Complexity of Basis-Restricted Local Hamiltonians

by

Henry Ma

B.S., University of California, Los Angeles (2021)

Submitted to the Department of Electrical Engineering and Computer Science in Partial Fulfillment of the Requirements for the Degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2025

Authored by: Henry Ma
Department of Electrical Engineering and Computer Science
May 16, 2025

Certified by: Anand Natarajan
Assistant Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by: Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Complexity of Basis-Restricted Local Hamiltonians

by

Henry Ma

Submitted to the Department of Electrical Engineering and Computer Science
on May 16, 2025 in Partial Fulfillment of the Requirements for the Degree of
Master of Science

Abstract

A major goal of quantum complexity theory is to understand which computational problems can be solved with access to certain quantum resources. The subfield of *Hamiltonian complexity* specifically considers computational problems that ask about properties of local Hamiltonians, which are of critical importance in quantum complexity because they can be viewed as quantum generalizations of classical constraint satisfaction problems.

In this work, we study the complexity of certain restricted variants of the QUANTUM-k-SAT problem, a quantum analog of the NP-complete k-SAT problem. We introduce new variants of QUANTUM-k-SAT which place a *basis restriction* on the input Hamiltonian $H = \sum_i h_i$. Each variant is defined by a fixed collection of bases $B_1, \ldots, B_r$ of n-qubit space. We require that each Hamiltonian term h_i must be diagonal in one of these bases.

Our results resolve the complexity of certaim basis-restricted variants of QUANTUM-k-SAT. First we show the QUANTUM-6-SAT problem with Hamiltonian terms restricted to be diagonal in an X/Z mixed basis is QMA_1-complete.

Second, we combine basis restriction with the restriction of commutativity, and show the following easiness result, which applies generally to higher-level quantum systems (qudits) and bases Q and R (which are real-valued and satisfy an overlap condition): The commmuting QUANTUM-SAT problem on qudits, where Hamiltonian terms are either diagonal in the Q basis, the R basis, or a single mixed Q/R basis, is in NP.

Thesis supervisor: Anand Natarajan
Title: Assistant Professor of Electrical Engineering and Computer Science

Acknowledgements

First and foremost, I am extremely grateful to my advisor Anand Natarajan for his support during this project. This work would not have been possible without his many dedicated hours of mentorship.

I'm also grateful for the everlasting support of my family, especially the encouragement and love from my parents.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

A major goal of quantum complexity theory is to understand which computational problems can be solved with access to certain quantum resources. For a given quantum model of computation, we can define the corresponding *complexity class* of problems which can be solved under this model. The definitions and techniques in quantum complexity draw heavily upon the well-developed theory of classical complexity. For example, consider one of the most fundamental classical complexity classes, NP, which is the class of problems that can be *verified* by a polynomial-time deterministic algorithm. NP has a number of quantum generalizations, including the classes QMA and QMA_1, which capture the problems which can be verified by a polynomial-time *quantum* algorithm.

To study these quantum complexity classes, the subfield of *Hamiltonian complexity* has emerged, which specifically considers computational problems that ask about properties of local Hamiltonians. Local Hamiltonians are of critical importance in quantum complexity because they can be viewed as noncommutative quantum generalizations of classical constraint satisfaction problems (CSPs). Seminal work by Cook [Coo71] showed that a CSP problem known as 3-SAT is NP-complete. Morally speaking, this means that 3-SAT is the *hardest* problem which can be solved using the classical verification model of NP, and thus captures the computational power of NP. In a similar vein, a line of research in Hamiltonian complexity has tried to understand the class QMA_1 through the local Hamiltonian problem QUANTUM-k-SAT, culminating in work by Gosset and Nagaj [GN16] which showed that QUANTUM-3-SAT is QMA_1-complete.

1.1 Our contributions

In this work, we study the complexity of certain restricted variants of the QUANTUM-k-SAT problem. In order to state our main results, we first give an informal overview of the restrictions placed on QUANTUM-k-SAT; this will be discussed formally in chapter 2.

In the QUANTUM-k-SAT problem, we are given a k-local Hamiltonian H, and asked to determined whether its ground energy is 0. H is a Hermitian matrix acting on n qubits which can be written in the form $H = \sum_i h_i$, where each h_i is a projector which acts nontrivially on at most k qubits. As a limiting case, the classical k-SAT problem can be viewed as a special case of QUANTUM-k-SAT in which each h_i is diagonal in the computational (Z) basis.

We introduce new variants of Quantum-k-Sat which place a *basis restriction* on the input Hamiltonian H. Each variant is defined by a fixed collection of bases $B_1, \ldots, B_r$ of n-qubit space. We require that each Hamiltonian term h_i must be diagonal in one of these bases.

What motivates us to study basis-restricted problems? On the one hand, as the number of allowed bases decreases, the problem seemingly becomes more classical than general Quantum-k-Sat; indeed, in the extreme case of a single allowed basis, the problem reduces to the classical problem k-Sat. On the other hand, it is still plausible that a basis-restricted problem (with two or more distinct bases) is no easier to solve than general Quantum-k-Sat, despite its simpler definition. Thus, it is interesting to precisely nail down the complexity of basis-restricted problems: when is the problem still expressive enough to be a complete problem for QMA_1, and when does the problem become so simple that it has a classical verification procedure, putting it in NP?

One other important variant of Quantum-k-Sat, which has been studied extensively, places a restriction of *commutation* on the input Hamiltonian H: for all i and j, h_i and h_j must commute. Similar to basis-restricted problems, the commuting variant of Quantum-k-Sat is interesting because its complexity lies somewhere between NP and QMA_1 (note that a classical k-Sat instance, when viewed as a local Hamiltonian, has commuting h_i terms).

Our main results resolve the complexity of certain basis-restricted variants of Quantum-k-Sat. First, we prove the following hardness result:

- (Informal version of Theorem 3.1.2) The Quantum-6-Sat problem with Hamiltonian terms restricted to be diagonal in a X/Z mixed basis is QMA_1-complete.

Second, we combine basis restriction with the commutativity restriction, and show the following easiness result, which applies generally to higher-level quantum systems (qudits) and bases Q and R (satisfying some technical assumptions):

- (Informal version of Theorem 4.2.4) The commmuting Quantum-Sat problem on qudits, where Hamiltonian terms are either diagonal in the Q basis, the R basis, or a single mixed Q/R basis, is in NP.

The rest of the thesis is structured as follows. Chapter 2 introduces background from complexity theory and quantum computation. In chapter 3, we prove a QMA_1-completeness result for basis-restricted Quantum-k-Sat. In chapter 4, we prove a that an instance of basis-restricted commuting Quantum-k-Sat is in NP. Finally, we discuss open problems and future work in chapter 5.

Chapter 2

Preliminaries

2.1 Classical complexity theory

We begin by recalling basic notions from classical complexity theory, and refer the reader to [Sip96] for more detail. The computational problems we consider in this work are decision problems with classical input. In such a problem, we are given a classical description of an object (e.g. a boolean formula) and asked to decide a certain property of the object (e.g. whether it has a satisfying assignment). Formally, a problem is represented by a *language* L, which is a subset $L \subseteq \{0,1\}^\star$ of finite-length bitstrings. The problem corresponding to L is: given $x \in \{0,1\}^n$, is x in L? In other words, L contains the "yes" instances of its corresponding problem. We commonly use $n = |x|$ to denote the length of the input.

A language L is *efficiently solvable* if there is a polynomial-time deterministic algorithm A which decides L. That is, there is some polynomial p such that on any input x of length n, A halts in at most $p(n)$ steps, and the output $A(x)$ satisfies: $A(x) = 1$ if $x \in L$, and $A(x) = 0$ if $x \notin L$. The complexity class P consists of all languages which are efficiently solvable.

A language L is *efficiently verifiable* if there is a deterministic algorithm A which takes an input x of length n and an additional bitstring y of length $\mathrm{poly}(n)$, runs in $\mathrm{poly}(n)$ time, outputs $A(x,y) \in \{0,1\}$, and satisfies the following guarantees:

1. If $x \in L$, then there is some y such that $A(x,y) = 1$.

2. If $x \notin L$, then for all y, $A(x,y) = 0$.

We think of y as a proof that x is in L. The first condition above is known as *completeness* and the second is known as *soundness*. Together, they capture the notion that membership in L can be efficiently verified with a proof y: if x is indeed in L, there is some proof y which will quickly convince you of this fact; otherwise, no proof will convince you. The complexity class NP consists of all languages which are efficiently verifiable.

The *satisfiability problem* asks whether a boolean formula in conjunctive normal form (CNF) is satisfiable. A CNF formula φ on boolean variables $x_1, \ldots, x_n$ is a conjunction (logical "and") of clauses $\varphi = C_1 \wedge \cdots \wedge C_m$. Each clause C_i is itself a disjunction (logical "or") of literals $C_i = l_{i,1} \vee \cdots \vee l_{i,r}$, where each literal $l_{i,j}$ is a variable or its negation. φ is satisfiable if there is some assignment $x = x_1 \cdots x_n \in \{0,1\}^n$ to the variables which makes φ

evaluate to True. Formally, the language Sat contains all satisfiable CNF boolean formulas. It is clear that Sat is in NP: the verifying algorithm receives the input formula φ and a proof string x, then checks that x is a satisfying assignment to the variables of φ.

For an integer $k \geq 1$, the *k-satisfiability problem* asks whether a k-CNF formula is satisfiable. A k-CNF formula is a CNF formula in which each clause is an disjunction of exactly k literals. Let k-Sat be the language corresponding to the k-satisfiability problem. Surprisingly, it is known that 2-Sat is in P [Kro67]. In contrast, it is widely believed that there is no efficient algorithm which solves 3-Sat. This belief stems from the important notion of *NP-completeness*, which aims to characterize the "hardest" problems in NP.

A language K has an *efficient reduction* to a language L if there is an efficient deterministic algorithm A which maps a bitstring x to a bitstring $A(x)$ and satisfies: $x \in K$ if and only if $A(x) \in L$. Intuitively, K is no harder than L: deciding whether x is in K reduces to the problem of deciding whether $A(x)$ is in L. A language L is *NP-hard* if for every $K \in \mathsf{NP}$, there is an efficient reduction from K to L. A language L is *NP-complete* if L is in NP and L is NP-hard. Cook [Coo71] introduced the notion of NP-completeness and showed that 3-Sat is NP-complete. We can interpret L being NP-complete as evidence that it is not efficiently solvable. Indeed, under the widely held belief that there are problems which are efficiently verifiable but not efficiently solvable (i.e. $\mathsf{P} \neq \mathsf{NP}$), if there is an efficient algorithm which solves L, then *every* problem in NP is efficiently solvable.

2.2 Quantum computation

We now set up the basic formalism of quantum computation. For a more detailed exposition, we refer the reader to [NC10].

An n-qubit *quantum state* $|\psi\rangle$ is a unit vector in a complex Hilbert space $(\mathbb{C}^2)^{\otimes n}$. A k-qubit *quantum operation* is a unitary operator U acting on a k-qubit space $(\mathbb{C}^2)^{\otimes k}$.

Let $[n] = \{1, \ldots, n\}$. We can extend U to an operation $U \otimes I_{[n]\setminus S}$ on n-qubit space, where $S \subseteq [n]$ is the set of k qubits which U acts on, and $I_{[n]\setminus S}$ is the identity operator on the remaining qubits.

We give a few important examples and fix some notation. Define the single qubit states $|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $|1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, $|+\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, and $|-\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$. For multi-qubit states which are a tensor product of state on fewer qubits, we often use one of two notations which omits the tensor product, e.g. $|00\rangle = |0\rangle|0\rangle = |0\rangle \otimes |0\rangle$, $|{+}{+}\rangle = |+\rangle|+\rangle = |+\rangle \otimes |+\rangle$, and $|{+}0\rangle = |+\rangle|0\rangle = |+\rangle \otimes |0\rangle$ are all two-qubit states. Define the single qubit gates $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$, $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$, and $\hat{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ (we put the hat to avoid confusion with a Hamiltonian H). Let

$$CX = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad \mathrm{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{2.1}$$

$$CCX = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}, \quad CCZ = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix} \tag{2.2}$$

denote the CNOT, CZ, SWAP, Toffoli, and CCZ gates.

The *Z basis* (or computational basis) of n-qubit space is $\{|z\rangle\}_{z\in\{0,1\}^n}$. The *X basis* (or Hadamard basis) of n-qubit space is $\{|x\rangle\}_{x\in\{+,-\}^n}$, where each x is a formal string of $+$ and $-$ symbols.

For a state $|\psi\rangle$, we use $\langle\psi|$ to denote the conjugate transpose of $|\psi\rangle$. As examples, $|\psi\rangle\langle\psi|$ is the projector onto the space spanned by $|\psi\rangle$, and $\langle\psi|\varphi\rangle$ is the inner product between states $|\psi\rangle$ and $|\varphi\rangle$. We use $U^\dagger$ to denote the adjoint of an operator.

A *quantum circuit* on n qubits is an n-qubit operation U which can be decomposed into a sequence of gates $U = G_m \cdots G_1$, where each G_i is a quantum operation from some fixed *gate set*; typically, this gate set contains operations which each acts nontrivially on just a small number of qubits.

After applying a quantum circuit, we glean information about the resulting state through a process called *measurement*, which probabilistically yields an outcome and also changes the state. Formally, a (projective) measurement is a complete set of orthogonal projectors, i.e. a set $\{P_i\}_i$ which satisfies $\sum_i P_i = I$, $P_i^2 = P_i$ for all i, and $P_iP_j = 0$ for all $i \neq j$. Applying this measurement to a state $|\psi\rangle$ yields outcome i with probability $\| P_i|\psi\rangle \|^2 = \langle\psi|P_i|\psi\rangle$, and upon measuring outcome i, the state collapses to the state $\frac{P_i|\psi\rangle}{\sqrt{\langle\psi|P_i|\psi\rangle}}$. As an example, a (full) Z basis measurement on an n-qubit state $|\psi\rangle$ is defined by the projectors $\{|z\rangle\langle z|\}_{z\in\{0,1\}^n}$. The outcome z is measured with probability $\langle\psi|z\rangle\langle z|\psi\rangle = |\langle z|\psi\rangle|^2$, after which the state collapses to $|z\rangle$.

More generally, we can consider measurements of an n-qubit state which only measure a subset of the qubits. For example, a measurement of just the first qubit in the Z basis is given by the projectors $\{|0\rangle\langle 0|_1 \otimes I, |1\rangle\langle 1|_1 \otimes I\}$, where the subscript indicates the qubit register 1, and I acts on the remaining qubits.

2.2.1 Quantum complexity theory

We now define the model of quantum verification of interest in this work. A *promise problem* is a pair (L_{yes}, L_{no}) of subsets of bitstrings $L_{yes}, L_{no} \subseteq \{0,1\}^\star$ which satisfies $L_{yes} \cap L_{no} = \varnothing$. Promise problems generalize languages: a language is just a promise problem which additionally satisfies $L_{yes} \cup L_{no} = \{0,1\}^\star$. A *quantum verifier* is a quantum circuit whose goal to determine whether a given string x is in L_{yes} or L_{no}, promised that one is the case. The circuit is given the input x of length n, a quantum proof state $|\psi\rangle$, and has access to a register of ancilla qubits all initialized to $|0\rangle$. The circuit is efficient in the sense that the number of qubits in $|\psi\rangle$, the number of ancilla qubits, and the number of gates in the circuit are all polynomial in n. The circuit also comes from a uniform circuit family $\{V_n\}$, i.e. there is an efficient deterministic algorithm which outputs a classical description of circuit V_n, given input 1^n.

After the circuit is applied, the final step of the verification is to measure the first qubit in the Z basis. We say the verifier accepts if the measurement outcome is 1 and rejects if the outcome is 0. The probability that the verifier V_n accepts is then

$$\|(|1\rangle\langle 1|_1 \otimes I)V_n|x\rangle|\psi\rangle|0\cdots 0\rangle_{\mathrm{anc}}\|^2 . \tag{2.3}$$

We now formally define the classes QMA_1 and QMA; we will mainly be interested in QMA_1 in this work.

Definition 2.2.1. A promise problem (L_{yes}, L_{no}) is in QMA_1 if there is a uniform family of efficient quantum circuits $\{V_n\}$ such that for any $x \in \{0,1\}^n$,

1. If $x \in L_{yes}$, then V_n accepts with probability 1.
2. If $x \in L_{no}$, then V_n accepts with probability $\leq 1/3$.

Importantly, the definition of QMA_1 is not known to be independent of the choice of gate set for the circuits. In this work, we choose the gate set $\mathcal{G} = \{X, CX, CCX, \hat{H} \otimes \hat{H}\}$ for QMA_1. When needing to explicitly refer to QMA_1 with this gate set (e.g. in comparison with other gate sets), we use the notation ${\mathsf{QMA}_1}^{\mathcal{G}}$.

We make two remarks about the chosen gate set.

Remark 2.2.2. First, $\mathcal{G}$ is computationally universal, i.e. any quantum circuit can be simulated by a circuit which only uses gates from $\mathcal{G}$. This is because gates in $\mathcal{G}$ allow us to implement CCX and $\hat{H}$ (which can be implemented by applying $\hat{H} \otimes \hat{H}$ on the desired qubit and an otherwise unused ancilla qubit); these two gates already form a computationally universal gate set [Shi03, Aha03].

Second, $\mathcal{G}$ is studied in a work of Rudolph [Rud25] which makes progress on the interesting open problem of whether QMA_1 has a universal gate set ($\mathcal{G}$ is referred to as $\mathcal{G}_2$ in their work). Among other results related to $\mathcal{G}$, they show that the problem of Gapped Clique Homology on weighted graphs (introduced by King and Kohler) is ${\mathsf{QMA}_1}^{\mathcal{G}}$-complete. This result gives a surprising connection between an important problem in computational topology and $\mathsf{QMA}^{\mathcal{G}}$, motivating our study of other properties of ${\mathsf{QMA}_1}^{\mathcal{G}}$.

The class QMA is defined in the same way as QMA_1, except the completeness parameter is $2/3$, i.e. for $x \in L_{yes}$, the verifier must accept with probability at least $2/3$. In contrast to QMA_1, the definition of QMA is independent of the choice of gate set, since the Solovay-Kitaev theorem [DN06] shows that any quantum gate can be efficiently approximated using gates from a universal gate set. This argument does not straightforwardly extend to QMA_1:

approximation of a gate is not enough, since this may not preserve the perfect completeness required of a QMA_1 protocol.

NP-completeness generalizes to QMA_1 and QMA in a straightforward way. A promise problem $K = (K_{yes}, K_{no})$ has an *efficient reduction* to a promise problem $L = (L_{yes}, L_{no})$ if there is an efficient deterministic algorithm A which maps a bitstring x to a bitstring $A(x)$ such that $A(x) \in L_{yes}$ if $x \in K_{yes}$, and $A(x) \in L_{no}$ if $x \in K_{no}$. A promise problem L is *QMA_1-hard* if for every $K \in \mathsf{QMA}_1$, there is an efficient reduction from K to L. A promise problem L is *QMA_1-complete* if L is in QMA_1 and L is QMA_1-hard. The definitions for QMA are analogous.

2.2.2 Hamiltonian complexity

In this section, we define the Quantum-k-Sat problem and introduce its basis-restricted variant. An n-qubit *Hamiltonian* is a Hermitian operator acting on n qubits. A *k-local operator* on n qubits is an operator which acts nontrivially on at most k of the qubits. The *energy* of a state $|\psi\rangle$ with respect to Hamiltonian H is $\langle\psi|H|\psi\rangle$. Note that this quantity is always real and non-negative, since H is Hermitian.

Definition 2.2.3. Let $k \geq 1$ and let $\mathcal{S}$ be a set of Hermitian k-local projectors. Quantum-k-Sat is a promise problem whose input is a classical description of a Hamiltonian $H = \sum_{i=1}^m h_i$, where each term h_i acts on n qubits and belongs to $\mathcal{S}$. H is a "yes" instance if there is an n-qubit state $|\psi\rangle$ such that $\langle\psi|H|\psi\rangle = 0$. H is a "no" instance if for all n-qubit states $|\psi\rangle$, $\langle\psi|H|\psi\rangle \geq 1$.

Note that the definition of Quantum-k-Sat implicitly depends on the choice of allowed projectors $\mathcal{S}$; this will be relevant in defining its basis-restricted variant.

As an illustrating example, we first understand how Quantum-k-Sat generalizes the NP-complete problem k-Sat. For simplicity we consider the case $k = 3$: let $\varphi = C_1 \wedge \cdots \wedge C_m$ be a 3-Sat instance on n variables $x_1, \ldots, x_n$. We construct an instance $H = \sum_{i=1}^m h_i$ of Quantum-3-Sat. Recall that clause C has the form $C = l_1 \vee l_2 \vee l_3$ for some literals l_j. Set $a_j \in [n]$ so that l_j is the variable x_{a_j} or its negation (we can assume that the a_j's are distinct). Let $z = z_1z_2z_3 \in \{0,1\}^3$ be the unique assignment to l_1, l_2, and l_3 which does not satisfy C. Define a 3-local n-qubit projector h corresponding to C by $h = |z\rangle\langle z|_{a_1,a_2,a_3} \otimes I$, i.e. h acts as the projector $|z_j\rangle\langle z_j|$ on qubit a_j for $j = 1, 2, 3$, and acts trivally on the remaining qubits. H is constructed by defining each h_i from clause C_i in this manner.

Note that for any assignment $a \in \{0,1\}^n$ for φ, if a satisfies C_i then $\langle a|h_i|a\rangle = 0$, and otherwise $\langle a|h_i|a\rangle = 1$. Thus, If φ is satisfiable with some assignment a, the state $|a\rangle$ is a zero-energy state of H. If φ is not satisfiable, then any Z basis state $|z\rangle$ satisfies $\langle z|H|z\rangle \geq 1$, and thus $\langle\psi|H|\psi\rangle \geq 1$ for any n-qubit state $|\psi\rangle$. This shows how 3-Sat can be viewed as a special case of Quantum-3-Sat.

Note that the Hamiltonian which arises from this transformation take on a highly restricted form: each projector term is diagonal in the same, predetermined basis (namely, the Z basis). Let us refer to this as a *classical Hamiltonian*. This brings into view one of the guiding questions of this work: if we relax the restrictions placed on a classical Hamiltonian, at what point does the Hamiltonian become "quantum"? This question has been previously studied for the particular relaxation in which the Hamiltonian is no longer required to be classical, but instead the Hamiltonian terms must pairwise commute (note that this indeed is a generalization, since the terms of a classical Hamiltonian pairwise commute). We discuss commuting Hamiltonians further in chapter 4.

In this work, we introduce a new way of generalizing classical Hamiltonians which we call *basis restriction*.

Definition 2.2.4. For $n \geq 1$, let $\mathcal{B}_n$ be a set of bases of n-qubit space. Let $\mathcal{B} = \cup_{n \geq 1} \mathcal{B}_n$. Define $\mathcal{B}$-Quantum-k-Sat to be the problem Quantum-k-Sat where we choose $\mathcal{S}$ (the set of allowed projectors in the definition) to be the set of Hermitian k-local projectors which are diagonal in some basis in $\mathcal{B}$.

As an example, if for all n we let $\mathcal{B}_n$ include just the Z basis on n qubits, then an instance of $\mathcal{B}$-Quantum-k-Sat is a Hamiltonian $H = \sum_i h_i$ for which each h_i is diagonal in the Z basis. Then H is a classical Hamiltonian and $\mathcal{B}$-Quantum-k-Sat is in NP.

We are interested in the complexity of basis-restricted Quantum-k-Sat when we allow for more than one "type" of basis (for example, when each h_i is diagonal in either the Z or the X basis). In particular, is there some setting in which the problem already becomes QMA-hard, even though the number of basis types is small? We address this question in the following chapter.

Chapter 3

QMA_1-Completeness of Basis-Restricted Quantum-k-Sat

3.1 Overview

In this chapter, we prove a QMA_1-hardness result for Quantum-6-Sat with the restriction that each Hamiltonian term is diagonal in a X/Z "mixed" basis. First, we formally define this problem.

Definition 3.1.1. For $n \geq 1$, let $\mathcal{B}_n^{XZ}$ be the set of bases of n-qubit space defined by

$$\mathcal{B}_n^{XZ} = \{\{\hat{H}^a|z\rangle\}_{z\in\{0,1\}^n}\}_{a\in\{0,1\}^n}, \tag{3.1}$$

where $\hat{H}^a$ is shorthand for $\hat{H}^{a_1} \otimes \cdots \otimes \hat{H}^{a_n}$.

For example, $\mathcal{B}_1^{XZ} = \{\{|0\rangle, |1\rangle\}, \{|+\rangle, |-\rangle\}\}$, and

$$\mathcal{B}_2^{XZ} = \{\{|z\rangle\}_{z\in\{0,1\}^2}, \{|z\rangle|x\rangle\}_{z\in\{0,1\},x\in\{+,-\}}, \{|x\rangle|z\rangle\}_{z\in\{0,1\},x\in\{+,-\}}, \{|x\rangle\}_{x\in\{+,-\}^2}\}. \tag{3.2}$$

Define $\mathcal{B}^{XZ} = \cup_{n\geq 1}\mathcal{B}_n^{XZ}$. We now state the main result of this chapter.

Theorem 3.1.2. *$\mathcal{B}^{XZ}$-Quantum-6-Sat is $\mathsf{QMA}_1{}^{\mathcal{G}}$-complete.*

3.2 A new gate set $\mathcal{G}^{XZ}$

Recall that $\mathcal{G} = \{X, CX, CCX, \hat{H} \otimes \hat{H}\}$ is our chosen gate set for QMA_1. We first show that $\mathcal{G}$ can be exactly simulated by another gate set $\mathcal{G}^{XZ}$ which has useful properties. Note that *exact* simulation is important for our purposes: an approximating some operation using gates from another gate set might not preserve the perfect completeness of a QMA_1 protocol.

We introduce some notation specific to this section: for a k-qubit gate U, let $U_{i_1,\ldots,i_k}$ denote U acting on qubits i_1 through i_k. For example, $\hat{H}_i$ is a $\hat{H}$ gate on qubit i, and $CX_{i,j}$ denotes a CNOT gate with control qubit i and target qubit j.

Definition 3.2.1. Let $\mathcal{G}^{XZ} = \{X, CX, CCX, CZ, \hat{G}\}$, where we define the two-qubit gate

$$\hat{G}_{i,j} = \hat{H}_i CX_{i,j} \hat{H}_i. \tag{3.3}$$

Lemma 3.2.2. *Each gate in $\mathcal{G}$ can be written as a (constant length) sequence of gates from $\mathcal{G}^{XZ}$.*

Proof. This is obvious for gates X, CX, and CCX. To simulate $\hat{H} \otimes \hat{H}$ acting on qubits i and j, we use the identities $\mathrm{SWAP}_{i,j} = CX_{i,j}CX_{j,i}CX_{i,j}$ and $CZ_{j,i} = \hat{H}_i CX_{j,i}\hat{H}_i$ to compute

$$
\begin{aligned}
(\hat{H} \otimes \hat{H})_{i,j} &= (\hat{H} \otimes \hat{H})_{i,j}\mathrm{SWAP}_{i,j}\mathrm{SWAP}_{i,j} && (3.4)\\
&= \hat{H}_i \mathrm{SWAP}_{i,j}\hat{H}_i \mathrm{SWAP}_{i,j} && (3.5)\\
&= \hat{H}_i CX_{i,j}CX_{j,i}CX_{i,j}\hat{H}_i \mathrm{SWAP}_{i,j} && (3.6)\\
&= \hat{H}_i CX_{i,j}\hat{H}_i\hat{H}_i CX_{j,i}\hat{H}_i\hat{H}_i CX_{i,j}\hat{H}_i \mathrm{SWAP}_{i,j} && (3.7)\\
&= \hat{G}_{i,j}CZ_{j,i}\hat{G}_{i,j}CX_{i,j}CX_{j,i}CX_{i,j}, && (3.8)
\end{aligned}
$$

which expresses $\hat{H} \otimes \hat{H}$ in terms of the gate set $\mathcal{G}^{XZ}$. □

The simulation in the other direction holds as well, if we are allowed to use an extra qubit.

Lemma 3.2.3. *Each gate in $\mathcal{G}^{XZ}$ can be written as a sequence of gates from $\mathcal{G}$ (provided that we have an additional qubit, which can be in any state and will be left unchanged after the simulation).*

Proof. We only need to consider the gates CZ and $\hat{G}$ in $\mathcal{G}^{XZ}$. Using an additional qubit k, we can write $CZ_{i,j}$ in the gate set $\mathcal{G}$ as

$$CZ_{i,j} \otimes I_k = (\hat{H} \otimes \hat{H})_{j,k}CX_{i,j}(\hat{H} \otimes \hat{H})_{j,k}, \tag{3.9}$$

and $\hat{G}_{i,j}$ as

$$\hat{G}_{i,j} \otimes I_k = (H \otimes H)_{i,j}(\hat{H} \otimes \hat{H})_{j,k}CX_{i,j}(\hat{H} \otimes \hat{H})_{j,k}(\hat{H} \otimes \hat{H})_{i,j}. \tag{3.10}$$

□

These two lemmas immediately imply the following result.

Lemma 3.2.4. $\mathsf{QMA}_1{}^{\mathcal{G}} = \mathsf{QMA}_1{}^{\mathcal{G}^{XZ}}$.

Proof. Any QMA_1 verification circuit consisting of gates from $\mathcal{G}^{XZ}$ can be exactly written as a circuit with gates in $\mathcal{G}$, by Lemma 3.2.3 (use an ancilla for the additional circuit). This shows that $\mathsf{QMA}_1{}^{\mathcal{G}^{XZ}} \subseteq \mathsf{QMA}_1{}^{\mathcal{G}}$. The reverse inclusion holds analogously by Lemma 3.2.2. □

We state the properties of $\mathcal{G}^{XZ}$ which will be used in the proof of Theorem 3.1.2.

Lemma 3.2.5. *Every gate in $\mathcal{G}^{XZ}$ is Hermitian and is diagonal in some basis in $\mathcal{B}^{XZ}$.*

Proof. It is straightforward to check that each gate is Hermitian. Table 3.1 gives, for every gate $U \in \mathcal{G}$, the change of basis matrix V which diagonalizes U. Each of the corresponding bases is in $\mathcal{B}^{XZ}$. □

U	V	$V^\dagger U V$
X_i	$\hat{H}_i$	Z_i
$CX_{i,j}$	$\hat{H}_j$	$CZ_{i,j}$
$CCX_{i,j,k}$	$\hat{H}_k$	$CCZ_{i,j,k}$
$CZ_{i,j}$	$I_{i,j}$	$CZ_{i,j}$
$\hat{G}_{i,j}$	$\hat{H}_i\hat{H}_j$	$CZ_{i,j}$

Table 3.1: Diagonalizing matrices for gates in $\mathcal{G}^{XZ}$.

We also give a useful property of the original gate set $\mathcal{G}$. For any gate U, let $\Gamma(U)$ denote the controlled-U gate.

Lemma 3.2.6. *For every gate U in $\mathcal{G}$, $\Gamma(U)$ can be exactly implemented using gates in $\mathcal{G}$ (provided that we have an additional qubit, which can be in any state and will be left unchanged after the simulation).*

Proof. $\Gamma(X)$ and $\Gamma(CX)$ are already included in $\mathcal{G}$. Let a denote the ancilla qubit. $\Gamma(CCX)$ on control qubits i,j,k and target qubit l can be decomposed as

$$\Gamma(CCX)_{i,j,k,l} \otimes I_a = CCX_{i,j,a}CCX_{k,a,l}CCX_{i,j,a}CCX_{k,a,l}. \tag{3.11}$$

$\Gamma(H \otimes H)$ with control qubit i, and target qubits j,k, we first give a decomposition using the $\Gamma(\text{SWAP})$ gate.

$$\Gamma(H \otimes H)_{i,j,k} \otimes I_a = \Gamma(\text{SWAP})_{i,j,k}(H \otimes H)_{k,a}\Gamma(\text{SWAP})_{i,j,k}(H \otimes H)_{k,a}. \tag{3.12}$$

Then, since $\Gamma(AB) = \Gamma(A)\Gamma(B)$ for any gates A and B, the identity $\text{SWAP}_{i,j} = CX_{i,j}CX_{j,i}CX_{i,j}$ gives a decomposition of $\Gamma(\text{SWAP})$ in terms of CCX's. This completes the proof. We illustrate these constructions in Figure 3.1.

□

Figure 3.1: Decomposition of $\Gamma(CCX)$ and $\Gamma(H \otimes H)$ into gates from $\mathcal{G}$.

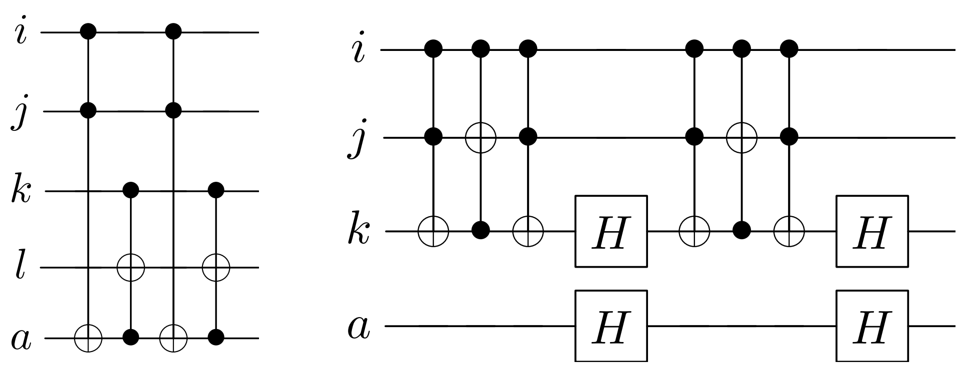

3.3 The Kitaev circuit-to-Hamiltonian construction

The proof of Theorem 3.1.2 uses the Kitaev circuit-to-Hamiltonian construction [KSV02]. which is a quantum analog of the fundamental construction used by Cook and Levin [Coo71, Lev73] to show that SAT is NP-complete. In this section, we review the relevant aspects of the Kitaev construction, referring the reader to [KSV02] for the full details.

The Kitaev construction is an efficient mapping from a quantum verification circuit to a Hamiltonian, such that the Hamiltonian a zero-energy ground state if and only if there is a proof state which makes the verification circuit accept with certainty. The mapping has the property that if each gate in the verification circuit acts on at most k qubits, then the resulting Hamiltonian is $(k+3)$-local. This was originally applied in the imperfect completeness setting to show the QMA-completeness of the 5-local Hamiltonian problem, which is a generalization of QUANTUM-5-SAT in which the goal is to decide whether a Hamiltonian has a *low*-energy state. We will see that same construction applies to the perfect completeness setting we are interested in.

We establish some notation which will be reused in the proof of Theorem 3.1.2. Let V be a quantum circuit for verifiying $x \in \{0,1\}^n$ given a proof state $|\psi\rangle$. Recall that the verification starts from initial state $|\text{init}\rangle = |x\rangle|\psi\rangle|0\cdots0\rangle_{\text{anc}}$; let $\mathcal{H}$ denote the Hilbert space in which this state lives. Let p and q be the number of proof and ancilla qubits respectively (both are $\text{poly}(n)$). Let $Q_{\text{input}} = [n]$, $Q_{\text{proof}} = \{n+1,\ldots,n+p\}$, and $Q_{\text{anc}} = \{n+p+1,\ldots,n+p+q\}$ denote the indices of the qubits in each register. Write the circuit as $V = U_T\cdots U_1$, where each U_i is a gate from the gate set.

The Kitaev construction works by building a Hamiltonian H such that if the verification circuit can be made to accept, H has a zero-energy state called a *history state* which encodes all of the circuit's computations. The Hamiltonian acts on the space $\mathcal{H}\otimes\mathcal{H}_{\text{clock}}$, where $\mathcal{H}_{\text{clock}} = (\mathbb{C}^2)^{\otimes T}$ is an additional T-qubit *clock register*. Let $|\hat{0}\rangle,\ldots,|\hat{T}\rangle \in \mathcal{H}_{\text{clock}}$ be *clock states* (we define these states exactly in a moment). The history state is

$$|\text{hist}\rangle = \frac{1}{\sqrt{T+1}}\sum_{t=0}^{T} U_t\cdots U_1|\text{init}\rangle\otimes|\hat{t}\rangle. \tag{3.13}$$

The Hamiltonian H takes the form

$$H = H_{\text{in}} + H_{\text{out}} + H_{\text{prop}} + H_{\text{clock}}. \tag{3.14}$$

Intuitively, each term in the Hamiltonian represents a constraint on the proof state (which is claimed to be the history state); an energy penalty is given if the constraint is not met.

The first term is

$$H_{\text{in}} = \sum_{i\in Q_{\text{input}}} |\overline{x_i}\rangle\langle\overline{x_i}|_i\otimes|\hat{0}\rangle\langle\hat{0}| + \sum_{i\in Q_{\text{anc}}} |1\rangle\langle1|_i\otimes|\hat{0}\rangle\langle\hat{0}|, \tag{3.15}$$

where $\overline{x_i}$ is the negation of the i-th bit of x, and $|\overline{x_i}\rangle\langle\overline{x_i}|_i$ denotes the projector on $\mathcal{H}$ which acts as $|\overline{x_i}\rangle\langle\overline{x_i}|$ on qubit i and as identity on the remaining qubits. H_{in} checks the constraint that when the clock state is $|\hat{0}\rangle$, the input register is x and the ancilla register is all 0's (note that $\langle\text{hist}|H_{\text{in}}|\text{hist}\rangle = 0$).

The next term

$$H_{\text{out}} = |0\rangle\langle 0|_1 \otimes |\hat{T}\rangle\langle\hat{T}| \tag{3.16}$$

checks that when the clock state is $|\hat{T}\rangle$, the first qubit is $|1\rangle$. This corresponds to V accepting after measuring the first qubit at the end of its computation.

The propagation term has the form $H_{\text{prop}} = \sum_{t=1}^{T} H_{\text{prop},t}$, with

$$H_{\text{prop},t} = \frac{1}{2}\left(I_{\mathcal{H}} \otimes |\widehat{t-1}\rangle\langle\widehat{t-1}| + I_{\mathcal{H}} \otimes |\hat{t}\rangle\langle\hat{t}| - U_t \otimes |\hat{t}\rangle\langle\widehat{t-1}| - U_t^\dagger \otimes |\widehat{t-1}\rangle\langle\hat{t}|\right), \tag{3.17}$$

where U_t and $U_t^\dagger$ act on $\mathcal{H}$. It is straightforward to verify that $\langle\text{hist}|H_{\text{prop}}|\text{hist}\rangle = 0$. It is also easy to check that $H_{\text{prop},t}$ is a projector.

Before writing H_{clock}, we specify the exact definition of the clock states. The state $|\hat{t}\rangle = |1^t 0^{T-t}\rangle$ is a unary encoding of t. The clock term

$$H_{\text{clock}} = \sum_{i=1}^{T-1} I_{\mathcal{H}} \otimes |01\rangle\langle 01|_{i,i+1} \tag{3.18}$$

enforces that the state on $\mathcal{H}_{\text{clock}}$ is only supported on the clock states, by giving an energy penalty to states which have a 0 before a 1 on the clock register.

Remark 3.3.1. We comment on the locality of H. By the definition of the clock states, the projectors onto the clock space in H_{in}, H_{out} and H_{prop} can be replaced with local checks which find the location of the first 0 in the clock register. For example, $|\hat{t}\rangle\langle\hat{t}|$ can be replaced with the 2-local projector $|10\rangle\langle 10|_{t,t+1}$ for $t = 1, \ldots, T-1$, $|\hat{0}\rangle\langle\hat{0}|$ with the 1-local projector $|0\rangle\langle 0|_1$, and $|\hat{T}\rangle\langle\hat{T}|$ with the 1-local projector $|1\rangle\langle 1|_T$.

Thus, each projector term in H_{in} andH_{out} is 2-local (one from $\mathcal{H}$ and one from the clock). The terms of H_{clock} are also 2-local. The locality of H then hinges on the locality of the remaining projector terms $H_{\text{prop},t}$. A 3-local check suffices for the clock in $H_{\text{prop},t}$, so if each gate in the gate set acts on at most k qubits, H is $(k+3)$-local.

3.4 Proof of Theorem 3.1.2

By the definition of QMA_1-completeness, Theorem 3.1.2 is proven by the following two propositions.

Proposition 3.4.1. *$\mathcal{B}^{XZ}$-Quantum-6-Sat is in $\mathsf{QMA}_1{}^{\mathcal{G}}$.*

Proof. First, without loss of generality, we can build a $\mathsf{QMA}_1{}^{\mathcal{G}}$ verification circuit which allows intermediate Z basis measurements and quantum operations conditioned on the measurement results. This is because a conditionally applied gate U can be replaced by the controlled operation $\Gamma(U)$, removing the intermediate measurement. Lemma 3.2.6 then says that $\Gamma(U)$ can then be decomposed back into the gate set $\mathcal{G}$. Thus, a circuit with intermediate measurements can be rewritten as a circuit with just a single measurement at the end of the computation, as in our definition of QMA_1. Also note that classical processing can be performed freely, since our gate set includes CCX, which is universal for classical computation.

We use the verification procedure of Gosset and Nagaj [GN16], which is correct as long as the following condition holds:

(*) There is an efficient quantum algorithm using gates in $\mathcal{G}$ which, given a projector $\Pi \in \mathcal{S}$, exactly performs the projective measurement $\{I - \Pi, \Pi\}$ on a given state $|\psi\rangle$.

Recall that $\mathcal{S}$ is the set of Hermitian 6-local projectors which are diagonal in a basis in $\mathcal{B}^{XZ}$.

We briefly sketch how the QMA_1 protocol works assuming (*) holds. First, choose a random projector term Π in the Hamiltonian H (random bits can be obtained by applying X to a $|0\rangle$ ancilla then measuring). Then, perform the projective measurement in (*) on the proof state $|\psi\rangle$. Accept when the outcome is $I - \Pi$, and reject otherwise. The probability of accepting is $1 - \langle\psi|\Pi|\psi\rangle$. When H has a zero-energy state, the protocol accepts this state with probability 1, showing perfect completeness. If instead $\langle\psi|H|\psi\rangle \geq 1$ for all $|\psi\rangle$, the protocol will reject with at least $1/\operatorname{poly}(n)$ probability, which can be amplified to give the desired soundness; the full analysis is given in [GN16].

We now prove that (*) holds. A projector $\Pi \in \mathcal{S}$ can be specified by a set of six qubits S on which it acts nontrivially, a change of basis matrix V (which is a tensor product of H and I), and a set of supported strings $S \subseteq \{0,1\}^6$, such that

$$\Pi = \sum_{z \in S} V|z\rangle\langle z|V. \tag{3.19}$$

The algorithm first applies V on $|\psi\rangle$, by applying $\hat{H}$ on each of the specified qubits (using the $\hat{H} \otimes \hat{H}$ gate, with the second $\hat{H}$ acting on an ancilla qubit). It then measures the qubits in Q in the Z basis. Let $z^\star$ be the measurement outcome. The algorithm returns outcome Π if $z^\star \in S$, and otherwise returns outcome $I - \Pi$. The probability of getting outcome Π is

$$\sum_{z \in S} \left\| |z\rangle\langle z|V|\psi\rangle \right\|^2 = \langle\psi|\Pi|\psi\rangle, \tag{3.20}$$

as desired. This completes the proof. □

Proposition 3.4.2. *$\mathcal{B}^{XZ}$-Quantum-6-Sat is $\mathsf{QMA}_1^{\mathcal{G}}$-hard.*

Proof. By Lemma 3.2.4, it suffices to show that any QMA_1 verifier V with gates in $\mathcal{G}^{XZ}$ can be efficiently mapped into an instance of $\mathcal{B}^{XZ}$-Quantum-6-Sat in a way which preserves yes and no instances. Let $H_V = H_{\text{in}} + H_{\text{out}} + H_{\text{prop}} + H_{\text{clock}}$ be the Kitaev Hamiltonian for V, as defined in section 3.3. If V accepts with probability 1 on some proof state $|\psi\rangle$, then the state $|\text{hist}\rangle$ (eq. 3.13) is a zero-energy state of H_V (since $\langle\text{hist}|H_{\text{out}}|\text{hist}\rangle = 0$), i.e. H_V is a yes instance of $\mathcal{B}^{XZ}$-Quantum-6-Sat. The analysis in the no case is the same as in Kitaev's proof that 5-local Hamiltonian is QMA-hard [KSV02], since QMA_1 and QMA have the same soundness.

We now show that H_V is indeed an instance of $\mathcal{B}^{XZ}$-Quantum-6-Sat. Each Hamiltonian term is a projector, and from Remark 3.3.1, the terms are 6-local, since each gate in $\mathcal{G}^{XZ}$ acts on at most three qubits.

Finally, we check the basis restriction condition. It is clear by definition that each term in H_{in}, H_{out}, and H_{clock} is diagonal in the Z basis. It remains to show that each projector term $H_{\text{prop},t}$ is diagonal in a basis in $\mathcal{B}^{XZ}$. Explicitly writing out the 3-local check on qubits $t-1$, t, and $t+1$ of $\mathcal{H}_{\text{clock}}$ gives

$$
\begin{aligned}
H_{\text{prop},t} = \frac{1}{2} \big(& I_{\mathcal{H}} \otimes |100\rangle\langle 100|_{t-1,t,t+1} + I_{\mathcal{H}} \otimes |110\rangle\langle 110|_{t-1,t,t+1} \\
& -U_t \otimes |110\rangle\langle 100|_{t-1,t,t+1} - U_t^\dagger \otimes |100\rangle\langle 110|_{t-1,t,t+1} \Big) \qquad (3.21) \\
= |1\rangle\langle 1|_{t-1} & \otimes |0\rangle\langle 0|_{t+1} \otimes \frac{1}{2} \Big(I_{\mathcal{H}} \otimes |0\rangle\langle 0|_t + I_{\mathcal{H}} \otimes |1\rangle\langle 1|_t - U_t \otimes |1\rangle\langle 0|_t - U_t^\dagger \otimes |0\rangle\langle 1|_t \Big) . \qquad (3.22)
\end{aligned}
$$

By Lemma 3.2.5 we have $U_t = U_t^\dagger$, so equation 3.22 becomes

$$
|1\rangle\langle 1|_{t-1} \otimes |0\rangle\langle 0|_{t+1} \otimes \frac{1}{2} \left(I_{\mathcal{H}} \otimes I_t - U_t \otimes X_t \right) \qquad (3.23)
$$

where I_t and X_t are the I and X operators on qubit t. In this form, the basis which diagonalizes $H_{\text{prop},t}$ is clear: the change of basis matrix is VH_t, where V is the change of basis matrix for U_t given in table 3.1. For example, if U_t is $CX_{i,j}$, the corresponding change of basis matrix acts as $\hat{H}$ on qubits i and j in $\mathcal{H}$, $\hat{H}$ on qubit t in $\mathcal{H}_{\text{clock}}$, and identity on all other qubits. Every such basis is in $\mathcal{B}^{XZ}$, which completes the proof. $\square$

Chapter 4

Commuting Basis-Restricted Qudit Hamiltonians in NP

4.1 Commuting Hamiltonians

In this chapter, we study the complexity of Hamiltonians which combine basis restriction with the restriction of *commutativity*. We say a Hamiltonian $H = \sum_i h_i$ is *commuting* if its terms pairwise commute, i.e. $h_i h_j = h_j h_i$ for all $i \neq j$. There is a large existing body of work on the complexity of commuting Hamiltonians, which, much like the basis-restricted Hamiltonian, is a system that is seemingly more classical than a general Hamiltonian. These results are almost all "easiness" theorems, i.e. NP-inclusions for commuting Hamiltonian problems. Among them is the following theorem of Bravyi and Vyalyi:

Theorem 4.1.1 ([BV05], Section 7)**.** *The following problem is in* NP*: given a commuting Hamiltonian $H = \sum_{i=1}^{m} h_i$ acting on n qubits, where each h_i is an orthogonal projector and is factorized (i.e. is a tensor product of single-qubit operations), decide whether H has a zero-energy state.*

We make a few remarks about this problem.

Remark 4.1.2.

1. In [BV05] this problem is referred to as the Factorized Projectors Common Eigenspace problem, and the notation is slightly different. We choose the formulation above to stress that the problem is a variant of Quantum-Sat where the projectors commute and are factorized.

2. The Hamiltonian in the problem does not need to be local. This is a stronger result than with the locality restriction, because it puts Quantum-Sat with a larger class of Hamiltonians into NP.

3. The factorized condition is similar to basis restriction if we consider bases with single-qubit tensor structure. Theorem [BV05] implies that commuting $\mathcal{B}$-Quantum-Sat is in NP if the bases vectors in each basis $\mathcal{B}$ are tensor products of single qubits states. For example, commuting $\mathcal{B}^{XZ}$-Quantum-k-Sat is in NP, in contrast to the QMA_1-hardness result shown in chapter 3 in the non-commuting case.

It is still an open question whether the same NP inclusion holds if the Hamiltonian acts on *qudits* with local dimension greater than 2, since the techniques of [BV05] use properties specific to qubits. In this chapter, we develop techniques which work for qudits, independent of local dimension, and show that a particular basis-restricted commuting qudit Hamiltonian problem is in NP.

4.2 Qudits and general bases

We introduce qudits and set some notation about bases, in order to state our main theorem in full generality. Throughout this chapter, we fix any integer $d \geq 2$ and define $\mathcal{D} = \{0, \ldots, d-1\}$. A *qudit* with local dimension d is a unit vector in $\mathbb{C}^d$. An n-qudit state is a unit vector in $(\mathbb{C}^d)^{\otimes n}$. A quantum operation on n qudits is a unitary operator acting on $(\mathbb{C}^d)^{\otimes n}$. We denote the computational basis on n qudits by $\{|z\rangle\}_{z \in \mathcal{D}^n}$. The other usual conventions for qubits carries over to qudits in a straightforward manner.

The bases we consider will have single-*qudit* tensor structure, where each qudit is in one of two bases $\mathcal{Q}$ or $\mathcal{R}$. We define $\mathcal{Q}$ in terms of its unitary change of basis matrix Q from the computational basis, i.e. $\mathcal{Q} = \{Q|z\rangle\}_{z \in \mathcal{D}}$. Likewise, we define $\mathcal{R} = \{R|z\rangle\}_{z \in \mathcal{D}}$ for some unitary matrix R.

Definition 4.2.1. For $n \geq 1$, let $\mathcal{B}_n^{QR}$ be the set of bases of n-qudit space defined by

$$\mathcal{B}_n^{QR} = \{\{Q^{a_1}R^{1-a_1} \otimes \cdots \otimes Q^{a_n}R^{1-a_n}|z\rangle\}_{z \in \mathcal{D}^n}\}_{a \in \{0,1\}^n}, \tag{4.1}$$

i.e. $\mathcal{B}_n^{QR}$ contains the bases where each qudit is either in the $\mathcal{Q}$ or $\mathcal{R}$ bases. We refer to these bases as *mixed bases*.

Definition 4.2.2. For $n \geq 1$, define the *pure bases*

$$B_n^Q = \{Q^{\otimes n}|z\rangle\}_{z \in \mathcal{D}^n}, \qquad B_n^R = \{R^{\otimes n}|z\rangle\}_{z \in \mathcal{D}^n}. \tag{4.2}$$

We now state the main theorem of this chapter.

Definition 4.2.3. We define the language d-(Q, R, QR)-Comm-QSat. The input is a (local dimension d) n-qudit commuting Hamiltonian, given as a classical description

$$\{\{h_i^Q\}_i, \{h_j^R\}_j, \{h_k^S\}_k\}, \tag{4.3}$$

where the all the terms are orthogonal projectors, the terms h_i^Q are each diagonal in the basis B_n^Q, the terms h_j^R are each diagonal in the basis B_n^R, and the terms h_k^S are each diagonal in the *same* basis in $\mathcal{B}_n^{QR}$. The Hamiltonian corresponding to this description is $H = \sum_i h_i^Q + \sum_j h_j^R + \sum_k h_k^S$. The decision problem is: does H have a zero-energy state?

Theorem 4.2.4. *Suppose the bases $\mathcal{Q}$ and $\mathcal{R}$ satisfy the following conditions.*

1. *For all $|q\rangle \in \mathcal{Q}$ and $|r\rangle \in \mathcal{R}$, $\langle q|r\rangle \neq 0$.*

2. *Q and R are real-valued matrices.*

Then d-(Q, R, QR)-Comm-QSat is in NP.

4.3 Trace calculation and examples

Before proving the theorem in its most general form, we sketch the proof idea and give two simple examples to illustrate key points. The starting point of our proof is the following nice expression for the ground space projector of a commuting Hamiltonian with projector terms.

Lemma 4.3.1. *Let $H = \sum_j h_j$ be a commuting qudit Hamiltonian with projector terms. Then the projector onto the zero-energy eigenspace of H can be written as*

$$\prod_j (I - h_j). \tag{4.4}$$

Proof. Since the h_j's commute (and are diagonalizable), they have a simultaneous eigenbasis, which we denote by $B = \{|\varphi_i\rangle\}_i$. The projector onto the zero-energy eigenspace of H is, by definition, the linear operator P which has the following action on the basis B: $P|\varphi_i\rangle = |\varphi_i\rangle$ if $H|\varphi_i\rangle = 0$, and $P|\varphi_i\rangle = 0$ if $H|\varphi_i\rangle \neq 0$.

We show that $\prod_j(I - h_j) = P$ by showing that it acts the same on B. Suppose $H|\varphi_i\rangle = 0$. Then $h_j|\varphi_i\rangle = 0$ for all j, so $\prod_j(I - h_j)|\varphi_i\rangle = |\varphi_i\rangle$.

On the other hand, if $H|\varphi_i\rangle \neq 0$, then there is some k such that $h_k|\varphi_i\rangle \neq 0$. But $|\varphi_i\rangle$ is an eigenstate of h_k, and it can only have eigenvalue 0 or 1 since h_k is a projector. Thus, it is a 1-eigenstate, and $h_k|\varphi_i\rangle = |\varphi_i\rangle$. Then $(I - h_k)|\varphi_i\rangle = 0$, so $\prod_j(I - h_j)|\varphi_i\rangle = 0$ by commuting $I - h_k$ to the right. This completes the proof. □

Central to our proof is the idea of using a *trace calculation* to give an efficient witness that there is a zero-energy state.

Lemma 4.3.2. *Let H be a qudit Hamiltonian and let P be the projector onto the zero-energy eigenspace of H. Then $\operatorname{tr} P \neq 0$ if and only if H has a zero-energy state.*

Proof. H has a zero-energy state if and only if the dimension of the zero-energy eigenspace of H is not 0. But $\operatorname{tr} P$ calculates the sum of the eigenvalues of P, which is just the dimension of the zero-energy eigenspace of H, since P is a projector. □

Example 1. Consider a n-qubit commuting Hamiltonian $H = H_Z + H_X$, where the projector terms in $H_Z = \sum_i h_i^Z$ are diagonal in the pure Z basis, and those in $H_X = \sum_j h_j^X$ are diagonal in the pure X basis. We wish to find a classical proof that H has a zero-energy state. We define $P_Z = \prod_i(I - h_i^Z)$ and $P_X = \prod_i(I - h_i^X)$ Combining Lemmas 4.3.1 and 4.3.2 gives us that H has a zero-energy state if and only if $\operatorname{tr}(P_Z P_X) \neq 0$.

Now, note that P_Z is also the projector onto the zero-energy eigenspace of H_Z. Defining the "good" Z strings $G_Z = \{z \in \{0,1\}^n : H_Z|z\rangle = 0\}$, we can write $P_Z = \sum_{z \in G_Z} |z\rangle\langle z|$.

Defining $G_X = \{x \in \{+,-\}^n : H_X|x\rangle = 0\}$ analogously, we can rewrite the trace as

$$\operatorname{tr}(P_Z P_X) = \sum_{z\in G_Z, x\in G_X} \operatorname{tr}(|z\rangle\langle z|x\rangle\langle x|) \tag{4.5}$$

$$= \sum_{z\in G_Z, x\in G_X} \operatorname{tr}(\langle z|x\rangle\langle x|z\rangle) \tag{4.6}$$

$$= \sum_{z\in G_Z, x\in G_X} |\langle z|x\rangle|^2 . \tag{4.7}$$

We now highlight the role of condition 1 from Theorem 4.2.4. Since every Z basis state has nonzero overlap with every X basis state, Each term in the sum 4.7 is positive. Thus, $\operatorname{tr}(P_Z P_X) \neq 0$ if and only if G_Z and G_X are both non-empty. We have arrived at our NP proof: give satisfying strings $z \in G_Z$ and $x \in G_x$.

Example 2. Now suppose we have a commuting Hamiltonian which is the Hamiltonian from the previous example, plus one additional mixed basis term $h = |ab\rangle\langle ab|_{1,2}$ acting on qubits 1 and 2, for some $a \in \{+,-\}$ and $b \in \{0,1\}$. We use M to denote the basis of h (X basis on the first qubit and Z on the rest, for example).

One natural candidate for an NP proof is a strings $z \in G_Z$ and $x \in G_X$ such that the "slice" of z and x corresponding to basis M satisfies constraint h, i.e. $h|x_1 z_2\rangle = 0$. By generalizing the trace computation and analyzing the slice operation, we can show that this candidate works.

Let $P_M = I - h$. The desired trace now becomes

$$\operatorname{tr}(P_Z P_X P_M) = \operatorname{tr}(P_Z P_M P_X P_M) \tag{4.8}$$

$$= \sum_{z\in G_Z, x\in G_X} \operatorname{tr}(|z\rangle\langle z|P_M|x\rangle\langle x|P_M) \tag{4.9}$$

$$= \sum_{z\in G_Z, x\in G_X} |\langle x|P_M|z\rangle|^2 . \tag{4.10}$$

where we used in (4.8) that $P_M^2 = P_M$ and the Hamiltonian terms commute.

We now wish to show that for all $z \in \{0,1\}^n$ and $x \in \{+,-\}^n$, $\langle x|P_M|z\rangle \neq 0$ if and only if the slice condition above holds. The slice condition says that $\langle a|x_1\rangle\langle b|z_2\rangle = 0$, i.e. either $x_1 \neq a$ or $z_2 \neq b$. In one direction, if the slice condition holds, then either $P_M|x\rangle = |x\rangle$, (if $x_1 \neq a$), or $P_M|z\rangle = |z\rangle$ (if $z_2 \neq b$). In both cases, we get $\langle x|P_M|z\rangle = \langle z|x\rangle \neq 0$, as desired.

The other direction is slightly more interesting. If the slice condition does not hold, then $x_1 = a$ and $z_2 = b$. We compute

$$\langle x|P_M|z\rangle = \langle x_1 \cdots x_n|(I-h)|z_1 \cdots z_n\rangle \tag{4.11}$$

$$= \langle x|z\rangle - \langle a x_2 \cdots x_n|ab\rangle_{1,2}\langle ab|_{1,2}|z_1 b z_3 \cdots z_n\rangle \tag{4.12}$$

$$= \langle x|z\rangle - \langle x_2|b\rangle\langle x_3 \cdots x_n|\langle a|z_1\rangle|z_3 \cdots z_n\rangle \tag{4.13}$$

$$= \langle x|z\rangle - \langle x|z\rangle \tag{4.14}$$

$$= 0. \tag{4.15}$$

This completes the argument for the NP witness.

4.4 Slicing qudits

We now generalize the argument in the previous example. First, we fix some notation and formally define the notion of slices for qudits. We will use $a \in \{0,1\}^n$ to refer to the mixed basis in $\mathcal{B}_n^{QR}$ given by

$$\{Q^{a_1}R^{1-a_1} \otimes \cdots \otimes Q^{a_n}R^{1-a_n}|z\rangle\}_{z\in\mathcal{D}^n}. \tag{4.16}$$

We use $q \in \mathcal{D}^n$ as the classical description of the basis state $|\hat{q}\rangle = Q^{\otimes n}|q\rangle \in B_n^Q$, where we use the hat notation simply to distinguish $|\hat{q}\rangle$ from a computational basis state. Likewise, $r \in \mathcal{D}^n$ will be the classical description of the basis state $|\breve{r}\rangle = R^{\otimes n}|r\rangle \in B_n^R$, with the breve notation.

Definition 4.4.1. Fix any mixed basis $a \in \{0,1\}^n$. For $q, r \in \mathcal{D}^n$, let $|\mathrm{slice}_a(q,r)\rangle$ be the n-qudit state which is a tensor product of n qudits, where the i-th qudit is $Q|q_i\rangle$ if $a_i = 1$, and $R|r_i\rangle$ if $a_i = 0$.

We prove a key technical lemma which will be used in conjunction with slicing. We index into a $d^n \times d^n$ matrix M with elements in $\mathcal{D}^n$, i.e. $M_{i_1\cdots,i_n,j_1\cdots j_n}$ is the element at row $i_1\cdots i_n$ and column $j_1\cdots j_n$, where $i_1\cdots i_n$ and $j_1\cdots j_n$ are in $\mathcal{D}^n$.

Definition 4.4.2. Let M be a $d^n \times d^n$ matrix. For $k \in [n]$, we define M^{Γ_k} to by the $d^n \times d^n$ matrix whose elements are given by

$$M^{\Gamma_k}_{i_1\cdots i_n,j_1\cdots j_n} = M_{i_1\cdots i_{k-1}j_k i_{k+1}\cdots i_n,j_1\cdots j_{k-1}i_k j_{k+1}\cdots j_n}. \tag{4.17}$$

For a set $T = \{t_1,\ldots,t_m\} \subseteq [n]$, define

$$M^{\Gamma_T} = (((M^{\Gamma_{t_1}})^{\Gamma_{t_2}})\cdots^{\Gamma_{t_m}}). \tag{4.18}$$

It is clear from the definition of M^{Γ_k} that M^{Γ_S} is well-defined (i.e. the order does not matter).

We remark that the operation $(\cdot)^{\Gamma_S}$ is called the *partial transpose* and has other applications in quantum information, for example in the positive positive transpose criterion in separability testing [Per96, HHH96].

Lemma 4.4.3. *Let P be an operator on n qudits which is diagonal in a basis in $\mathcal{B}^{QR}$. Assume that Q and R are real-valued matrices. Then for any $T \subseteq [n]$, $P^{\Gamma_T} = P$.*

Proof. By the definition of $(\cdot)^{\Gamma_T}$, it suffices to show that $P^{\Gamma_m} \in [n]$ for a single $m \in [n]$. Without loss of generality, we show it for $m = 1$.

Since P is diagonal in a basis in $\mathcal{B}^{QR}$, we can write P as

$$P = ((U^{(1)})^\dagger \otimes \cdots \otimes (U^{(n)})^\dagger)D(U^{(1)} \otimes \cdots \otimes U^{(n)}), \tag{4.19}$$

where $U^{(m)} \in \{Q, R\}$ is the change of basis matrix acting on qubit m, and D is a diagonal n-qudit operator.

We show that P and P^{Γ_1} have the same matrix elements. Fix any row $i = i_1 \cdots i_n$ and column $j = j_1 \cdots j_n$. Since D is diagonal, we can write

$$P_{i,j} = \sum_{k \in \mathcal{D}^n} (U^{(1)\dagger} \otimes \cdots \otimes U^{(n)\dagger})_{i,k} D_{k,k} (U^{(1)} \otimes \cdots \otimes U^{(n)})_{k,j} \tag{4.20}$$

$$= \sum_{k \in \mathcal{D}^n} \left(\prod_{m=1}^{n} U^{(m)\dagger}_{i_m,k_m} \right) D_{k,k} \left(\prod_{m=1}^{n} U^{(m)}_{k_m,j_m} \right). \tag{4.21}$$

Next, we have

$$P^{\Gamma_1}_{i,j} = P_{j_1 i_2 \cdots i_n, i_1 j_2 \cdots j_n} \tag{4.22}$$

$$= \sum_{k \in \mathcal{D}^n} U^{(1)\dagger}_{j_1,k_1} \left(\prod_{m=2}^{n} U^{(m)\dagger}_{i_m,k_m} \right) D_{k,k} U^{(1)}_{k_1,i_1} \left(\prod_{m=2}^{n} U^{(m)}_{k_m,j_m} \right). \tag{4.23}$$

But, denoting complex conjugate by $(\cdot)^\star$, we have $U^{(1)\dagger}_{j_1,k_1} = (U^{(1)})^\star_{k_1,j_1} = U^{(1)}_{k_1,j_1}$, since Q and R are assumed to be real-valued. Likewise, $U^{(1)}_{k_1,i_1} = (U^{(1)\dagger})^\star_{i_1,k_1} = (U^{(1)\dagger})_{i_1,k_1}$. So (4.21) and (4.23) are equal, which completes the proof. □

We remark that the matrix element $M_{i_1 \cdots i_n, j_1 \cdots j_n}$ in our standard notation, which we now return to, is just $\langle i_1 \cdots i_n | M | j_1 \cdots j_n \rangle$. Next, we show a straightforward property of the partial transpose.

Lemma 4.4.4. *Let M be an operator on n qudits and suppose $|u_1\rangle, \ldots, |u_n\rangle, |v_1\rangle, \ldots, |v_n\rangle$ are real-valued single-qudit states. Then for any $k \in [n]$,*

$$\langle u_1 | \cdots \langle u_n | M | v_1 \rangle \cdots | v_n \rangle = \langle u_1 | \cdots \langle u_{k-1} | \langle v_k | \langle u_{k+1} | \cdots \langle u_n | M^{\Gamma_k} | v_1 \rangle \cdots | v_{k-1} \rangle | u_k \rangle | v_{k+1} \rangle \cdots | v_n \rangle. \tag{4.24}$$

Proof. Without loss of generality, we prove it for $k = 1$. To simplify notation, let $|u'\rangle = |u_2\rangle \cdots |u_n\rangle$ and let $|v'\rangle = |v_2\rangle \cdots |v_n\rangle$. In the computational basis, $|u_1\rangle = \sum_{i \in \mathcal{D}} b_i |i\rangle$, $|v_1\rangle = \sum_{i \in \mathcal{D}} c_i |i\rangle$, $|u'\rangle = \sum_{i_2 \cdots i_n \in \mathcal{D}^{n-1}} b_{i_2 \cdots i_n} |i_2 \cdots i_n\rangle$, and $|v'\rangle = \sum_{i_2 \cdots i_n \in \mathcal{D}^{n-1}} c_{i_2 \cdots i_n} |i_2 \cdots i_n\rangle$. Then, by the definition of partial transpose and using that $|u_1\rangle$ and $|v_1\rangle$ are real-valued,

$$\langle u_1 | \cdots \langle u_n | M | v_1 \rangle \cdots | v_n \rangle = \sum_{i_1, j_1, i_2 \cdots i_n, j_2 \cdots j_n} b^\star_{i_1} b^\star_{i_2 \cdots i_n} c_{j_1} c_{j_2 \cdots j_n} \langle i_1 i_2 \cdots i_n | M | j_1 j_2 \cdots j_n \rangle \tag{4.25}$$

$$= \sum_{i_1, j_1, i_2 \cdots i_n, j_2 \cdots j_n} b_{i_1} b^\star_{i_2 \cdots i_n} c^\star_{j_1} c_{j_2 \cdots j_n} \langle j_1 i_2 \cdots i_n | M^{\Gamma_1} | i_1 j_2 \cdots j_n \rangle \tag{4.26}$$

$$= \langle v_1 | \langle u_2 | \cdots \langle u_n | M^{\Gamma_1} | u_1 \rangle | v_2 \rangle \cdots | v_n \rangle. \tag{4.27}$$

□

4.5 Proof of Theorem 4.2.4

Proof of Theorem 4.2.4. We first set some notation. Let $a \in \{0,1\}^n$ correspond to the mixed basis in which all terms h^S_k are diagonal. For $\square = Q, R, S$ define $H_\square = \sum_i h^\square_i$ and $P_\square = \prod_i (I - h^\square_i)$. Let $G_Q = \{q \in \mathcal{D}^n : H_Q |\hat{q}\rangle = 0\}$ and $G_R = \{r \in \mathcal{D}^n : H_R |\check{r}\rangle = 0\}$.

The NP proof is a pair of strings $q, r \in \mathcal{D}^n$. The verification accepts if and only if $q \in G_Q$, $r \in G_r$, and $H_S|\mathrm{slice}_a(q,r)\rangle = 0$. It is straightforward to see that this check can be done efficiently and deterministically. The check $q \in G_Q$ (and similarly $r \in G_r$) is just a classical verification of constraints, because every term in H_Q is diagonal in the pure Q basis, and $|\hat{q}\rangle$ is a pure Q basis state. Next, note that $|\mathrm{slice}_a(q,r)\rangle$ is a basis state in the a basis, whose classical description is easily recovered from q, r, and a, by definition. Each term in H_S is diagonal in the a basis, so the check $H_S|\mathrm{slice}_a(q,r)\rangle = 0$ is classical as well.

We now show correctness. Combining Lemmas 4.3.1 and 4.3.2, H has a zero-energy state if and only if $\mathrm{tr}(P_Q P_R P_S) \neq 0$. P_Q is the projector onto the zero-energy eigenspace of H_Q, which is spanned by the states $\{|\hat{q}\rangle : q \in G_Q\}$ (and similarly for P_R). So, we can rewrite the trace as

$$\begin{aligned}
\mathrm{tr}(P_Q P_R P_S) &= \mathrm{tr}(P_Q P_S P_R P_S) && (4.28)\\
&= \sum_{q \in G_Q, r \in G_R} \mathrm{tr}\left(|\hat{q}\rangle\langle\hat{q}|P_S|\breve{r}\rangle\langle\breve{r}|P_S\right) && (4.29)\\
&= \sum_{q \in G_Q, r \in G_R} \left|\langle\hat{q}|P_S|\breve{r}\rangle\right|^2. && (4.30)
\end{aligned}$$

Note that $|\hat{q}\rangle = Q|q_1\rangle \otimes \cdots \otimes Q|q_n\rangle = |\hat{q_1}\rangle \cdots |\hat{q_n}\rangle$, and similarly $|\breve{r}\rangle = |\breve{r_1}\rangle \cdots |\breve{r_n}\rangle$.

We show that for any $q, r \in \mathcal{D}^n$, $\langle\hat{q}|P_S|\breve{r}\rangle = \langle\mathrm{slice}_{\bar{a}}(q,r)|P_S|\mathrm{slice}_a(q,r)\rangle$, where $\bar{a}$ is the bitwise complement of a. First, note that $\langle\hat{q}|P_S|\breve{r}\rangle$ is just $\langle\mathrm{slice}_{1^n}(q,r)|P_S|\mathrm{slice}_{0^n}(q,r)\rangle$, by definition. Let us first demonstrate how to swap the first qudit. Since Q and R are real-valued, $|\hat{q_1}\rangle$ and $|\breve{r_1}\rangle$ are as well. We apply Lemma 4.4.4 then Lemma 4.4.3 to get

$$\begin{aligned}
\langle\hat{q}|P_S|\breve{r}\rangle &= \langle\breve{r_1}|\langle\hat{q_2}|\cdots\langle\hat{q_n}|P_S^{\Gamma_1}|\hat{q_1}\rangle|q_2\rangle\cdots|q_n\rangle && (4.31)\\
&= \langle\breve{r_1}|\langle\hat{q_2}|\cdots\langle\hat{q_n}|P_S|\hat{q_1}\rangle|q_2\rangle\cdots|q_n\rangle. && (4.32)\\
&= \langle\mathrm{slice}_{01^{n-1}}(q,r)|P_S|\mathrm{slice}_{10^{n-1}}(q,r)\rangle. && (4.33)
\end{aligned}$$

It is clear that this swapping procedure can be repeated on any subset of the n qudits. This shows that $\langle\hat{q}|P_S|\breve{r}\rangle = \langle\mathrm{slice}_{\bar{b}}(q,r)|P_S|\mathrm{slice}_b(q,r)\rangle$ for any $b \in \{0,1\}^n$. In particular, it holds for a.

Finally, H has a zero-energy state if and only if $\mathrm{tr}(P_Q P_R P_S) \neq 0$ if and only if there exists $q, r \in \mathcal{D}^n$ such that $q \in G_Q$, $r \in G_R$, and $\langle\mathrm{slice}_{\bar{a}}(q,r)|P_S|\mathrm{slice}_a(q,r)\rangle$, by (4.30) and our last observation. This is almost the check made by the NP verifier: it remains to show that $\langle\mathrm{slice}_{\bar{a}}(q,r)|P_S|\mathrm{slice}_a(q,r)\rangle \neq 0$ if and only if $H_S|\mathrm{slice}_a(q,r)\rangle = 0$. If $H_S|\mathrm{slice}_a(q,r)\rangle = 0$, then $h_i^S|\mathrm{slice}_a(q,r)\rangle$ for all i, and thus $P_S|\mathrm{slice}_a(q,r)\rangle = |\mathrm{slice}_a(q,r)\rangle$. Then

$$\begin{aligned}
\langle\mathrm{slice}_{\bar{a}}(q,r)|P_S|\mathrm{slice}_a(q,r)\rangle &= \langle\mathrm{slice}_{\bar{a}}(q,r)|\,\mathrm{slice}_a(q,r)\rangle && (4.34)\\
&= \langle\mathrm{slice}_{1^n}(q,r)|\,\mathrm{slice}_{0^n}(q,r)\rangle && (4.35)\\
&= \langle\hat{q}|\breve{r}\rangle && (4.36)\\
&\neq 0, && (4.37)
\end{aligned}$$

by the first condition of the theorem.

On the other hand, if $H_S|\mathrm{slice}_a(q,r)\rangle \neq 0$, then there is some h_k^S such that $h_k^S|\mathrm{slice}_a(q,r)\rangle \neq 0$. h_k^S is diagonal in the a basis, so by definition $|\mathrm{slice}_a(q,r)\rangle$ is an eigenstate of h_k^S.

Since h_k^S is a projector, it must be a 1-eigenstate. Then $(I - h_k^S)|\text{slice}_a(q,r)\rangle = 0$, so $\langle \text{slice}_{\bar{a}}(q,r)|P_S|\text{slice}_a(q,r)\rangle = 0$ by commuting $I - h_k^S$ to the right. This completes the proof.

$\square$

Chapter 5

Conclusion

We give some remarks and open questions which are raised by our results.

$\mathsf{QMA_1}$-completeness. The main question left open by our work is whether QUANTUM-k-SAT restricted to the *pure* X and Z bases is still QMA_1-hard. This would be quite an interesting result if true, since this class of Hamiltonians, at least in the basis-restriction setting, is the most restricted instance which is not trivially classical.

The main bottleneck for showing such a result using our proof method is that the propagation term in the Kitaev Hamiltonian are naturally projectors which are mixed-basis diagonal. It is likely that a gadget reduction would be needed to make the given a pure-basis diagonal Hamiltonian. This brings up another bottleneck which is commonly found when analyzing QMA_1-hardness: known gadget reductions for QMA do not carry over nicely, since they break perfect completeness. One other interesting direction is to further explore QMA_1 with the gate set $\mathcal{G}$, since $\mathsf{QMA}^{\mathcal{G}}$ is now known to be complete for two seemingly very different problems, $\mathcal{B}^{XZ}$-QUANTUM-6-SAT and weighted graph Gapped Clique Homology [Rud25].

Inclusion in NP. Some interesting open questions arise when considering the various conditions built into Theorem 4.2.4. Is the problem still in NP after removing one of these conditions? The main open question is whether the method of trace computation and slicing extend to the case in which the Hamiltonian can include more than one (strictly) mixed basis. If the problem is still in NP when all mixed Q and R bases are allowed, this would compare directly to our hardness result (Theorem 3.1.2) in the non-commuting case.

One bottleneck in showing this is in the term $\langle \hat{q}|P_S|\breve{r}\rangle$ which arises during the trace computation. It is straightforward to see that this term becomes $\langle \hat{q}|P_S P_T|\breve{r}\rangle$ when we allow for two strictly mixed bases S and T. However, a precondition of 4.4.3 is that the middle matrix is diagonal in a mixed basis, which is not true of $P_S P_T$. It is unclear if other techniques can allow us to overcome this.

Also worth addressing are the two conditions on the bases stated in 4.2.4. Condition 1 seems fundamental, and is unlikely to be removed when using a proof method which involves trace computation. On the other hand, the second condition is not as critical. We require that our states and gates be real-valued because 4.4.3 involves taking the partial transpose, an operation which does not play nicely with complex values. Still, it may be possible to generalize 4.4.3 and the slicing argument in a way which removes condition 2.

We also remark that Theorem 4.2.4 straightforwardly extends to the case of qubit Hamiltonians which have tensor structure but over *blocks* of qubits (under the same conditions in Theorem 4.2.4). Like qudits, this is a case which is left open by the result of Bravyi and Vyalyi [BV05]. For example, let one basis be a tensor product of the Bell basis, and the other basis a tensor product of the "X_1Z_2" basis. These bases satisfy condition 1 of Theorem **??**, so the associated commuting Hamiltonian problem is in NP.

Bibliography

[Aha03] Dorit Aharonov. A Simple Proof that Toffoli and Hadamard are Quantum Universal, January 2003.

[BV05] Sergey Bravyi and Mikhail Vyalyi. Commutative version of the local Hamiltonian problem and common eigenspace problem. *Quantum Info. Comput.*, 5(3):187–215, May 2005.

[Coo71] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, pages 151–158, New York, NY, USA, May 1971. Association for Computing Machinery.

[DN06] Christopher M. Dawson and Michael A. Nielsen. The Solovay-Kitaev algorithm. *Quantum Info. Comput.*, 6(1):81–95, January 2006.

[GN16] David Gosset and Daniel Nagaj. Quantum 3-SAT Is QMA1-Complete. *SIAM Journal on Computing*, 45(3):1080–1128, January 2016.

[HHH96] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. Separability of mixed states: Necessary and sufficient conditions. *Physics Letters A*, 223(1):1–8, November 1996.

[Kro67] M. R. Krom. The Decision Problem for a Class of First-Order Formulas in Which All Disjunctions Are Binary. *Zeitschrift fur mathematische Logik und Grundlagen der Mathematik*, 13(1-2):15–20, 1967.

[KSV02] A. Yu. Kitaev, A. Shen, and M. N. Vyalyi. *Classical and Quantum Computation.* American Mathematical Society, USA, June 2002.

[Lev73] L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.

[NC10] Michael A. Nielsen and Isaac L. Chuang. Quantum Computation and Quantum Information: 10th Anniversary Edition, December 2010.

[Per96] Asher Peres. Separability Criterion for Density Matrices. *Physical Review Letters*, 77(8):1413–1415, August 1996.

[Rud25] Dorian Rudolph. Towards a universal gateset for QMA1, April 2025.

[Shi03] Yaoyun Shi. Both Toffoli and controlled-NOT need little help to do universal quantum computing. *Quantum Info. Comput.*, 3(1):84–92, January 2003.

[Sip96] Michael Sipser. *Introduction to the Theory of Computation*. International Thomson Publishing, 1st edition, November 1996.