, H_f all with large spectral gaps.

We end with a more general open problem which seems related to the results presented here:

Open Question 8 Does the quantum analog of the PCP theorem hold? Can we prove hardness of approximation for quantum computation?

Hopefully, the results and open questions presented in this survey will provide easy access to research in what we view as a fascinating area.

8 Acknowledgements

One of us (D.A.) wishes to thank Avi Wigderson for insightful comments during Kitaev's first lecture about the subject[9].

References

- [1] L. Adleman, J. DeMarrais, and M. Huang. Quantum computability. *SIAM J. Computing* 26 (1997) 1524-1540.
- [2] Dorit Aharonov and Amnon Ta-Shma, Quantum sampling and statistical zero knowledge (tentative name), in preparation.
- [3] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Michael Sipser, Quantum Computation by Adiabatic Evolution, quant-ph/0001106
- [4] Richard. P. Feynman, Quantum Mechanical Computers, *Optic News*, **11**, February 1985, p. 11
- [5] O. Goldreich, Lecture notes for the course "Introduction to Complexity Theory", 1999. <http://www.wisdom.weizmann.ac.il/mathusers/oded/cc99.html>
- [6] L. Fortnow and J. Rogers. Complexity limitations on quantum computation. *Journal of Comput. and Syst. Sci.* 59(2) (1999), 240-252
- [7] Oded Goldreich and David Zuckerman. Another proof that $BPP \subseteq PH$ (and more). Electronic Colloquium on Computational Complexity Technical Report TR97-045, September 1997.
- [8] Iordanis Kerenidis and Ronald de Wolf, Exponential Lower Bound for 2-Query Locally Decodable Codes, quant-ph/0208062
- [9] A. Yu Kitaev, Lecture given in Hebrew University, Jerusalem, Israel, 1999.
- [10] A. Yu Kitaev, A. Shen, M. N. Vyalyi Classical and Quantum Computation, American Mathematical Society, 2002
- [11] A. Kitaev and J. Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. *Proceedings of the 32nd ACM Symposium on Theory of Computing*, 608-617, 2000.
- [12] L. Lovasz: Random Walks on Graphs: A Survey [in: *Combinatorics, Paul Erds is Eighty*, Vol. 2 (ed. D. Mikls, V. T. Ss, T. Sznyi), Jnos Bolyai Mathematical Society, Budapest, 1996, 353-398.]
- [13] Motwani and Raghavan, *Randomized algorithms*, Cambridge University Press, 1995
- [14] C. Papadimitriou, *computational complexity*, Addison-Wesley, 1994

- [15] Alistair Sinclair, Algorithms for random generation and counting: a Markov chain approach, Birkhauser Verlag, Basel, Switzerland, 1993
- [16] J. Watrous. PSPACE has constant-round quantum interactive proof systems. Proceedings of the 40th Annual Symposium on Foundations of Computer Science, pages 112-119, 1999.
- [17] J. Watrous. Succinct quantum proofs for properties of finite groups. Proceedings of the 41st Annual Symposium on Foundations of Computer Science, pages 537-546, 2000.
- [18] W.K.Wootters and W.H.Zurek, Nature (London), **299**, 802, 1982
- [19] Stathis Zachos, Martin Furer: Probabilistic Quantifiers vs. Distrustful Adversaries. *Proc. FSTTCS*, Springer-Verlag, Lecture Notes in Computer Science (Vol. 287) pages 443-455, 1987.