CS 231 Final Project: MA and Stoquastic Hamiltonians in Quantum Complexity Theory

Benji Kan Vassilios Kaxiras

May 14, 2022

1 Introduction

Complexity theory quantifies the difficulty of solving computational problems and classifies such problems by comparative difficulty. With recent experimental and theoretical advances in quantum computation, there is renewed interest in studying the complexity of various quantum models of computation.

In this project, we survey three essential classes in the study of quantum complexity and highlight their relationships: NP in Section 3, QMA in Section 4, and MA in Section 5. Our selected progression is designed to showcase the intuitive connections and interplay between definition and characteristics that build off of each class. One focus is on situating the complexity class MA in the landscape between the better-studied NP and QMA classes; to this end, we define stoquastic Hamiltonians and other preliminaries in Section 2. As a preview, we will see that stoquastic Hamiltonians lie on the presumed boundary between classical and quantum complexity, as quantum objects which define a complete problem for a classical complexity class.

As our goal is to survey the landscape of complexity at a high-level, we omit many crucial technical details (which can be found in the original papers), opting instead for proof sketches and intuitive descriptions of main ideas. The purpose of this document, then, is to emphasize and elucidate the relationships between these classes at an intuitive level, distilling some of the magic of quantum in the world of complexity theory.

2 Preliminaries

Definition 2.1 (k-local Hamiltonian)**.** A k-local Hamiltonian is a Hermitian operator H over n qubits that is the sum of M k-local operators H_a:

$$H = \sum_{a=1}^{M} H_a \tag{1}$$

Each k-local H_a acts non-trivially only on $\leq k$ qubits.

Definition 2.2 (Stoquastic k-local Hamiltonian)**.** A k-local Hamiltonian is called stoquastic with respect to a basis if every off-diagonal element of it in that basis is non-positive.
In this paper, we will use the shorthand "stoquastic Hamiltonian" to refer to a stoquastic k-local Hamiltonian in the computational basis.

One particularly interesting property of stoquastic Hamiltonians is that their ground eigenstates can always be taken to have non-negative components. To show this, we invoke a theorem:

Theorem 2.3 (Perron-Frobenius)**.** *If all entries of a square matrix A are positive, then it has a unique maximal eigenvalue. Its eigenvector has positive entries. [12]*

Now given a stoquastic Hamiltonian H, consider the Hamiltonian $A = I - \beta H$. We can choose β small enough such that all elements of A are non-negative. Note that the eigenvector of the largest eigenvalue of A is the ground state of H. Then by the Perron-Frobenius theorem, the ground state of H has positive entries. This result indicates that the ground states of stoquastic Hamiltonians have a natural interpretation: that of a probability distribution.
As a direct example of this probability distribution interpretation, one class of stoquastic Hamiltonians are those with thermal ground states, of the form

$$p(x) = |\langle\psi|x\rangle|^2 = \frac{1}{Z}e^{-\beta E_x} \implies |\psi\rangle = \frac{1}{\sqrt{Z}}\sum_x e^{-\beta E_x/2}|x\rangle \tag{2}$$

where Z is the partition function of the (classical) system represented by this Hamiltonian, with energy spectrum $\{E_x\}$.
Like this example, stoquastic Hamiltonians are commonly found in real-world problems. For instance, all Hamiltonians in the Transverse Ising Model (TIM), which are of the form

$$H = -J\sum_{\langle i,j\rangle} Z_i Z_j - g\sum_i X_i \tag{3}$$

are stoquastic. TIM Hamiltonians represent nearest-neighbor spin interactions plus an external magnetic field in an atomic lattice. The physical relevance of this setup is based on its connections to quantum optics experiments and quantum simulators.

3 NP Complexity Class

The complexity class NP, nondeterministic polynomial time, is a complexity class capturing a set of decision problems where the proofs for "yes" instances are verifiable in polynomial time by a Turing machine. Alternatively, this class can be interpreted as the set of problems that can be solved in polynomial time by a nondeterministic Turing machine. Formally, we define NP below.

Definition 3.1 (NP)**.** The class NP is the set of languages L that can be verified in *deterministic* polynomial time. To determine if $x \in L$, the prover sends a proof y with $|y| = poly(|x|)$ to the verifier. Then,

$$\begin{aligned} x \in L &\implies \exists y \text{ s.t. } V(x,y) = 1, \\ x \notin L &\implies \forall y; V(x,y) = 0. \end{aligned} \tag{4}$$

Observe that these statements are entirely deterministic in nature; there is no probabilistic component to the two cases described.

3.1 Cook-Levin Theorem and NP-completeness

One of the most classical results in complexity theory is the Cook-Levin Theorem, which gives a fundamental NP-complete problem.

Theorem 3.2 (Cook-Levin Theorem)**.** *The Boolean satisfiability problem (SAT) is NP-complete.*

We sketch the proof for this theorem, which proceeds in two parts, first showing that SAT is in NP, and then that SAT is NP-hard. First, we see that any claimed satisfying assignment of Boolean variables to a given expression can be efficiently verified by a Turing machine, so SAT is indeed in NP.

Now, we consider any problem L supposedly in NP. Then, there is a Turing machine M which checks if $x \in L$ using y. We will construct a reduction of the operation of this machine to k-SAT. First, we can represent the finite possibilities of the state of the machine and its time as an assignment; this assignment represents, in essence, the *history* of a computation. Then, we can construct a k-SAT formula which checks that the evolution of the machine is valid: it will check if the ith tape location at time $t+1$ follows logically from the $i-1, i, i+1$st tape locations at time t. If every step in the computation is valid, as are its start and end states, then we have a valid computation history, which can be verified efficiently.

In particular, we need only perform *local verification* of the evolution of the machine; thus, we can represent each verification as a k-clause. Indeed, this reduction actually produces a k-SAT instance, which is a subproblem of the more general SAT problem. We will soon see how this local property leads naturally to the quantum analog in other complexity classes.

4 QMA Complexity Class

The idea of quantum nondeterminism was first introduced and studied by Knill [11], and further explored in the complexity class QMA by Kitaev [10]. QMA, also known as BQNP, is intuitively understood as the quantum analog of the class NP in a probabilistic setting. Thus, the relationship between QMA and BQP can be understood as the quantum version of the relationship between NP and P. Formally, we define QMA below.

Definition 4.1 (QMA)**.** The class QMA is the set of languages L that can be verified in a polynomial size *quantum* circuit. To determine if $x \in L$, the prover sends a proof $|\psi\rangle$ with $||\psi\rangle| = poly(|x|)$ qubits to the verifier. Then, the verifier runs a quantum circuit V with number of gates also $poly(|x|)$.

$$\begin{aligned} x \in L &\implies \exists\, |\psi\rangle \text{ s.t. } \mathbb{P}[V(x, |\psi\rangle] = 1] \geq \frac{2}{3}, \\ x \notin L &\implies \forall\, |\psi\rangle\,; \mathbb{P}[V(x, |\psi\rangle] = 1] \leq \frac{1}{3}. \end{aligned} \tag{5}$$

A diagram displaying an example verification circuit is shown in Appendix A.1.

Note the two key differences between this definition of QMA and the previous definition of NP: 1. the classical witness proof y has been replaced by a quantum witness state $|\psi\rangle$, and 2. the conditions are framed probabilistically. These two changes are crucial for any quantum analogy, due to the inherently probabilistic nature of the quantum world.

4.1 QMA-completeness

We will see that the Local Hamiltonian problem is QMA-complete. First, we define the Local Hamiltonian problem, which is a quantum version of the classical Boolean k-SAT problem.

Definition 4.2 (Quantum k-SAT (Local Hamiltonian))**.** Quantum k-SAT (also known as the Local Hamiltonian problem) is the quantum analog of the Boolean satisfiability problem k-SAT. The input is a set of positive semi-definite k-local Hamiltonians $H_1, H_2, \ldots, H_M$, where each H_a represents a

clause that must be satisfied, just as the clauses in classical k-SAT. Satisfying a clause H_a means finding an eigenvector with eigenvalue 0. Thus for "yes" instances which satisfy all clauses, there exists some ground state $|\psi\rangle$ such that

$$\forall a,\ H_a |\psi\rangle = 0 \implies H |\psi\rangle = 0, \quad H \equiv \sum_{a=1}^{M} H_a \tag{6}$$

where we have defined the combined Hamiltonian H. For "no" instances, we have $\forall |\psi\rangle,\ H |\psi\rangle \geq \epsilon$ for some separation parameter ϵ.

We provide a comparison of the analogies between the classical and quantum k-SAT problems in Table 1, provided in Appendix A.1.

In a classic result, Kitaev first showed that the 5-Local Hamiltonian problem is QMA-complete [10]. As full proofs of the Local Hamiltonian problem being QMA-complete are discussed in detail elsewhere [3], we simply highlight the main ideas in this completeness construction.

In the first step to showing QMA-completeness, we verify that k-SAT is indeed in QMA.

Theorem 4.3. *The k-Local Hamiltonian problem is in QMA.*

The main idea behind the proof to this theorem relies on the observation that if $x \in L$, then there exists a ground state with small eigenvalues; otherwise, any witness state used has large eigenvalues. If all Hamiltonians H_i were simple projections $H_i = |\alpha_i\rangle\langle\alpha_i|$, then we could easily verify by randomly selecting an i as the basis and measuring $|\eta\rangle$ in the basis $\{|\alpha_i\rangle, |\alpha_i^{\perp}\rangle\}$. In general, we rotate the qubits of each term of the local Hamiltonian H_i based on its spectral decomposition, and then proceed as in the projections case by measuring in a basis of a randomly chosen local Hamiltonian. It can be shown that this procedure produces the desired output with high probability.

Next, we complete the proof for QMA-completeness of k-SAT.

Theorem 4.4. *The k-Local Hamiltonian problem is QMA-hard.*

That is, we wish to show a reduction of any QMA problem to a k-Local Hamiltonian problem. Recall that for a QMA language L, an input $x \in L$ can be verified by a polynomial size quantum circuit with a witness $|\psi\rangle$. Drawing inspiration from the proof of the classical Cook-Levin Theorem, we aim to locally verify the history of our computation, which can be written as a sequence of states.

However, quantum computation is not locally checkable, as a reduced density matrix may be indistinguishable for two different states; the CAT state is well-known as an illuminating example of this technicality. In order to address this issue, we verify the computation with a superposition over the entire history, including the clock register. This idea of moving from checking each step in time-evolution to an overall time-independent local Hamiltonian was originated in part by Feynman, and is shown by the superposition history state:

$$|\text{history}\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} U_t \dots U_1 |x, \psi\rangle |t\rangle \tag{7}$$

Another key difference from the classical reduction involves the continuous nature of the quantum model; while showing that the classical eigenvalue is nonzero is sufficient for soundness (as discreteness ensures it is at least 1), the quantum eigenvalue needs to be polynomially bounded away from 0. Kitaev shows this stronger condition with a geometric application of Jordan's Lemma.

In his original result, Kitaev proved that this reduction holds for 5-local Hamiltonians, as he uses a unary representation of time with 3 qubits. Later, Kempe, Kitaev, and Regev improved this result to show that k-local Hamiltonians for any $k \geq 2$ are also QMA-complete, using linear algebra and perturbation theory arguments, the details of which exceed the scope of this comparison [9]. Further QMA-complete problems are listed in [5].

5 MA Complexity Class

5.1 Preliminaries

The class MA was first introduced by Babai [4] as one of a set of classes "just above NP", defined by a wizard Merlin that attempts to convince a king Arthur that a string x belongs to a language L using statistical evidence. MA refers to the case where Merlin sends only one message to Arthur, and thus can be thought of as the probabilistic analogue of the class NP.

Definition 5.1 (MA)**.** Concretely, MA consists of all problems for which there is a probabilistic polynomial-time verifier V that takes problem inputs x and witnesses w such that

$$\begin{aligned} x \in L &\implies \exists\, w \text{ s.t. } \mathbb{P}[V(x,w)=1] \geq p_{yes}, \\ x \notin L &\implies \forall\, w;\ \mathbb{P}[V(x,w)=1] \leq p_{no}. \end{aligned} \tag{8}$$

and $p_{yes} - p_{no} = 1/poly(|x|)$ [8].

Curiously, the only natural example of an MA-complete problem is quantum in nature. This is known as the stoquastic k-SAT problem. To understand this problem, we build off of the more general quantum k-SAT defined in Section 4, also known as the Local Hamiltonian problem.

Stoquastic k-SAT is simply a restricted version of quantum k-SAT:

Definition 5.2 (Stoquastic k-SAT)**.** Stoquastic k-SAT is an instance of quantum k-SAT where every H_a is stoquastic; that is each local H_a has non-positive off-diagonal elements.

Theorem 5.3. *Stoquastic k-SAT is MA-complete.*

We prove this theorem by showing that stoquastic k-SAT is both in MA and MA-hard.

5.2 Stoquastic k-SAT is in MA

To show stoquastic k-SAT is in MA, Bravyi and Terhal construct a probabilistic verifier that, for a given $|\psi\rangle$, checks if $|\psi\rangle$ satisfies the stoquastic Hamiltonian H [6]. Instead of being given the complete description of $|\psi\rangle$, the verifier only requires a single basis element $|x_0\rangle$ with nonzero amplitude in $|\psi\rangle$ as a witness. Note that $|\psi\rangle$ satisfying H necessarily implies $|x_0\rangle$ will have a non-zero projection into the 0-eigenspace of every H_a (which is the space of satisfying solutions). In other words, denoting the projector into the 0-eigenspace of H_a as Π_a, this condition means

$$\langle x_0| \Pi_a |x_0\rangle > 0 \tag{9}$$

The verifier then simulates a random walk, starting at the basis element it is given and sampling the probability distribution at each step to progress to another basis element. At every step, it obtains the current state $|x\rangle$ and checks that the random walk distribution is properly normalized, and that $\langle x| \Pi_a |x\rangle > 0$ still holds. If H is satisfiable, the random walk will be well-defined and will proceed always within the set of all basis elements with non-zero amplitude in $|\psi\rangle$. This means the above two conditions will always be met. If H is not satisfiable, then the random walk will contain a component in the 0-eigenspaces Π_a with exponentially decreasing (in the number of walk steps) probability. Eventually, the random walk will run into a state for which

$$\langle x| \Pi_a |x\rangle = 0 \tag{10}$$

and the verifier knows to reject the witness.
As this result is one of the core topics of this review, we provide the details of this proof in Appendix A.2.

5.3 Stoquastic k-SAT is MA-hard

To prove the completeness condition of stoquastic k-SAT within MA, it remains to show that any problem in MA can be reduced to it. Bravyi and Terhal show this in the following sequence:

1. Define a quantum analogue to MA, denoted MA_q, where the verifier V_x for a particular problem input x is a classical reversible circuit that takes as input a quantum witness $|\chi\rangle$ and a set of ancilla qubits initialized to $|0\rangle$. A $|+\rangle$ state substitutes for the randomness in MA.

2. Show $\mathrm{MA}_q = \mathrm{MA}$. Since V_x is classical, one can show that a satisfying quantum witness $|\chi\rangle$ exists if and only if a satisfying single basis state witness exists; the replaced randomness from the $|+\rangle$ state does not change the outcome. Thus both V_x and $|\chi\rangle$ can be taken to be classical, which means $\mathrm{MA}_q = \mathrm{MA}$.

3. Using the clock construction, show that we can define a 6-local Hamiltonian whose ground state energy is 0 if and only if x is a "yes" instance. Specifically, consider the set of gates in V_x, and construct a clock Hamiltonian that verifies the correct evaluation of V_x gate-by-gate, with V_x outputting 1 at the end.

4. Show that the above clock Hamiltonian is stoquastic, which follows trivially from writing its description down. Thus one can determine if it is satisfiable using stoquastic k-SAT, which means we can reduce any problem in MA to stoquastic 6-SAT.

5.4 Comparison to QMA

While quantum k-SAT is significant for being QMA-hard, the key idea for stoquastic k-SAT is that it resides in MA. Thus, intuitively, k-SAT and stoquastic k-SAT lie on either side of the border between QMA and MA, establishing stoquasticity as critical to understanding the distinction between quantum and classical complexity.

6 Conclusion and Future Work

In this report, we compare the complexity classes NP, QMA, and MA, and discuss the differences in their respective completeness problems. These classes are related to each other via the following inclusions, which may or may not be unconditionally strict:

$$\mathrm{P} \subseteq \mathrm{NP} \subseteq \mathrm{MA} \subseteq \mathrm{QMA} \subseteq \mathrm{PSPACE} \tag{11}$$

Of particular interest is the inclusion $\mathrm{MA} \subseteq \mathrm{QMA}$, as this represents the nominal transition from classical to quantum computation. By investigating complete problems for both MA and QMA, we have discovered that stoquastic k-SAT is a problem that lies on the border between classical and quantum complexity.

As an extension to the described hierarchy, work has been done relating the complexity class QMA to QCMA, which involves QMA restricted to a classical witness $|\alpha\rangle$ [1]. These investigations aim to reveal where quantum proofs offer a computational advantage over classical proofs, if any.

As alluded to in our project, stoquastic Hamiltonians also provide other crucial roles in the study of complexity theory. One incredible result is the introduction of problems in the Tranverse field Ising Model (TIM), a very physically relevant model, and showing that estimating the ground state energy of TIM is StoqMA-complete, an extension of MA [7]. A recent result connects the gap amplification of stoquastic Local Hamiltonians to progress on the quantum PCP conjecture, a long-standing and important question in the field [2].

Acknowledgements We would first like to express our gratitude to Anurag Anshu for teaching an incredible and eye-opening course on Quantum Complexity, and inspiring our study of quantum complexity classes. We also thank Chi-Ning Chou and Prayaag Venkat for helpful discussions regarding the landscape of quantum complexity in theoretical computer science more generally. Finally, we thank our peers and colleagues in CS 231 for cultivating a supportive and collaborative environment to study quantum information.

References

[1] Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 115–128, 2007.

[2] Dorit Aharonov and Alex Bredariol Grilo. Stoquastic PCP vs. randomness. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1000–1023. IEEE Computer Society, 2019.

[3] Dorit Aharonov and Tomer Naveh. Quantum np - a survey. *arXiv: Quantum Physics*, 2002.

[4] L Babai. Trading group theory for randomness. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC '85, page 421–429, New York, NY, USA, 1985. Association for Computing Machinery.

[5] Adam D. Bookatz. Qma-complete problems. *Quantum Info. Comput.*, 14(5 amp; 6):361–383, apr 2014.

[6] Sergey Bravyi, Arvid J. Bessen, and Barbara M. Terhal. Merlin-arthur games and stoquastic complexity, 2006.

[7] Sergey Bravyi and Matthew Hastings. On complexity of the quantum ising model. 2014.

[8] Sergey Bravyi and Barbara Terhal. Complexity of stoquastic frustration-free hamiltonians. *SIAM Journal on Computing*, 39(4):1462–1485, 2010.

[9] Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006.

[10] A.Y. Kitaev, A. Shen, M.N. Vyalyi, and M.N. Vyalyi. *Classical and Quantum Computation*. Graduate studies in mathematics. American Mathematical Society, 2002.

[11] E. Knill. Quantum randomness and nondeterminism, 1996.

[12] Oliver Knill. Lecture 34: Perron frobenius theorem, 2011.

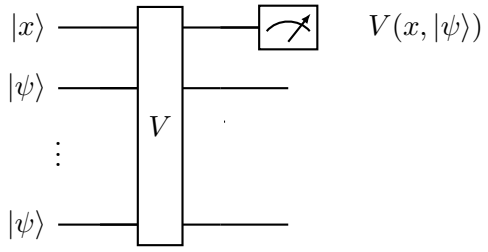

Figure 1: Example verification circuit for QMA.

Table 1: Comparison of classical and quantum k-SAT problems.

Classical	Quantum
Constraint Satisfiability Problem (CSP) ϕ	Hamiltonian $H = \sum_{a=1}^{M} H_a$
Variable x_i	Qubit $\|x_i\rangle$
k-Constraint $(x_1 \vee x_2 \vee \cdots \vee x_k)$	k-Local Hamiltonian H_a
Assignment y	Quantum state $\|\psi\rangle$
Number satisfied constraints for y	Energy of $\|\psi\rangle$ for Hamiltonian $\langle\psi\| H \|\psi\rangle = \lambda$
Satisfying assignment	Hamiltonian ground state $\|\eta\rangle$

A Appendix

A.1 Miscellaneous figures and tables

In Figure 1, we display an example of the QMA verification circuit.

In Table 1, we compare the analogous components of the classical and quantum k-SAT problems.

A.2 Stoquastic k-SAT is in MA

Consider an instance of stoquastic k-SAT. We can find the projectors $\{\Pi_a\}$ onto the 0-eigenspaces of the Hamiltonians $\{H_a\}$ efficiently, since each H_a is local. Now suppose H is satisfiable by $|\psi\rangle$, so by definition $\Pi_a |\psi\rangle = |\psi\rangle$ for all a. Define $S \subseteq \{0,1\}^n$ as the set of basis states with non-zero amplitudes in $|\psi\rangle$:

$$S \equiv \{x \,:\, \langle x|\psi\rangle > 0\} \tag{12}$$

Note that $\langle x|\psi\rangle > 0 \iff \langle x|\psi\rangle \neq 0$ since H is stoquastic.
Now we define a random walk over this set S, with the transition probability

$$P_{xy} = \frac{\langle y|\psi\rangle}{\langle x|\psi\rangle} \frac{1}{M} \sum_{a=1}^{M} \langle x| \Pi_a |y\rangle \tag{13}$$

To see this is a well-defined random walk if H is satisfiable, first note

$$\sum_y P_{xy} = \frac{1}{M} \sum_a \langle x| \Pi_a \sum_y |y\rangle \langle y| \frac{|\psi\rangle}{\langle x|\psi\rangle} = \frac{1}{M} \sum_a \frac{\langle x| \Pi_a |\psi\rangle}{\langle x|\psi\rangle} = \frac{1}{M} \sum_a \frac{\langle x|\psi\rangle}{\langle x|\psi\rangle} = 1 \tag{14}$$

Secondly, $P_{xy} \neq 0 \iff \langle y|\psi\rangle \neq 0$, so if the random walk starts in S it can never leave S.
So how can we compute P_{xy}? Consider the 0-eigenspace of H_a, spanned by orthonormal $\{|\psi_i\rangle\}$. Then

$$\Pi_a = \sum_i |\psi_i\rangle \langle\psi_i| \tag{15}$$

Since $\Pi_a |\psi\rangle = |\psi\rangle$, we have that $|\psi\rangle = |\psi_i\rangle$ for some i. Furthermore, since $\langle\psi_i|\psi_j\rangle = \delta_{ij}$ and the elements of $|\psi_i\rangle$ are all non-negative, we have $\langle x|\psi_i\rangle > 0 \implies \langle x|\psi_j\rangle = 0$ for $i \neq j$. So if $\langle x|\psi\rangle, \langle y|\psi\rangle > 0$, then

$$\frac{\sqrt{\langle x| \Pi_a |x\rangle}}{\sqrt{\langle y| \Pi_a |y\rangle}} = \frac{\sqrt{\langle x|\psi\rangle \langle\psi|x\rangle}}{\sqrt{\langle y|\psi\rangle \langle\psi|y\rangle}} = \frac{\langle x|\psi\rangle}{\langle y|\psi\rangle} \tag{16}$$

Thus for any arbitrary a,

$$P_{xy} = \frac{\sqrt{\langle x| \Pi_a |x\rangle}}{\sqrt{\langle y| \Pi_a |y\rangle}} \frac{1}{M} \sum_{a'=1}^{M} \langle x| \Pi_{a'} |y\rangle \tag{17}$$

To compute this, all we have to do is access the elements of Π_a for a given a. This is efficient, because we can effectively diagonalize H_a in its non-trivial subspace and find the complete matrix description of Π_a in that subspace. Furthermore, $|S| \leq 2^k M$, since each component in $|\psi\rangle$ must be acted upon non-trivially by H to send its amplitude to 0, and H acts on at most $2^k M$ terms non-trivially. This means one state can only transition to at most $2^k M$ other states. So if we simulate the random walk only storing the current state (in terms of a computational basis state), to perform one step of this simulation we only need to perform a constant (with respect to the number of qubits n of H) number of operations.

A.2.1 Completeness

After every step of the walk which gives the current state $|x\rangle$, we check that $|x\rangle$ is still in the ground eigenspace of every H_a by evaluating

$$\langle x| \Pi_a |x\rangle > 0 \tag{18}$$

for all a. If at any point this condition is not satisfied, we reject the witness.
If $|\psi\rangle$ satisifies H, the above check will always be true, because x will remain in S throughout the walk and thus

$$\langle x| \Pi_a |x\rangle = \langle x|\psi\rangle \langle\psi|x\rangle > 0 \tag{19}$$

for all a.

A.2.2 Soundness

What if H is not satisfiable? Then there is no state $|\psi\rangle$ such that $H|\psi\rangle = 0$. This means that the random walk may have unnormalized probabilities, since the normalization relies on $\langle x| \Pi_a |\psi\rangle = \langle x|\psi\rangle$. This provides one check that we can apply as we're simulating the random walk: if all of a sudden the probabilities become unnormalized, then we know to reject the witness. This does not detect every unsatisfiable H, however. We also need to analyze the aforementioned condition $\langle x| \Pi_a |x\rangle > 0$.
First define

$$G \equiv \frac{1}{M} \sum_{a=1}^{M} \Pi_a \tag{20}$$

as used in the definition of P_{xy}. H being unsatisfiable means for any $|\psi\rangle$

$$H|\psi\rangle \geq \epsilon \implies \langle\psi| G |\psi\rangle \leq 1 - \delta \tag{21}$$

for some δ. So the largest eigenvalue of G is $1 - \delta$. Now define D as the set of all "good" random walk states:

$$D \equiv \{x \ : \ \forall\, a\,, \langle x| \Pi_a |x\rangle > 0\} \tag{22}$$

We wish to find the probability that the walk remains in D, as once it leaves D we know to reject. This can be found by simply enumerating every possible path the walk can take from an initial state x_0 to a state x_L after L steps:

$$\mathbb{P}[x_1, x_2, \ldots, x_L \in D] = \sum_{x_1, x_2, \ldots, x_L \in D} P_{x_0 \to x_1} P_{x_1 \to x_2} \ldots P_{x_{L-1} \to x_L} \tag{23}$$

Now, note

$$P_{x_i \to x_{i+1}} = \frac{r_i}{r_{i+1}} \langle x_i | G | x_{i+1} \rangle, \quad r_i \equiv \sqrt{\langle x | \Pi_a | x \rangle} \tag{24}$$

So

$$\begin{aligned} \mathbb{P}[x_1, x_2, \ldots, x_L \in D] &= \sum_{x_1, x_2, \ldots, x_L \in D} \frac{r_0}{r_L} \langle x_0 | G | x_1 \rangle \langle x_1 | G | x_2 \rangle \langle x_2 | \ldots | x_{L-1} \rangle \langle x_{L-1} | G | x_L \rangle \\ &= \sum_{x_1, x_2, \ldots, x_L \in D} \frac{r_0}{r_L} \langle x_0 | G^L | x_L \rangle \end{aligned} \tag{25}$$

Since the maximum eigenvalue of G is $1 - \delta$, G^L will exponentially approach 0. Since the rest of the above expression does not change with increasing L, this means the probability of remaining in D decreases exponentially with L, so we know to reject the witness quickly.