



Commuting Local Hamiltonians Beyond 2D

John Bostanci^{*1} and Yeongwoo Hwang^{†2}

¹Columbia University

²Harvard University

April 2025

Abstract

Local Hamiltonians are a quantum analogue to SAT in classical complexity. As a model of “intermediate” complexity, *commuting* local Hamiltonians provide a testing ground for studying many of the most interesting open questions in quantum information theory, including the quantum PCP conjecture and the nature of entanglement, while possibly being more amenable to analysis. Despite its simpler nature, the exact complexity of the commuting local Hamiltonian problem (CLH) remains elusive. A number of works have shown that increasingly expressive families of commuting local Hamiltonians admit completely classical verifiers. Despite intense work, proofs placing CLH in NP rely heavily on an underlying 2D lattice structure, or a very constrained local dimension and locality.

In this work, we present a new technique to analyze the complexity of various families of commuting local Hamiltonians: *guided reductions*. Intuitively, these are a generalization of typical reduction where the prover provides a guide so that the verifier can construct a simpler Hamiltonian. The core of our reduction is a new rounding technique based on a combination of Jordan’s Lemma and the Structure Lemma. Our rounding technique is much more flexible than previous work and allows us to remove constraints on local dimension in exchange for a rank-1 assumption. Specifically, we prove the following two results:

1. 2D-CLH for rank-1 instances are contained in NP, independent of the qudit dimension. It is notable that this family of commuting local Hamiltonians has no restriction on the local dimension *or* the locality of the Hamiltonian terms.
2. 3D-CLH for rank-1 instances are in NP. To our knowledge this is the first time a family of 3D commuting local Hamiltonians has been contained in NP.

Our results apply to Hamiltonians with large qudit degree and remain non-trivial despite the quantum Lovász Local Lemma.

*Email: johnb@cs.columbia.edu

†Email: yeongwoohwang@g.harvard.edu

Contents

1	Introduction	3
1.1	On the rank-1 constraint	3
1.2	Related results	4
1.3	Our contributions	4
1.4	Discussion and future work	6
2	Technical overview	7
2.1	Rounding schemes via the Structure Lemma	7
2.2	Our Techniques	10
3	Preliminaries	12
3.1	Quantum preliminaries and notation	12
3.2	Local Hamiltonian problems	12
3.3	C^* -algebras and the Structure Lemma	13
3.4	Geometrically constrained commuting local Hamiltonians	14
4	Rounding commuting local Hamiltonians	16
4.1	Previous works recasted as rounding schemes	17
4.2	An improved rounding scheme	19
5	Tools for rank-1 commuting operators	21
6	$2D^*$-CLH⁽¹⁾ is in NP	23
6.1	Extending to degree $k > 4$	26
7	Commuting Hamiltonians in three dimensions	27
7.1	3D cubic lattice	28
7.2	General 3D lattices	32
A	Verifiers for commuting local Hamiltonians	35
B	Lemmas for commuting local Hamiltonians	36

1 Introduction

The local Hamiltonian problem [KSV02] is one of the most important problems in quantum complexity theory, and the central problem of Hamiltonian complexity. Understanding the properties of local Hamiltonians, including the properties of their spectrum, ground states, and Gibbs states, has been a central question of study for the past 20 years, and has greatly improved our understanding of both quantum mechanics and quantum computation.

For those aiming to understand the properties of local Hamiltonians, commuting local Hamiltonians (CLHs) represent a kind of half-way point between classical constraint satisfaction problems and general local Hamiltonians. On one hand, the ground states of commuting Hamiltonians can still exhibit high entanglement and can be non-trivial, as exhibited by the Toric code and the recent proof of the no low-energy trivial states (NLTS) theorem [ABN23]. On the other hand, commutation often simplifies the analysis of quantum algorithms, as measurements in a shared basis do not experience the uncertainty principle when multiple measurements are applied. In addition to this, commuting local *projectors* always have an integral spectrum, and satisfy perfect completeness, making them robust to small perturbations and thus making them prime candidates for studying constructions of quantum PCPs. These properties make commuting local Hamiltonians a popular test bed for physicists and computer scientists alike; many advances in Gibbs state sampling come from analyzing the Gibbs states of commuting local Hamiltonians [KB16; Bar+23; Din+24; HJ24], and work towards the quantum PCP conjecture has largely focused on quantum LDPC codes, and related code Hamiltonians [ABN23; Cob+23].

Thus, by sharing some properties of classical constraint satisfaction problems and some properties of local Hamiltonians, the commuting local Hamiltonian problem represents an interesting lens through which to study the following problem: what properties of quantum computation make it more powerful than classical computation? The problem of determining the complexity of commuting local Hamiltonians has been studied explicitly by a number of authors, starting with the work of [BV04], who shows that the ground energy of 2-local commuting Hamiltonians are classically verifiable. Since then, a long line of works have shown that determining the ground energy of commuting Hamiltonians that are qubit 3-local or almost Euclidean [AE15], on locally expanding graphs [AE11], 2D over qubits [AKV18], and on a square lattice over qutrits [LJ23] are also problems contained in NP. In our work, we continue this line of inquiry and study the complexity of commuting local Hamiltonians. However, whereas previous work largely relied either on (a) Hamiltonian terms being extremely local, or (b) the local qudit dimension being extremely small, we study a class of commuting Hamiltonians with local terms of arbitrarily high locality, and unrestricted qudit dimension. The catch is that we impose strong restrictions on the rank of each term; namely, all Hamiltonian terms should be rank-1 projections.

A central theme of our work is the construction of *reductions* between CLH instances. In the classical world, reductions have provided an incredibly general framework for relating the computational complexity of two families of problems are through reductions. Starting with the seminal Cook-Levin Theorem [Coo71] which placed Boolean satisfiability in NP, reductions have paved the way for a rich set of tools for placing increasingly exotic decision tasks in NP. For QMA, the quantum analogue of NP, similar efforts have been made, the most successful of which are *perturbation gadgets*, which were original used to reduce the locality of QMA-hard local Hamiltonians from 3 to 2 [KKR05]. Since then, perturbation gadgets have been used extensively to show many other, often much more constrained, Hamiltonians are also QMA-complete [CM16]. Still, perturbation gadgets are not universally applicable and, in particular, cannot be used to construct *commuting* Hamiltonians. To address this deficit, we introduce a framework of *guided reductions* which preserve commutativity.

1.1 On the rank-1 constraint

All of the results in the paper apply to a restriction of the commuting local Hamiltonian problem where all the Hamiltonian terms are rank-1 projectors. It is already known [LJ23] that restricting to projectors does to change the complexity of the commuting local Hamiltonian problem, but the rank-1 version of this problem has not been studied in detail before. We believe this to be a worthwhile model to consider for a number of reasons. First, we note that in related complexity classes like NP, QMA, and QCMA, the rank-1 versions of canonical complete problems retain their hardness. Taking circuit SAT [Coo71; Lev73] as an example,

k -SAT can be viewed as a rank-1 local Hamiltonian version of SAT, where the Hamiltonian is a sum over clauses of the assignment that makes the clause false. For QMA, [CM16] showed that the Heisenberg local Hamiltonian problem is QMA complete. The Heisenberg local Hamiltonian problem is the local Hamiltonian problem where each local term is (a constant times) the projector onto the singlet state, $\frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$. One important reason for studying the complexity of commuting local Hamiltonians is that they provide a simplified setting in which one can study the complexity of the (non-commuting) local Hamiltonian problem. Since rank-1 local Hamiltonians are QMA complete, we believe that studying the rank-1 version of commuting local Hamiltonians can provide insights into the nature of the local Hamiltonian problem, and provides a useful test-bed for studying reductions between local Hamiltonian problems.

On the other hand, one should be careful that this problem remains non-trivial. Specifically, the quantum Lovász Local Lemma [AKS12] shows that if we have a Hamiltonian H on d -dimensional qudits with a uniform bound r on the rank, k on the locality, and g on the number of terms acting non-trivially on any qudit, then the Hamiltonian H is always satisfiable, as long as

$$g \leq \frac{d^k}{re},$$

where e is Euler’s number. This implies that if we restrict to a rank-1 qubit Hamiltonian on the 2D square lattice (so $g = 4$ and $d^k/re > 5.9$), the problem becomes trivial! Therefore, we necessarily need to consider more complex geometries, such as the 2D quasi-Euclidean complexes studied by [AKV18], where g can be any $c \in \mathcal{O}(1)$.

1.2 Related results

Our results follow a series of works studying the complexity of the commuting local Hamiltonian problem, under the restriction that all terms are commuting. The line of study was initiated by [BV04] who showed that the 2-local CLH problem was contained in NP. They also considered the specialized case of *factorized* commuting local Hamiltonians, where each Hamiltonian term h can be written as a tensor product of single-qudit operators. In this specialized setting, they showed that qubit CLH is contained in NP, regardless of geometry and locality. A key contribution of their work was the introduction of the *Structure Lemma*. This lemma characterizes the local algebra of commuting operators and has been an integral part of nearly all works on CLH since then. In fact, the lack of tools beyond the Structure Lemma has been a stumbling block on extending results to higher localities and qudit dimension, as the “structure” induced by the lemma becomes increasingly complex as the qudit dimension and locality increase.

The next major step along this line was due to [AE11] who extended the result of [BV04] to showed that 3-local qubit CLHs are in NP. For qutrits, the assumption that the underlying interaction graph is “nearly Euclidean” also yields containment in NP. Following this work, other authors moved onto stronger geometric restrictions. [Sch11] considered the 2D square lattice and showed qubit Hamiltonians defined in the lattice can be placed in NP. A follow up work by [AKV18] showed that Schuch’s work can even be made *constructive*, i.e. the verifier is able to efficiently prepare a ground state of the local Hamiltonian, rather than just be convinced about its existence. Their work can be viewed as reducing 2D lattice CLHs to a generalization of the Toric code, whose ground state can be prepared efficiently in classical polynomial time. Finally, [IJ23] extended the ideas contained in [Sch11] to show that CLHs with qutrits on the 2D square lattice are also in NP. Like the proof Schuch, this result is non-constructive.

A natural question is whether regardless of qudit dimension or geometric structure, CLHs are generically in NP. One evidence that this might not be the case comes from [GMV16], who show that the Ground State Connectivity (GSC) problem for commuting local Hamiltonians is QCMA-hard, which matches the complexity of GSC for general local Hamiltonians.

1.3 Our contributions

Our main contribution is to demonstrate a family of commuting local Hamiltonians on which a guided reduction can be performed to the 2-local CLH problem; rank-1 commuting Local Hamiltonians in 2D and 3D. We define a guided reduction as follows.




Figure 1: An example of a guided reduction. On the left is a Hamiltonian over registers $R_1 \otimes \cdots \otimes R_5$, with the green regions denoting individual Hamiltonian terms. A guided reduction is composed of a series of applications of a rounding schemes, which iteratively simplify the Hamiltonian. In the first round, the register \tilde{R}_0 is decomposed into registers \tilde{R}_0^a and \tilde{R}_0^b (while the non-tilde'd registers are unchanged). In the second round, \tilde{R}_4 is entirely removed.

Definition 1 (Guided reduction). *We say that the CLH family \mathcal{H} has an (ϵ, δ) guided reduction to another CLH family \mathcal{H}' if there is a mapping $\mathcal{M} : \mathcal{H} \times \{0, 1\}^* \rightarrow \mathcal{H}'$ such that for a n -qudit Hamiltonian $H \in \mathcal{H}$,*

- *If $\lambda_0(H) = 0$ then there exists a string s with $|s| \in \text{poly}(n)$ such that $\lambda_0(\mathcal{M}(H, s)) = 0$.*
- *If $\lambda_0(H) > \epsilon$, then for all strings s , $\lambda_0(\mathcal{M}(H, s)) > \delta$.*

Here $\lambda_0(H)$ is the minimum eigenvalue of H . When both ϵ and δ are constants in $(0, 1]$, we simply call this a guided reduction.

This definition is reminiscent of the “NP”-reductions referred to in [IJ23], where the reduction can not be performed in polynomial time, but the prover can attach an additional proof that allows the verifier to efficiently transform the instance to an equivalent one. The important property of this type of reduction is that such a valid guiding string only exists for “yes” instances of the problem. The result of [AKV18] can also be interpreted as providing a guided reduction from 2D qubit CLHs to 2-local CLHs (see Section 2.2).

Our main results are the following two theorems.

Theorem 1.1 (Guided reduction from rank-1 2D CLH to 2-local CLH (informal)). *There is a guided reduction from the family of rank-1 2D CLHs to the family of 2-local CLHs.*

Theorem 1.2 (Guided reduction from rank-1 3D* CLH to 2-local CLH (informal)). *There is a guided reduction from the family of rank-1 3D* CLHs to the family of 2-local CLHs.*

In the above theorem, the “*” refers to some constraints on the types of 3D Hamiltonians for which our reduction applies; we cover these details in Section 3.4.

New techniques for rounding commuting local Hamiltonians Our guided reductions will be constructed using the building block of “rounding schemes” (Definition 19) in which a Hamiltonian H over a set of registers $R = R_1 \otimes \cdots \otimes R_n$ is “rounded” to a Hamiltonian \tilde{H} over a possibly smaller space, as in Figure 1. In particular, \tilde{H} is defined over registers $\tilde{R} = \tilde{R}_{i_1} \otimes \cdots \otimes \tilde{R}_{i_\ell}$, with \tilde{R} is a subspace of R . As long as we can ensure \tilde{H} remains commuting and has a zero-energy ground state if and only if H does, this step will allow us to iteratively *simplify* the Hamiltonian. This idea of rounding projectors is not new (prior works such as [BV04; IJ23] can be viewed as implementing a rounding scheme). Our contribution, however, is to formalize this notion and introduce more general rounding schemes.

In [IJ23], the authors devised a rounding scheme in the case where all but a single Hamiltonian term commutes with a set of projectors. In our work, we construct a rounding scheme which works even when there are two non-commuting terms, increasing the generality of this approach.

Theorem 1.3 (Rounding pairs of projectors (informal version of Theorem 4.7)). *Let $H = \sum h$ be a commuting local Hamiltonian and π_1, π_2 be a pair of projectors such that for all but one term h that does not commute with π_1 , it is removed by π_2 (in the sense that $\pi_2 h \pi_2 = 0$), and vice versa. Then H can be rounded*

to a new commuting local Hamiltonian \tilde{H} that has a zero-energy ground state if and only if H has a 0 energy ground state.

Our main insight behind this theorem is that in the case when only two local projectors survive, we can recover a commuting local Hamiltonian by projecting onto the Jordan blocks of the pair that are non-zero (see Lemma 3.1). We can view this as an extension of the rounding scheme from [IJ23] to pairs of projectors, where we look at blocks of size 2 instead of size 1. The use of Jordan’s Lemma requires a much more careful and detailed analysis of the commutative properties of Jordan blocks.

Characterization of commutation for rank 1 Hamiltonians Our second technical result is a characterization of rank 1 commuting Hamiltonians. The Structure Lemma can be viewed as providing a general scaffolding for how two operators can commute. When both operators are restricted to be rank 1, we can strengthen the Structure Lemma to show that commutation occurs in one of two specific ways.

Theorem 1.4 (Characterization of commuting rank 1 projectors (informal version of Theorem 5.1)). *Let P and Q be rank 1 projectors so that P acts non-trivially only on registers A, B and Q acts non-trivially only on B, C . Further assume that P and Q commute. Then, restricted to the subspace B , either*

1. (Singular commutation) both P and Q are projections onto the same state ψ , or
2. (Reducing commutation) P and Q are projections onto orthogonal spaces.

This result allows us to employ a strategy of “puncturing and repairing” paths through local Hamiltonians. In particular, we can partition the local terms adjacent to a qudit into two sets such that terms in each set interact *only* on the qudit. Within each set, if we find that the local terms commute in a singular way, we can show that the qudit is a so-called classical qudit and can be removed. In the case when the local terms commute in a reducing way, the prover can provide a pair of projector that removes all but two of the terms. The prover’s projectors essentially cuts a small patch into the local Hamiltonian, but leaves terms that no longer commute. We patch these holes using the rounding technique described previously, and by repeating this protocol we can cut long strings into geometrically local commuting Hamiltonians.

Cubulation and triangulation techniques Our final contribution is a more detailed application of the triangulation technique of [AKV18]. In particular, we use a technique we call “cubulation”, where we super-impose a 3D cubic lattice on top of a general polyhedral lattice. We subsequently use our cutting and repairing technique to remove Hamiltonian terms intersecting both the vertices and edges of the super-imposed cubic lattice, allowing us to reduce a 3D commuting local Hamiltonian problem to a 2-local problem, which is known to be in NP.

We note that applying the cubulation technique to a general polyhedral lattice requires dealing with many different cases related to how the Hamiltonian terms intersect with the edges of the super-imposed lattice.

Verifiers for commuting local Hamiltonians In Appendix A, we provide a simple model of verifiers that captures the complexity of commuting local Hamiltonians. As we discuss in the next subsection, we believe that demonstrating hardness for commuting local Hamiltonian problems is one of the most interesting directions of research related to commuting Hamiltonians. We admit that the model we provide does not make any progress towards this goal, but we hope that providing a tangible “complexity class” associated with commuting local Hamiltonians, we will at least provide a language that can be used to describe this problem.

1.4 Discussion and future work

In this paper, we present a number of new results related to the geometric commuting local Hamiltonian problem. We show that in 2D, verifying the ground energy of rank-1 commuting local Hamiltonians of *any* local dimension are in NP. We do this by combining the triangulation technique from [AKV18] with a novel commuting Hamiltonian rounding scheme with a characterization of commutation for rank-1 commuting local Hamiltonians to puncture paths in 2D lattices. The rounding technique we provide works for much

more generic conditions than rank-1 commuting local Hamiltonians, and we believe it is of independent interest.

Secondly, we show that a wide family of 3D commuting local Hamiltonians are in NP. We show how to extend the triangulation technique to a cubulation technique. Using our puncturing technique, we can remove local terms that lie on either edges or vertices of the cubic lattice, again reducing the problem to the 2-local commuting local Hamiltonian problem. We note that neither of our reductions are “constructive,” in that it is not clear how to recover a succinct description of the commuting local Hamiltonian, even given the proof that it is in NP.

There are several questions that could be considered direct follow ups to our work. For example, what is the complexity of general rank-1 commuting local Hamiltonians. Can any tools be developed that allow for analysis of commuting Hamiltonians, even those that are rank-1 that do not have geometric locality? We note that our techniques ultimately rely heavily on triangulation or cubulation. Although we could potentially apply our rounding technique to puncture holes in more complicated commuting local Hamiltonians, without geometric locality it is very difficult to see how puncturing holes or paths through a Hamiltonian could be useful for reducing the problem to the 2-local commuting local Hamiltonian problem.

Another natural question to ask is whether our results, and those of [LJ23] can be turned into constructive proofs. Opposed to this question, the authors wonder if commuting local Hamiltonians could be an instance where state or unitary complexity differ from decision complexity, for example, if it was shown that the complexity of preparing ground states of commuting local Hamiltonians in 2D or 3D is outside of $\text{stateBQP}^{\text{NP}}$. This would constitute one of the first instances where these theories differ.

Finally, the authors note that almost all results pertaining to commuting local Hamiltonians (this one included) have been focused on showing that increasingly complicated versions of the commuting local Hamiltonian problem are classically verifiable. Despite having limited evidence, from the work of [GMV17], that commuting local Hamiltonians should be harder than NP, there has been no formal evidence that commuting local Hamiltonians are hard for any class larger than NP. The authors believe providing *any* evidence that commuting local Hamiltonians go beyond NP is a fascinating open question. A natural class to study is “next biggest class”, MA, although we think that such a reduction to commuting local Hamiltonian problems might entirely capture the difficulty of this problem. Specifically, we can think of a MA protocol as starting from some state $|x\rangle|0\rangle$, first measuring some qubits in the $\{|+\rangle|+\rangle, |-\rangle|-\rangle\}$ basis, and then performing some reversible quantum circuit. Thus, simulating MA with a commuting local Hamiltonian would probably require dealing with a single non-commuting measurement. It is known that adding a single non-commuting measurement into a commuting local Hamiltonian yields a class of local Hamiltonians that is complete for QMA.

Acknowledgements We thank Sandy Irani and Jiaqing Jiang for helpful discussions related to the Structure Lemma and its limitations. We also thank Chinmay Nirkhe for recommending to us the problem of commuting local Hamiltonians.

YH is supported by the National Science Foundation Graduate Research Fellowship under Grant No. 2140743. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. JB is supported by Henry Yuen’s AFORS (award FA9550-21-1-036) and NSF CAREER (award CCF2144219). This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing, supported by NSF QLCI Grant No. 2016245.

2 Technical overview

2.1 Rounding schemes via the Structure Lemma

As noted in the introduction, the central tool used in the study of the commuting local Hamiltonian problem has been the Structure Lemma of [BV04]. An easy consequence of this lemma is the following corollary,

Corollary 3.6 (Structure of two commuting operators). *Let $\mathcal{A}_h \subseteq \mathcal{L}(\mathbb{R})$ be the C^* -algebra induced by a Hermitian h and $\mathcal{A}_{h'}$ be the C^* -algebra on \mathbb{R} induced by h' that commutes with h . Let \mathbb{R}_i^J be the decomposition*

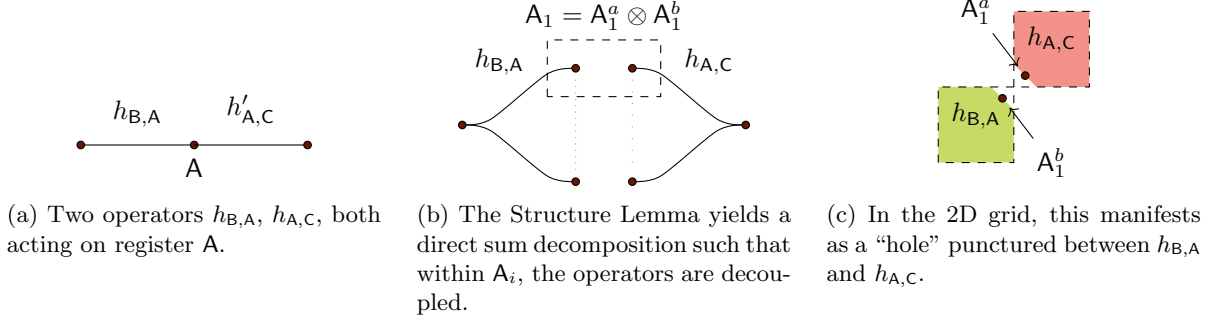


Figure 2: Puncturing holes via the Structure Lemma.

induced by lemma 3.5 applied to \mathcal{A}_h . Then the following holds:

$$\mathcal{A}_h = \bigoplus_i \mathcal{L}(\mathbb{R}_i^1) \otimes \text{id}(\mathbb{R}_i^2) \quad (9)$$

$$\mathcal{A}_{h'} \subseteq \bigoplus_i \text{id}(\mathbb{R}_i^1) \otimes \mathcal{L}(\mathbb{R}_i^2). \quad (10)$$

Crucially, all operators keep the decomposition $\mathbb{R} = \bigoplus_i \mathbb{R}_i$ invariant.

Informally, this states that by projecting onto a subspace spanned by \mathbb{R}_i , the operators h and h' can be thought of as decoupled, with h acting non-trivially only on \mathbb{R}_i^1 , and h' acting non-trivially only on \mathbb{R}_i^2 .

2-local setting. In the 2-local setting of [BV04], every pair of terms h, h' either do not interact, or interact *only* on a single-qudit register \mathbb{R} . This means that for each register, a single decomposition $\mathbb{R} = \bigoplus_i \mathbb{R}_i$ can be constructed so that every operators acts invariantly on the corresponding projectors $\pi_i \triangleq (\pi_i)_{\mathbb{R}_i}$. Concretely, each term h is *block diagonal* with respect to the subspaces \mathbb{R}_i . This decomposition naturally lends itself to a rounding scheme, where the projector is chosen to be one of the π_i 's. The resulting Hamiltonian is

$$\tilde{H}_i = \sum_{h \in H} \pi_i h \pi_i$$

Certainly $\lambda_0(\tilde{H}_i) = 0 \implies \lambda_0(H) = 0$. For the other direction, we use that for any CLH instance, $\lambda_0(H) = 0 \iff \text{Tr}[\prod_j (\mathbb{I} - h_j)] = \sum_i \text{Tr}[\prod_h \pi_i (\mathbb{I} - h) \pi_i] = 0$. Moreover, since all h commute with π_i , each trace in the sum is non-negative. Thus we get the following equivalent condition to the existence of a ground state:

$$\lambda_0(H) = 0 \iff \exists i \text{ s.t. } \text{Tr} \left[\prod_j (\pi_i - \pi_i h_j \pi_i) \right] > 0. \quad (1)$$

The expression on the right is equivalent to existence of a zero-energy eigenstate for the Hamiltonian $\tilde{H}_i = \sum_h \pi_i h \pi_i$, defined over registers $(\mathbb{R}_1, \dots, \pi_i \mathbb{R}_i \pi_i, \dots, \mathbb{R}_n)$, where $\pi_i \mathbb{R}_i \pi_i \leq \mathbb{R}_i$ is the sub-space generated by projecting vectors of \mathbb{R}_i onto π_i . Thus, mapping $H \rightarrow \tilde{H}_i$ yields our desired rounding scheme. Crucially, the resulting Hamiltonian is also *simpler* than the original; the Structure Lemma tells us that $\mathbb{R}_i = \otimes_h \mathbb{R}_i^{(h)}$ where each $\pi_i h \pi_i$ acts on non-trivially only on $\mathbb{R}_i^{(h)}$. This effectively decouples all terms on \mathbb{R} , as depicted in Figure 2.

k -local setting. For higher localities, the situation becomes more complicated. In general, if one has a 2D CLH instance with a k -wise interaction on a register \mathbb{R} , then commutation between adjacent terms is not entirely determined by the terms' structure on \mathbb{R} . For instance, imagine $h = |1\rangle\langle 1| \otimes \tilde{h}_R$ and $h' = |0\rangle\langle 0| \otimes \tilde{h}'_R$; these terms commute regardless of our choice of \tilde{h} and \tilde{h}' ! However, a variation of the above argument can be made to work. Consider when $k = 4$, as in Figure 3. Then the pairs h_2 and h_4 induce some decomposition

$R = \oplus_i R_i$, and pairs h_1 and h_3 induce a decomposition $R = \oplus_i \tilde{R}_i$. In general these decompositions are not the same. Let the projectors for these subspaces be $\{\pi_i\}_i$ and $\{\tilde{\pi}_i\}_i$. Then, rather than Equation (1), we get

$$\lambda_0(H) = 0 \iff \exists i, j \text{ s.t. } \text{Tr} \left[\tilde{\pi}_j (\mathbb{I} - h_1) \tilde{\pi}_j \tilde{\pi}_j (\mathbb{I} - h_3) \tilde{\pi}_j \pi_i (\mathbb{I} - h_2) \pi_i \pi_i (\mathbb{I} - h_4) \pi_i \prod_{j>4} (\mathbb{I} - h_j) \right] > 0 \quad (2)$$

As in the 2-local case, this equivalence only holds when the traces on the RHS are non-negative. Fortunately, the two matrices

$$A \triangleq \tilde{\pi}_j (\mathbb{I} - h_1) \tilde{\pi}_j \tilde{\pi}_j (\mathbb{I} - h_3) \tilde{\pi}_j \quad (3)$$

$$B \triangleq \pi_i (\mathbb{I} - h_2) \pi_i \pi_i (\mathbb{I} - h_4) \pi_i \prod_{j>4} (\mathbb{I} - h_j) \quad (4)$$

are both PSD, and $\text{Tr}[AB] \geq 0$. Note that this is sensitive to ordering and the same may not hold if we order the terms as h_1, h_2, h_3, h_4 . As this equivalence will be heavily used going forward, we formalize it in the following lemma.

Lemma 2.1 (Equivalence under projectors, implicit in [JJ23]). *Suppose we have a CLH instance $H = \sum_{i \in [m]} h_i$ and two POVM measurements $\Pi_1 = \{\pi_1^1, \dots, \pi_{k_1}^1\}$ and $\Pi_2 = \{\pi_1^2, \dots, \pi_{k_2}^2\}$ and disjoint sets of Hamiltonian terms $S, T \subseteq [m]$ such that,*

- All terms h_i with $i \in [m] \setminus T$ commute with Π_1 .
- All terms h_i with $i \in [m] \setminus S$ commute with Π_2 .
- All remaining terms h_i with $i \in [m] \setminus (T \cup S)$ commute with both Π_1 and Π_2 .

Then,

$$\lambda_0(H) = 0 \iff \exists i \in [k_1], j \in [k_2] \text{ such that } \text{Tr} \left[\prod_{i \in S} (\pi_i^1 - \pi_i^1 h_1 \pi_i^1) \prod_{i \in T} (\pi_j^2 - \pi_j^2 h_2 \pi_j^2) \bar{H}_{rest} \right], \quad (5)$$

where $\bar{H}_{rest} \triangleq \prod_{i \in [m] \setminus (S \cup T)} (\text{id} - h_i)$. Note that the ordering of the terms on the RHS matters.

The proof is deferred to Appendix B. Equation (2) is recovered by taking $S = \{1, 2\}$ and $T = \{3, 4\}$. This lemma therefore yields a statement equivalent to Equation (1). However, even with this equivalence it is not clear how to convert the right-hand-side into a CLH instance. Not only do the terms $\tilde{\pi}_j h_1 \tilde{\pi}_j$ and $\pi_i h_2 \pi_i$ not commute, we do not have a consistent subspace onto which can restrict R .

Previous works dealt with this in a couple different ways. In [Sch11], any non-trivial decomposition (i.e. $|\{R_i\}_i| > 1$) on qubits (so dimension 2 registers R) requires $\dim(R_i) = 1$. Therefore,

$$\tilde{\pi}_j = |\psi\rangle\langle\psi| \quad \text{and} \quad \pi_i = |\varphi\rangle\langle\varphi|.$$

For h_2 and h_4 we perform the same restriction as before:

$$\begin{aligned} h_2 &\rightarrow |\varphi\rangle\langle\varphi| h_2 |\varphi\rangle\langle\varphi| = |\varphi\rangle\langle\varphi| \otimes \tilde{h}_2 & h_1 &\rightarrow |\psi\rangle\langle\psi| h_1 |\psi\rangle\langle\psi| = |\psi\rangle\langle\psi| \otimes \tilde{h}_1 \\ h_4 &\rightarrow |\varphi\rangle\langle\varphi| h_4 |\varphi\rangle\langle\varphi| = |\varphi\rangle\langle\varphi| \otimes \tilde{h}_4 & h_3 &\rightarrow |\psi\rangle\langle\psi| h_3 |\psi\rangle\langle\psi| = |\psi\rangle\langle\psi| \otimes \tilde{h}_3. \end{aligned}$$




Figure 3: A degree 4 register R .




Figure 4: The process of reducing a 2D CLH instance to a 2-local CLH instance. Notice in the final figure that all terms besides those containing the center of each triangle are 2-local. In [AKV18] the center terms are simply removed and corrected later.

After this step, \tilde{h}_1 and \tilde{h}_3 may no longer commute. However, notice that every term acts as $|\varphi\rangle\langle\varphi|$ on \mathbb{R} and thus this register (which [Sch11] calls a *removable* qudit) can be traced out from the system. [Sch11] shows that “puncturing” out all removable qudits yields a 1-dimensional Hamiltonian. Thus, the author obtains a guided reduction (where the prover has provided $\tilde{\pi}_j$ and π_i) from 2-local qubit CLHs on the 2D lattice to 1-dimensional (non-commuting) local Hamiltonians. Such Hamiltonians can be easily seen to be contained in NP.

It is not clear how to extend this argument beyond qudit dimension $d = 2$, as the subspaces from the Structure Lemma are no longer 1 dimensional and thus the Hamiltonian terms may act non-trivially within each subspace. Still, [IJ23] show that in the case of qutrits, Schuch’s techniques can be extended. One of their key insights is identifying *semi-separable* qutrits, which means that there is a non-trivial decomposition on which *all but one* term acts invariantly. In this case, [IJ23] demonstrate a rounding technique which reduces the original Hamiltonian H into another CLH instance H' with all semi-separable qutrits removed (the authors call this a “self-reduction”). Following this step, the Structure Lemma is applied and a careful casework of possible resulting decompositions of H' allow the authors to recover a 1D structure.

2.2 Our Techniques

In this work, we combine an improved rounding scheme with the idea of “triangulations” from [AKV18].

Guided reduction via triangulation. The key insight from [AKV18] is that assuming enough terms are removed from a 2D CLH, the remaining system can be viewed as two local. For now, consider a 2D grid (where terms are on the faces and qubits are on edges), as all the key ideas can be seen from this restricted case. First, a triangulation is constructed over the 2D lattice (Figure 4b). This step *does not depend* on the underlying Hamiltonian, although the *centers* of the triangles should be selected as to reside within some Hamiltonian term. Then, paths are drawn from the center, through each edge of the triangle, and connecting an adjacent center. By construction, besides the center of the triangle, all terms are 2-local. [AKV18] deal with the 3-local center terms by showing that arbitrary 2D qubit Hamiltonians are equivalent to a (generalization of) the Toric code. Using properties of the Toric code, [AKV18] argue that the Hamiltonian contains enough “correctable” terms so that the triangle centers can be chosen to coincide with these correctable terms. These terms are removed, then corrected after preparing the ground state of the 2-local Hamiltonian.

Our rounding scheme. In the reduction to 2-local Hamiltonians of [AKV18], all that is really needed is a way create holes within each triangle. They accomplish this primarily through the removal of correctable terms. At a high level, we use our rounding scheme to puncture holes in the Hamiltonian in more varied ways, allowing us to apply the techniques from [AKV18] in more general settings than qubit 2D CLH instances.

We accomplish this through analyzing the possible local algebras of rank-1 commuting operators via the Structure Lemma. In the simplest case, imagine there are only two operators acting non-trivially on a register R . This is essentially the *2-local* setting of [BV04]. There a simple consequence of the Structure Lemma is the following lemma.

Lemma 2.2 (Informal version of Corollary 4.3). *There exists a subspace (sub-register) $R_i = R_i^a \otimes R_i^b \subseteq R$ such that when h and h' are restricted to R_i , h acts non-trivially only on R_i^a and h' acts non-trivially only on R_i^b .*

As long as this subspace has non-zero overlap with the ground space, restricting our Hamiltonian to R_i punctures a “hole” in the geometric structure of the Hamiltonian, as in Figure 2. But like in the discussion in Section 2.1, this argument does not directly apply when interactions are more than 2-local. To generalize this argument, we show that the local algebra of rank-1 commuting operators can be classified as *reducing* or *singular* where intuitively,

- If two operators commute in a *reducing* way, then in R_i at least one operator becomes the 0-operator.
- If two operators commute in a *singular* way, then $\dim R_i = 1$ and both act as a scalar.

For each combination of reducing and singular cases, we show that our rounding scheme yields a Hamiltonian where a hole has been punctured near R and thus a corner of the co-triangulation can be placed in this hole. A crucial part of this argument is that rounding schemes yield *another* CLH instance. This fact will be crucial in the 3D setting, where after an initial application of our rounding scheme, there are still “obstructions” in the 3D surface. However, these obstructions turn out to be 2-local and, since our intermediate Hamiltonian is commuting, we may apply the Structure Lemma via Lemma 2.2 to remove the obstructions.

In comparison to previous techniques,

1. The semi-separable self-reduction of [LJ23] also retains commutativity. However, our reduction is more general and (assuming the rank-1 constraint) can be applied generically to single qudit register in the Hamiltonian.
2. [Sch11] reduction is also quite general, however, the final Hamiltonian generated is no longer commuting. This makes it hard to imagine how one would generalize their techniques to more complex structures, where one might want to perform the reduction *recursively*.
3. The result of [AKV18] is in some sense closest to ours; for any triangulation and center, they can find a term (or adjacent term) which is removable. However, their ideas rely on the exact characterization of 2D qubit Hamiltonians as related to the defected Toric code.¹

Paper Overview The remainder of the paper is laid out as follows.

- In Section 3 we introduce some notation and review the basics of the local Hamiltonian problem and the Structure Lemma, both of which will play a major part in this paper. Furthermore, we introduce some geometric variations of the commuting local Hamiltonian problem.
- In Section 4, we give the key technical tools we will need in this paper. In Section 4.2, we describe our *rounding scheme*, which allows us to reduce from a CLH instance to a “simpler” CLH instance. To apply this rounding scheme to rank-1 CLH instances, we will need our characterization of local algebras of rank-1 CLHs as singular and reducing. This is done in Section 5.
- The bulk of our results are contained in Sections 6 to 7.
 - In Section 6, we prove that 2D*-CLH⁽¹⁾ is in NP (Theorem 6.1)
 - In Section 7, we prove that 3D*-CLH⁽¹⁾ is in NP (Theorem 7.1).

¹The defected Toric code is a generalization of the Toric code where real coefficients are permitted in front of each term.

3 Preliminaries

3.1 Quantum preliminaries and notation

A register R is a named finite-dimensional complex Hilbert space. If A , B , and C are registers, then ABC denotes the tensor product of the associated Hilbert spaces. We write $\mathcal{L}(R)$ to denote the set of all linear transformations over R . When R appears alongside other registers, for a linear transformation $L \in \mathcal{L}(R)$ we sometimes write L_R to make explicit that L acts non-trivially on R and as identity on other registers. Additionally, we write ρ_R to indicate that ρ is a state on R . Generally, one may think of a register R as corresponding to a qudit.

We denote the set of positive semidefinite operators on a register R by $\text{pos}(R)$. The set of density matrices on R is denoted $S(R)$. For a pure state $|\varphi\rangle$, we write φ to denote the density matrix $|\varphi\rangle\langle\varphi|$. We denote the identity operator as id . For an operator $X \in \mathcal{L}(R)$, we define $\|X\|_\infty$ to be its operator norm, and $\|X\|_1 = \text{Tr}(|X|)$ to denote its trace norm, where $|X| = \sqrt{X^\dagger X}$. For Hermitian matrices A, B , we use the notation $A \succeq B$ to indicate that $A - B$ is positive semi-definite (PSD). Similarly, $A \preceq B$ indicates that $B - A$ is PSD. We also define the commutator, $[A, B] = AB - BA$.

We will also use Jordan's Lemma, which characterizes the decomposition of a subspace induced by a pair of projectors.

Lemma 3.1 (Jordan's Lemma [Jor75]). *For any two Hermitian projectors Π_P and Π_Q on a register R , there exists a orthogonal decomposition of $R = \bigoplus_b R_b$ (the Jordan decomposition with respect to Π_P, Π_Q) into one- and two-dimensional subspaces $\mathcal{B} = \{S_b\}_b$ (Jordan subspaces, or blocks) where each R_b is invariant under both Π_P and Π_Q . Moreover*

- In each one-dimensional subspace, Π_P and Π_Q act as identity or rank-0 projectors.
- In each two-dimensional subspace, Π_P and Π_Q are rank-1 projectors. In particular, there exist distinct orthogonal bases $\{|p_b^0\rangle, |p_b^1\rangle\}$ and $\{|q_b^0\rangle, |q_b^1\rangle\}$ for S_b such that Π_P projects onto $|p_b^1\rangle$ and Π_Q projects onto $|q_b^1\rangle$.

3.2 Local Hamiltonian problems

Definition 2 (Local Hamiltonian). *A k -local Hamiltonian over registers $\mathcal{R} = \{R_1, \dots, R_n\}$ is a Hermitian operator $H \in \mathcal{L}(R_1 \dots R_n)$ such that H can be written as $H = \sum_{i=1}^\ell h_i$, where each h_i is PSD and $\|h_i\|_\infty \leq 1$. Each h_i acts non-trivially on registers $\mathcal{S} \subseteq \mathcal{R}$ with $|\mathcal{S}| \leq k$. We say that a k -local Hamiltonian acts on d -dimensional qudits if each register R_i corresponds to the Hilbert space \mathbb{C}^d .*

Definition 3 (Local Hamiltonian problem). *Given a family \mathcal{H} of k -local Hamiltonians and parameters c, s with $s - c \geq \frac{1}{\text{poly}(n)}$, the local Hamiltonian problem is a promise problem so that, promised that $H \in \mathcal{H}$ has minimum eigenvalue $\lambda_0(H) \geq s$ or $\lambda_0(H) \leq c$, to decide which is the case.*

Definition 4 (Commuting local Hamiltonian problem). *The commuting local Hamiltonian (CLH) problem is identical to the local Hamiltonian problem, except that the Hamiltonian family \mathcal{H} is composed of commuting Hamiltonians H . A Hamiltonian $H = \sum_i h_i$ is commuting if for all i, j , $[h_i, h_j] = 0$.*

Remark 1 (Projective CLH). In the context of showing that the local Hamiltonian problem for a family of commuting local Hamiltonians is in NP, it's sufficient to restrict to the case where each h_i is a projector. This is because the prover may first identify an eigenspace Π_i and eigenvalue λ_i of each term such that $\sum_i \lambda_i \leq c$, the completeness parameter of the CLH problem. The prover then helps verify that $\sum_i (\mathbb{I} - \Pi_i)$ has a zero-energy eigenspace. Going forward, we'll assume that each h_i is a projector.

This remark has the important consequence that we may always take $c = 0$ and $s = 1$. Next, a particularly important notion in this work is a *rank-constrained* CLH.

Definition 5 (Rank constrained commuting local Hamiltonian). *A rank- r constrained, k -local CLH, denoted $CLH^{(r)}$, is a k -local CLH where each h_i is a matrix of rank at most r .*

We also give some simple definitions related to local Hamiltonians.

Definition 6 (Degree of a register). Suppose $H = \sum_i h_i$ is a CLH over n qudits, corresponding to registers R_1, \dots, R_n , the degree of a register R_i , denoted $\deg_H(R_i)$ is the number of terms in H which act non-trivially on R_i . If H is clear from context, we write this as simply $\deg(R_i)$.

Then, the *degree of a Hamiltonian* is defined simply as $\deg(H) = \max_i \deg_H(R_i)$. Note that this notion only makes sense once we identify a partitioning of the full space into registers $R_1 \dots R_n$. If unspecified, we associate registers with each qudit of the Hamiltonian.

When we have bounds on the degree, locality, and qudit dimension of a Hamiltonian, one can show that a wide range of Hamiltonians are always satisfiable.

Lemma 3.2 (Quantum Lovász Local Lemma (qLLL) [AKS12]). Let $H = \sum_i h_i$ be a Hamiltonian with locality k , degree g , qudit dimension d , and $\max_i \text{rank}(h_i) \leq r$. Then if $g \leq \frac{d^k}{re}$ then $\lambda_0(H) = 0$.

In our work, since we generally take $r = 1$, this implies only instances with large maximum degree are interesting.

3.3 C^* -algebras and the Structure Lemma

In this section we review the notion of a C^* -algebra, the Structure Lemma from [BV04] and the connection to commuting local Hamiltonians.

Definition 7 (C^* -algebra). Let R be a register, a C^* -algebra is any complex algebra $\mathcal{A} \subseteq \mathcal{L}(R)$ that is closed under the \dagger operation and includes the identity.

Definition 8 (Commuting algebras). Let \mathcal{A} and \mathcal{A}' be two C^* -algebras on R . We say \mathcal{A} and \mathcal{A}' commute if $[h, h'] = 0$ for all $h \in \mathcal{A}$ and $h' \in \mathcal{A}'$.

The connection between local Hamiltonians and algebras is made through the concept of an “induced algebra”.

Definition 9 (Induced algebra). Let h be a Hermitian operator acting on AB , and consider a decomposition of h into

$$h = \sum_{i,j} (h_{ij})_A \otimes (|i\rangle\langle j|)_B. \quad (6)$$

Here $\{|i\rangle_B\}_i$ is an orthogonal basis of B . Then the induced algebra of h on A , denoted \mathcal{A}_h^A , is defined to be the algebra generated by $\{(h_{ij})_A, \text{id}_A\}$.

The following lemma tells us that the induced algebra is independent of the choice of basis for B .

Lemma 3.3 (Claim B.3 of [AKV18]). Let h be a Hermitian operator and consider two decompositions of h

$$h = \sum_{i,j} (h_{ij})_A \otimes (g_{ij})_B = \sum_{i,j} (\widehat{h}_{ij})_A \otimes (\widehat{g}_{ij})_B, \quad (7)$$

where both the sets $\{g_{ij}\}$ and $\{\widehat{g}_{ij}\}$ are linearly independent. Then the C^* -algebra generated by $\{h_{ij}\}_{ij}$ and $\{\widehat{h}_{ij}\}_{ij}$ are the same.

The induced algebra gives us a tool which we can use to determine if two Hermitian operators commute via the following lemma.

Lemma 3.4. Let $(h_1)_{AB}$ and $(h_2)_{BC}$ be two Hermitian operators. Then $[h_1, h_2] = 0$ if and only if $\mathcal{A}_{h_1}^B$ commutes with $\mathcal{A}_{h_2}^B$.

Finally, we recall the Structure Lemma from [BV04], which is the main tool prior work used to characterize the commutation between adjacent commuting terms. At a high level, the Structure Lemma says that algebras can be block-diagonalized, and that within each of these blocks, the algebra takes on a tensor product structure. For a proof of the lemma, see [Gha+15, Section 7.3].

Lemma 3.5 (The Structure Lemma). *Let $\mathcal{A} \subseteq \mathcal{L}(\mathbb{R})$ be a C^* -algebra on \mathbb{R} . Then there exists a direct sum decomposition $\mathbb{R} = \bigoplus_i \mathbb{R}_i$ and a tensor product structure $\mathbb{R}_i = \mathbb{R}_i^1 \otimes \mathbb{R}_i^2$ such that*

$$\mathcal{A} = \bigoplus_i \mathcal{L}(\mathbb{R}_i^1) \otimes \text{id}_{\mathbb{R}_i^2}. \quad (8)$$

A corollary of this lemma, and the reason why it is so useful for characterizing the properties of commuting local Hamiltonians, is that in order to commute with an algebra, another algebra must live entirely within the \mathbb{R}_i^2 subspaces. Formally, we have the following

Corollary 3.6 (Structure of two commuting operators). *Let $\mathcal{A}_h \subseteq \mathcal{L}(\mathbb{R})$ be the C^* -algebra induced by a Hermitian h and $\mathcal{A}_{h'}$ be the C^* -algebra on \mathbb{R} induced by h' that commutes with h . Let \mathbb{R}_i^j be the decomposition induced by lemma 3.5 applied to \mathcal{A}_h . Then the following holds:*

$$\mathcal{A}_h = \bigoplus_i \mathcal{L}(\mathbb{R}_i^1) \otimes \text{id}(\mathbb{R}_i^2) \quad (9)$$

$$\mathcal{A}_{h'} \subseteq \bigoplus_i \text{id}(\mathbb{R}_i^1) \otimes \mathcal{L}(\mathbb{R}_i^2). \quad (10)$$

Crucially, all operators keep the decomposition $\mathbb{R} = \bigoplus_i \mathbb{R}_i$ invariant.

An easy application of the Structure Lemma can be seen in Figure 2; there, 2-local interactions can be simplified to generate 1-local terms. Finally, we introduce the notion of a ‘‘classical’’ register \mathbb{C} . Intuitively, this means that under some local unitary $U_{\mathbb{C}}$, the register ‘‘looks classical,’’ in that every eigenstate is can be written as $|x\rangle_{\mathbb{C}} \otimes |\psi\rangle$, for a classical string x .

Definition 10 (Classical register). *For a Hamiltonian H , a register \mathbb{C} is classical if there exists a decomposition of \mathbb{C} into 1-dimensional spaces $\mathbb{R}_1 \oplus \dots \oplus \mathbb{R}_d$ which is invariant under all terms h of H .*

Classical registers will be very useful for us; given any classical register \mathbb{C} , we can induce a ‘‘hole’’ in the Hamiltonian H on that register. This is because if each Hamiltonian term h acts invariants w.r.t. the decomposition $(\mathcal{H}_{\mathbb{R}})_1 \oplus \dots \oplus (\mathcal{H}_{\mathbb{R}})_d$, this implies that each $h = \sum_i |i\rangle\langle i|_{\mathbb{R}} \otimes h_i$, where $|i\rangle\langle i|$ is the projector onto the 1-dimensional subspace $(\mathcal{H}_{\mathbb{R}})_i$, and the entire Hamiltonian can be written as

$$H = \sum_i |i\rangle\langle i|_{\mathbb{R}} \otimes H_i \quad \text{with} \quad H_i = \sum_h h_i,$$

and thus every eigenstate, and in particular the ground state, is of the form $|i\rangle_{\mathbb{R}} \otimes |\psi\rangle$. Therefore, the prover can provide the subspace containing the ground space and the Hamiltonian $\tilde{H} \triangleq H_i$ consists of terms which all act trivially on \mathbb{R} .

Observation 3.7 (Classical restriction). *Suppose \mathbb{C} is a classical qubit for the Hamiltonian H . Then there exists a one-dimensional projector $\pi_{\mathbb{R}}$ such that,*

$$\lambda_0(H) = 0 \iff \lambda_0(\pi H \pi) = 0. \quad (11)$$

Moreover, $\pi H \pi = \pi \otimes \tilde{H}$, where \tilde{H} is a CLH instance where each term acts trivially on \mathbb{C} . Thus the local Hamiltonian of H reduces to solving the local Hamiltonian problem on \tilde{H} , where the verifier implicitly keeps track of the 1-dimensional assignment on \mathbb{C} .

3.4 Geometrically constrained commuting local Hamiltonians

Now we formally define the notion of 2D and 3D commuting local Hamiltonians in terms of chain complexes. We refer to a chain complex in 2D as a polygonal complex.

Definition 11 (Polygonal complex). *A polygonal complex \mathcal{K} is a collection of points, lines, and polygons such that*

1. Every side of a polygon in \mathcal{K} is a line in \mathcal{K} , and every endpoint of a line in \mathcal{K} is a point in \mathcal{K} .
2. For every pair of polygons, their intersection is either the empty set, or a single line in \mathcal{K} , along with its endpoints. For every pair of lines in \mathcal{K} , their intersection is either the empty set, or it is a single point in \mathcal{K} .

In 3D, we refer to a similar chain complex as a polyhedra complex, although it is clear how this would generalize to arbitrarily high dimensions via the concept of a chain complex (which, in reality, we are just defining in geometric terms).

Definition 12 (Polyhedral complex). *A polyhedral complex \mathcal{K} is a collection of points, lines, polygons, and polyhedra such that*

1. Every face of a polyhedral in \mathcal{K} is a polygon in \mathcal{K} , and every side of a polygon in \mathcal{K} is a line in \mathcal{K} , and every endpoint of a line in \mathcal{K} is a point in \mathcal{K} .
2. For every pair of polyhedral, their intersection is either the empty set or a single polygon in \mathcal{K} , along with its sides and their endpoints. For every pair of polygons in \mathcal{K} , their intersection is either the empty set, or a single line in \mathcal{K} , along with its endpoints. For every pair of lines in \mathcal{K} , their intersection is either the empty set, or a single point in \mathcal{K} .

There are many ways to define a commuting local Hamiltonian instance on a chain complex, depending on which terms correspond to Hamiltonian terms and which correspond to qudits. We present a few of the notable versions of the commuting local Hamiltonian problem that have been studied in the past.

Definition 13 (2D-CLH). *A (k, d) -2D-CLH instance is an instance of the commuting k -local Hamiltonian problem where the d -dimensional qudits can be mapped to vertices and Hamiltonian terms can be mapped to faces in a polygonal complex in such a way that the Hamiltonian term on a face only acts non-trivially on the qudits associated with endpoints of the sides of the face.*

This is the most general way to define a two dimensional commuting local Hamiltonian problem, but often times more constraints are required to produce non-trivial results. [AKV18] define another way to associate a polygonal complex with a Hamiltonian, which they call 2D-CLH*. They also show that every 2D-CLH* instance can be transformed into a 2D-CLH instance without changing the local dimension, making this problem strictly easier than the general 2D-CLH problem.²

Definition 14 (2D*-CLH). *A (k, d) -2D*-CLH instance is an instance of the commuting k -local Hamiltonian problem where the d -dimensional qudits can be mapped to edges and Hamiltonian terms can be mapped to faces and vertices of a polygonal complex in such a way that the Hamiltonian term associated with a vertex acts on qudits associated with the edges adjacent to the vertex, and similarly for Hamiltonian terms on faces.*

The works of [Sch11; Jia23] consider an even more constrained version of the problem in two dimensions, where the polygonal complex is a square lattice.

Definition 15 (2D-CLH-Grid). *The d -2D-CLH-Grid problem is the $(4, d)$ -2D*-CLH problem where the polygonal complex is always a square lattice.*

Similar to the definition of 2D-CLH, we can define the 3D-CLH problem as follows.

Definition 16 (3D-CLH). *A (k, d) -3D-CLH instance is an instance of the commuting k -local Hamiltonian problem where d -dimensional qudits can be mapped to vertices and Hamiltonian terms can be mapped to volumes in a polyhedral complex such that the Hamiltonian term associated with a volume only acts non-trivially on qudits associated with vertices that lie on the volume.*

As in the 2D case, we also define a “starred” version, where qudits are placed on edges, rather than vertices.

²The authors of [AKV18] also show that 2D-CLH* instances can be turned into 2D-CLH instances, but at a cost of expanding the local dimension. Since their result only holds for *qubits*, it cannot be extended to general 2D-CLH instances.

Definition 17 (3D*-CLH). *A (k, d) -3D*-CLH instance is an instance of the commuting k -local Hamiltonian problem where d -dimensional qudits can be mapped to edges and Hamiltonian terms can be mapped to volumes in a polyhedral complex such that the Hamiltonian term associated with a volume only acts non-trivially on qudits associated with edges that lie on the volume.*

Finally, we can define a very constrained version of 3D*-CLH where the polyhedral complex is constrained to be a cubic lattice. This case turns out to be trivial due to the quantum Lovász Local Lemma, but is useful for pedagogical purposes.

Definition 18 (3D*-CLH-Grid). *A d -3D*-CLH-Grid instance is an instance of the commuting 12-local Hamiltonian problem where the d -dimensional qudits can be mapped to edges in a 3D cubic lattice and Hamiltonian terms can be mapped to cubes in the same lattice such that the Hamiltonian term associated with a volume only acts on qudits associated with edges of faces of the volume.*

Remark 2. For the CLH families 2D*-CLH, 2D-CLH-Grid, 3D*-CLH, and 3D*-CLH-Grid, if the parameters (k, d) are left unspecified, we take this to mean there is some $k, d \in \text{poly}(n)$ for which every Hamiltonian in the family is a (k, d) -CLH instance.

Finally, we remark that for our results, we make some mild assumptions on the 3D geometric structure.

Assumption 1 (Uniformly high degree). *Given a Hamiltonian H over registers R_1, \dots, R_n , we assume that $\min_{i \in [n]} \deg_H(R_i) \geq 4$.*

This assumption is necessary in order to apply our rounding scheme. This assumption is somewhat justified by Lemma 3.2; if the degree of *all* registers was strictly less than 4, then the resulting (rank-1) Hamiltonian would be trivial. However, even a single high degree term makes qLLL inapplicable and we need this assumption to hold for all registers.

Lastly the following assumption enforces a nice structure on the 3D surface.

Assumption 2 (Generalized cubic lattice). *We maintain some connection with the simpler cubic lattice case. This connection is through two restrictions on the lattice.*

1. *If there are two Hamiltonian terms h_1, h_2 sharing a face, then there is no term which shares a face with both h_1 and h_2 . Intuitively, this prevents “Jenga-like” structures, which will make it hard to draw an analogue to the tunnel picture (Figure 9) from the simpler case.*
2. *Each Hamiltonian term should have at least 5 faces. This condition will ensure that when we punctures holes for the vertex, there are enough faces so that each boundary path (e.g. the lines in Figure 11a) can exit through a distinct face.*

4 Rounding commuting local Hamiltonians

In this section, we describe the rounding schemes used in our paper. As stated previously, these rounding schemes serve as a single step in a guided reduction; the existence of a rounding scheme asserts that there exists a projector π (or a set of projectors $\{\pi_i\}_i$) which can be provided by the prover to help the verifier simplify the Hamiltonian. At a high level, our rounding schemes emerge when Lemma 2.1 not only provides an equivalent condition to the existence of a ground state, but the resulting trace expression can be “rounded” back to an equivalent (and hopefully simplified) CLH instance.

Formally, a rounding scheme for Hamiltonians is defined as follows:

Definition 19 (Rounding scheme for commuting local Hamiltonians). *Let $H = (h_1)_{A_1 B} + \dots + (h_k)_{A_k B}$ be a CLH instance defined over registers A_1, \dots, A_k, B . The registers A_i, A_j are allowed to overlap, in the sense that $A_i \cap A_j \neq \emptyset$. Let $\Pi = \{(\pi_i)_B\}_i$ be a set of projectors acting only on B . A rounding scheme is an efficient classical algorithm that takes as input a description of h_1, \dots, h_k , and $\{\pi_i\}_i$, and outputs a CLH instance \tilde{H}_Π such that H has a 0-energy ground state if and only if \tilde{H}_Π does.*

h_{12}	h_{13}	h_{14}	h_{15}
h_8	h_9	h_{10}	h_{11}
h_4	h_5	h_6	h_7
h_0	h_1	h_2	h_3

(a) Terms surrounding register R.

h_{12}	h_{13}	h_{14}	h_{15}
h_8	h_9	h_{10}	h_{11}
h_4	h_5	h_6	h_7
h_0	h_1	h_2	h_3

(b) When register R is classical, there is a choice of projector such that the terms in the resulting Hamiltonian act *trivially* on R.

Figure 5: Result of Corollary 4.1

Although in the above definition, we write \tilde{H}_Π to emphasize that the resulting projector is constructed using Π , we'll usually drop the Π subscript for ease of notation, and refer to the resulting Hamiltonian as \tilde{H} . This definition may seem a bit opaque in that it does not specify *how* \tilde{H} should be constructed. Naively, one might assume we sequentially apply each $\pi \in \Pi$ to H ; in 2-local setting described in Section 2.1, this is exactly what we did. In general, however, these projectors may not commute with all terms and more complicated rounding schemes are often necessary.

In addition to 2-local rounding, we have seen another example of a rounding scheme. Observation 3.7 yields the following,

Corollary 4.1 (Rephrasing of Observation 3.7). *Let $H = \sum_h h$ be a CLH instance over (R_1, \dots, R_n) with classical register $C = R_i$. Let $\{\pi_i = |\psi\rangle\langle\psi|_i\}_i$ be the projectors on to the corresponding 1-dimensional subspaces. Then there is a rounding scheme for H and $\{\pi_i\}_i$. Moreover, the rounding scheme yields a simplified CLH instance $\tilde{H} = \sum_h \tilde{h}$ on $(R_1, \dots, R_{i-1}, R_{i+1}, \dots, R_n)$ such that,*

- If h acts trivially on C , then $\tilde{h} = h$.
- Otherwise, $\tilde{h} = (\text{id} \otimes \langle\psi_i|_C) h (\text{id} \otimes |\psi_i\rangle_C)$.

The result of this rounding is depicted in Figure 5.

In this section, we will first show how the key technical tools in previous papers (e.g. [BV04; IJ23]) can be phrased in terms of a rounding scheme. Then, we will give analyze the local algebra of commuting, rank-1 operators, then show how this yields new rounding scheme for rank-1 CLH instances. By showing that these rounding schemes significantly simplify the Hamiltonian, applying them iteratively (with corresponding projectors supplied by the prover) will yield a *guided reduction* from 2D and 3D rank-1 CLH instances to 2-local CLH instances, known to be contain in NP.

4.1 Previous works recasted as rounding schemes

Rounding scheme of Bravyi-Vyalyi A central result of [BV04] is that *2-local* CLH instances are contained in NP. At the heart of this result is the following rounding scheme. Consider a Hamiltonian H defined over registers $\{R_1, \dots, R_n\}$. Let $B = R_k$ to be any qudit register, and let S_B be the set of Hamiltonian terms acting non-trivially on B . Fix some $h \in S_B$ and consider the pair of commuting operators h and $\bar{h} = \sum_{h' \neq h \in S_B} h'$. Applying the Structure Lemma yields a decomposition into a direct sum of tensor product spaces in which \bar{h} and h act on different particles. Recursively applying this argument to \bar{h} yields a decomposition

$$B = \bigoplus_i B_i = \bigoplus_i \bigotimes_{h \in S_B} B_i^h, \quad (12)$$

where the algebra of any term satisfies

$$\mathcal{A}_h^{\mathbb{B}} \subseteq \bigoplus_i \text{id} \otimes \cdots \otimes \mathcal{L}(\mathbb{B}_i^h) \otimes \cdots \otimes \text{id}. \quad (13)$$

In particular, h acts non-trivially only on the subspaces $\{\mathbb{B}_i^h\}_i$ and is block-diagonal with respect to the subspaces $\{\mathbb{B}_i\}_i$. The rounding scheme is then to set each π_i to be the projector onto \mathbb{B}_i , yielding the Hamiltonians $H_i = \pi_i H \pi_i = \sum_j \pi_i h_j \pi_i$. We see that this is a valid rounding scheme as

- Each pairs of terms $h_j, h_k \in H_i$ is commuting since both are block-diagonal across the subspace corresponding to π_i (either because both projectors act trivially on \mathbb{B} , or if they do not act trivially then because π_i corresponds to the direct sum induced by the Structure Lemma).
- In the statement of Lemma 2.1 set $S = \{h\}_{h \in H}$, $\mathcal{M} = \{\pi_i\}_i$, $T = \emptyset$, and $\widetilde{M} = \{\text{id}\}$. Then,

$$\lambda(H) = 0 \iff \exists i \text{ s.t. } \text{Tr} \left[\prod_j (\pi_i - \pi_i h_j \pi_i) \right] > 0. \quad (14)$$

But the RHS is equivalent to the existence of a ground state of the Hamiltonian $\widetilde{H} = H_i = \sum_j \pi_i h_j \pi_i$, over the reduced space $\{\mathbb{R}_1, \dots, \pi_i \mathbb{R}_k \pi_i, \dots, \mathbb{R}_n\}$. Thus, there exists a H_i such that $\lambda_0(H_i) = 0 \iff \lambda(H) = 0$.

This implies the following rounding scheme:

Lemma 4.2 (Rounding scheme for 2-local Hamiltonians). *Suppose there exists a register \mathbb{B} such that all terms acting non-trivially on \mathbb{B} interact on no other registers (i.e. their corresponding \mathbb{A}_i 's are orthogonal). Then there exists a rounding scheme for H .*

Moreover, this rounding scheme significantly simplifies the Hamiltonian.

Corollary 4.3 (Product structure from 2-local rounding). *Let H be a Hamiltonian with a register \mathbb{B} satisfying the requirements of Lemma 4.2. After rounding to generate \widetilde{H} , a register \mathbb{B} gets mapped to $\widetilde{\mathbb{B}} \subseteq \mathbb{B}$. Moreover, $\widetilde{\mathbb{B}}$ can be represented as a tensor product of registers $\widetilde{\mathbb{B}} = \otimes_i \widetilde{\mathbb{B}}_i$ such that each term of \widetilde{H} either act trivially act on $\widetilde{\mathbb{B}}$, or act on a unique sub-register $\widetilde{\mathbb{B}}_i$.*

For two terms, this is demonstrated in Figure 2. If H is 2-local, each register \mathbb{R}_i satisfies Lemma 4.2. Thus, after applying the rounding scheme to each register, Corollary 4.3 implies that the resulting Hamiltonian is made 1-local, and the ground state can be verified in classical polynomial time. Since the prover must provide the projector required for each step of the rounding scheme, this yields containment in NP. In fact, this implies a *guided reduction* for 2-local Hamiltonians:

Corollary 4.4 (Guided reduction for 2-local Hamiltonians (restatement of [BV04])). *There exists a guided reduction from 2-local Hamiltonians to 1-local Hamiltonians.*

Rounding scheme of Irani-Jiang In [IJ23], a rounding scheme is introduced to remove what the authors call “semi-separable qudits,” which are qudits for which there exists a non-trivial decomposition, induced by the Structure Lemma, such that all but a *single* term acts invariantly with respect to this decomposition. In a sense, this is a weakening of the assumption of Lemma 4.2, where we required a decomposition under which all terms acted invariantly. Let the projectors corresponding to this decomposition be $\{\pi_i\}_i$. Then, [IJ23] show,

Theorem 4.5 (Rounding scheme for 2D-CLH-Grid [IJ23]). *Let H an instance of 2D-CLH-Grid and $\{\pi_i\}_{i \in \ell}$ (with $\ell > 1$) be a set of projectors on a register \mathbb{B} such that all but one term commute with each projector. Then there is a rounding scheme for H and $\{\pi_i\}_i$, yielding another 2D-CLH-Grid instance $\widetilde{H} = \pi_i H \pi_i$.*

Since $\ell > 1$, the result of this rounding scheme is a reduction in the local dimension of \mathbb{B} . Thus, by iteratively applying this rounding scheme, we can remove all semi-separable qudits from H , yielding a guided reduction from general 2D-CLH-Grid instances to 2D-CLH-Grid-without-semi-separable-qudits instances.

Corollary 4.6. *There is a guided reduction from 2D-CLH-Grid to 2D-CLH-Grid-without-semi-separable-qudits, which iteratively applies Theorem 4.5.*

To obtain a containment in NP, the authors need to consider qutrits, rather than general qudit instances. They show that for *qutrit* Hamiltonians with no semi-separable qutrits, their single-register algebras can be well characterized and a careful case analysis shows that the resulting Hamiltonian can be made to have a 1D structure and thus solveable in NP.

4.2 An improved rounding scheme

In this section, we demonstrate a rounding scheme whenever there exists projectors π_1, π_2 , each of which commute with all but one term. We note that while the setting in which we apply our rounding scheme is incomparable to that of [IJ23], the improved rounding scheme we supply applies in all the situations where the rounding scheme of [IJ23] would apply, taking one of the projectors to be id, the identity operator. In Section 6, using tools developed in Section 5, we will see how this can be applied to the rank-1 setting to yield a guided reduction for 2D-CLH⁽¹⁾-Grid and 2D*-CLH⁽¹⁾. In its most general form, however, Theorem 4.7 works for general CLH instances.

Definition 20. *We say that a matrix P survives a projector π if $\pi P \pi \neq 0$.*

Theorem 4.7 (Rank-1 rounding scheme). *Let H be a sum of commuting local projectors and let π_1, π_2 be a pair of projectors such that for both π_1 and π_2 , there is at most a single term that both does not commute with it and survives the other projector; call this term π_1 's non-commuting operator P and π_2 's operator Q . Then there is a rounding scheme for H and $\{\pi_1, \pi_2\}$. Moreover, in the rounded Hamiltonian \tilde{H} , P and Q are replaced by a single merged term R , spanning the supports of both P and Q .*

Intuitively, this means that if both π_1 and π_2 are over a single register B , applying *both* π_1 and π_2 kills off all but two terms which act non-trivially on B .

Proof. Let P and Q be as defined in the theorem statement. Let H_{rest} be the remaining terms of H that survive both π_1 and π_2 (and thus commute with both). For clarity, we will re-label $\pi_2 = \pi_P$ and $\pi_1 = \pi_Q$ so it is clear which operators commute. We also define $\overline{H}_{\text{rest}} = \prod_{h \in H_{\text{rest}}} (\text{id} - h)$. First, from Lemma 2.1 we have that,

$$\lambda_0(H) = 0 \iff \text{Tr}[(\pi_P - \pi_P P \pi_P)(\pi_Q - \pi_Q Q \pi_Q) \overline{H}_{\text{rest}}] > 0. \quad (15)$$

Now, define $\tilde{P} = \pi_P - \pi_P P \pi_P$ and $\tilde{Q} = \pi_Q - \pi_Q Q \pi_Q$. Note that because P and π_P commute, \tilde{P} is a projector, and similarly for \tilde{Q} . By Jordan's lemma (Lemma 3.1), we can express $\tilde{P}\tilde{Q}$ as

$$\tilde{P}\tilde{Q} = \sum_b \eta_b |p_b\rangle\langle q_b|, \quad (16)$$

where b indexes the Jordan blocks on $(A_Q \cup A_P) \otimes B$. Here, since A_P and A_Q might overlap, we write the union to denote the Hilbert space spanned by both. Formally, each block $B \in \mathcal{B}$ is spanned by $|p_b\rangle$ and $|q_b\rangle$; B is 1-dimensional if $|p_b\rangle = |q_b\rangle$ and 2-dimensional otherwise. Let B' be the following subspace:

$$B' = \text{span}\{|p_b\rangle, |q_b\rangle : \eta_b > 0\}, \quad (17)$$

and Δ be the projection onto B' . Further define $\Pi_{B'}^P = \sum_{b:\eta_b>0} |p_b\rangle\langle p_b|$ and $\Pi_{B'}^Q = \sum_{b:\eta_b>0} |q_b\rangle\langle q_b|$. The rounded commuting local projector instance that our algorithm will output is $\tilde{H} = (\text{id}_{B'} - \Delta) + H_{\text{rest}}$. We prove that the Hamiltonian satisfies the following two properties:

1. (Commutation): $[\Delta, \overline{H}_{\text{rest}}] = 0$.
2. (Completeness and soundness): $\text{Tr}[\Delta \overline{H}_{\text{rest}}] > 0 \iff \text{Tr}[(\pi_P - P)(\pi_Q - Q) \overline{H}_{\text{rest}}] > 0$.

These two properties guarantee that the new Hamiltonian we output is a commuting projector instance, and that it is a valid rounding scheme. Let's begin by showing commutation.

Claim 1. Δ and H_{rest} commute.

Proof. Within each Jordan block of dimension 2, let $|\tilde{q}_b\rangle = \frac{(\text{id} - |p_b\rangle\langle p_b|)|q_b\rangle}{\|(\text{id} - |p_b\rangle\langle p_b|)|q_b\rangle\|_2}$ so that $|\tilde{q}_b\rangle \perp |p_b\rangle$. Define $\Pi_{\mathbb{B}'}^P = \sum_{b:\eta_b>0} |p_b\rangle\langle p_b|$ and $\Pi_{\mathbb{B}'}^Q = \sum_{b:\eta_b>0} |\tilde{q}_b\rangle\langle \tilde{q}_b|$. Orthogonality of each $|\tilde{q}_b\rangle$ and $|p_b\rangle$ implies $\Delta = \Pi_{\mathbb{B}'} = \Pi_{\mathbb{B}'}^P + \Pi_{\mathbb{B}'}^Q$. Then to prove the claim, it suffices to show that for all $h \in H_{\text{rest}}$, h commutes with both $\Pi_{\mathbb{B}'}^P$ and $\Pi_{\mathbb{B}'}^Q$.

By assumption, $h \in H_{\text{rest}}$ commutes with π_P and π_Q , and because we started with a commuting local Hamiltonian instance, h commutes with P and Q . Therefore, h commutes with any polynomial in \tilde{P} and \tilde{Q} . Specifically, consider $\tilde{P}\tilde{Q}\tilde{P}$. We can write this operator in terms of its Jordan blocks as:

$$\tilde{P}\tilde{Q}\tilde{P} = \sum_b \eta_b |p_b\rangle\langle p_b|. \quad (18)$$

Via Lemma B.1 we conclude that h commutes with $\text{round}(\tilde{P}\tilde{Q}\tilde{P})$, which is exactly $\Pi_{\mathbb{B}'}^P$.

Now we show that h commutes with $\Pi_{\mathbb{B}'}^Q$. We can always write $|q_b\rangle = \sqrt{\eta_b}|p_b\rangle + \sqrt{1-\eta_b}|\tilde{q}_b\rangle$. Because $\tilde{Q}\tilde{P}\tilde{Q}$ commutes with every h , we also have that for all polynomials $p: \mathbb{C} \rightarrow \mathbb{C}$ acting individually on each eigenvalue,

$$[h, p(\tilde{Q}\tilde{P}\tilde{Q})] = 0. \quad (19)$$

This essentially follows from Lemma B.1. Letting $p(x) = \frac{1}{\sqrt{x}}$, we evaluate $c(\tilde{Q}\tilde{P}\tilde{Q})$ as follows:

$$\begin{aligned} \sum_b \sqrt{\eta_b} |q_b\rangle\langle q_b| &= \sum_b \frac{1}{\sqrt{\eta_b}} \left(\eta_b |p_b\rangle\langle p_b| + \sqrt{\eta_b(1-\eta_b)} (|p_b\rangle\langle \tilde{q}_b| + |\tilde{q}_b\rangle\langle p_b|) + (1-\eta_b) |\tilde{q}_b\rangle\langle \tilde{q}_b| \right) \\ &= \sum_b \sqrt{\eta_b} |p_b\rangle\langle p_b| + \sqrt{1-\eta_b} (|p_b\rangle\langle \tilde{q}_b| + |\tilde{q}_b\rangle\langle p_b|) + \frac{1-\eta_b}{\sqrt{\eta_b}} |\tilde{q}_b\rangle\langle \tilde{q}_b|. \end{aligned} \quad (20)$$

Then we can evaluate the following function of \tilde{Q} and \tilde{P} :

$$p(\tilde{Q}\tilde{P}\tilde{Q}) - \tilde{Q}\tilde{P} - \tilde{P}\tilde{Q} + \sqrt{\tilde{P}\tilde{Q}\tilde{P}} = \sum_{b:\eta_b \neq 0} \frac{1-\eta_b}{\sqrt{\eta_b}} |\tilde{q}_b\rangle\langle \tilde{q}_b|. \quad (21)$$

Applying the rounding map to the above expression yields exactly $\Pi_{\mathbb{B}'}^Q$, and, therefore, all $h \in \overline{H}_{\text{rest}}$ commutes with $\Pi_{\mathbb{B}'}^Q$. Since Δ is the sum of $\Pi_{\mathbb{B}'}^Q$ and $\Pi_{\mathbb{B}'}^P$, we have that Δ commutes with $\overline{H}_{\text{rest}}$ as desired. This completes the proof of Claim 1. \square

Claim 2. $\text{Tr}[\Delta \overline{H}_{\text{rest}}] > 0$ if and only if $\text{Tr}[(\pi_P - \pi_P P \pi_P)(\pi_Q - \pi_Q Q \pi_Q) \overline{H}_{\text{rest}}] > 0$.

Proof. First assume that $\text{Tr}[\tilde{P}\tilde{Q}\tilde{H}_{\text{rest}}] > 0$. We know that \tilde{P} is a projector, so we have that $\tilde{P} = \tilde{P}^2$. Using the cyclic property of the trace and the fact that \tilde{P} commutes with $\overline{H}_{\text{rest}}$, we have that this condition is equivalent to

$$\text{Tr}[\tilde{P}\tilde{Q}\tilde{P}\overline{H}_{\text{rest}}] > 0. \quad (22)$$

Then because $\Delta \succeq \tilde{P}\tilde{Q}\tilde{P}$ we have that $\text{Tr}[\Delta \overline{H}_{\text{rest}}] \geq \text{Tr}[\tilde{P}\tilde{Q}\tilde{P}\overline{H}_{\text{rest}}] > 0$.

Now assume that $\text{Tr}[\tilde{P}\tilde{Q}\tilde{H}_{\text{rest}}] = 0$. We will show that both $\text{Tr}[\Pi_{\mathbb{B}'}^Q \overline{H}_{\text{rest}}] = \text{Tr}[\Pi_{\mathbb{B}'}^P \overline{H}_{\text{rest}}] = 0$. Let η_{\min} be the smallest non-zero eigenvalue of $\tilde{P}\tilde{Q}\tilde{P}$ so that $\Pi_{\mathbb{B}'}^P \leq \frac{1}{\eta_{\min}} \tilde{P}\tilde{Q}\tilde{P}$. It immediately follows that $\text{Tr}[\Pi_{\mathbb{B}'}^P \overline{H}_{\text{rest}}] \leq \frac{1}{\eta_{\min}} \text{Tr}[\tilde{P}\tilde{Q}\tilde{P}\overline{H}_{\text{rest}}] = 0$. Similarly, because \tilde{Q} is a projector, we have the following:

$$\begin{aligned} \text{Tr} \left[\sum_{b:\eta_b \neq 0} |q_b\rangle\langle q_b| \overline{H}_{\text{rest}} \right] &\leq \frac{1}{\eta_{\min}} \text{Tr}[\tilde{Q}\tilde{P}\tilde{Q}\overline{H}_{\text{rest}}] \\ &= \frac{1}{\eta_{\min}} \text{Tr}[\tilde{P}\tilde{Q}\overline{H}_{\text{rest}}] \end{aligned} \quad (23)$$

$$= 0.$$

Here the second line uses the definition of $\tilde{Q}\tilde{P}\tilde{Q}$, the third line uses the fact that \tilde{Q} commutes with $\overline{H}_{\text{rest}}$, the cyclic property of the trace, and that \tilde{Q} is a projector. As in Claim 2 we really want a bound on $\text{Tr}[\sum_b |\tilde{q}_b\rangle\langle\tilde{q}_b| \overline{H}_{\text{rest}}]$. Again we use Equation (21) to write,

$$\text{Tr} \left[\sum_{b: \eta_b \neq 0} \frac{1 - \eta_b}{\sqrt{\eta_b}} |\tilde{q}_b\rangle\langle\tilde{q}_b| \overline{H}_{\text{rest}} \right] = \text{Tr}[c(\tilde{Q}\tilde{P}\tilde{Q})\overline{H}_{\text{rest}}] - \text{Tr}[\tilde{P}\tilde{Q}\overline{H}_{\text{rest}}] - \text{Tr}[\tilde{Q}\tilde{P}\overline{H}_{\text{rest}}] + \text{Tr}[\sqrt{\tilde{Q}\tilde{P}\tilde{Q}\overline{H}_{\text{rest}}}],$$

and the above argument shows that each expression on the RHS is 0. Finally, we let $\eta'_{\min} = \min_{b: \eta_b \neq 0} \frac{1 - \eta_b}{\sqrt{\eta_b}}$ and bound,

$$\text{Tr}[\Pi_{\mathbb{B}'}^{\tilde{Q}} \overline{H}_{\text{rest}}] \leq \frac{1}{\eta'_{\min}} \text{Tr} \left[\sum_{b: \eta_b \neq 0} \frac{1 - \eta_b}{\sqrt{\eta_b}} |\tilde{q}_b\rangle\langle\tilde{q}_b| \overline{H}_{\text{rest}} \right] = 0,$$

Again noting that $\Delta = \Pi_{\mathbb{B}'}^P + \Pi_{\mathbb{B}'}^{\tilde{Q}}$, we conclude that $\text{Tr}[\Delta \overline{H}_{\text{rest}}] = 0$, completing the proof of Claim 2. \square

With the proofs of the two claims, the only remaining fact to check is that $\text{id}_{\mathbb{B}'} - \Delta$ is a projector. However since Δ is a sum of orthogonal projectors and all of the vectors are contained entirely in \mathbb{B}' , it must be a projector. Thus, $H_{\text{rest}} + (\text{id}_{\mathbb{B}'} - \Delta)$ is a rounding scheme for H and $\{\pi_1, \pi_2\}$. \square

5 Tools for rank-1 commuting operators

In this section we develop a characterization of the local algebras of rank-1 commuting operators. In particular, we show that rank 1 commuting projectors commute in one of two ways, either in a so called *singular*, or *reducing* way, which we will define below.

Definition 21 (Reducing commutation for rank 1 projectors). *Two rank 1 projectors P_{AB} and Q_{BC} commute in a reducing way if there exists a projector Π such that*

$$\begin{aligned} \Pi P \Pi &= P \text{ and} \\ \Pi Q \Pi &= 0. \end{aligned} \tag{24}$$

Definition 22 (Singular commutation for rank 1 projectors). *Two rank 1 projectors P_{AB} and Q_{BC} commute in a singular way if the following holds.*

$$\begin{aligned} P &= |\psi\rangle\langle\psi|_{\mathbb{B}} \otimes \tilde{P}_{\text{A}} \text{ and} \\ Q &= |\psi\rangle\langle\psi|_{\mathbb{B}} \otimes \tilde{Q}_{\text{C}}. \end{aligned} \tag{25}$$

Our goal for this section will be to prove the following lemma, although in reality we will prove a more general lemma than we need for the rest of the paper.

Theorem 5.1 (Commutation of rank 1 projectors). *Let P and Q be two rank 1 projectors. Then P and Q either commute in a singular (Definition 22) or reducing way (Definition 21).*

For the proof of this lemma, we will consider restricted algebras, where the restricting projector must also commute with the algebra. Formally,

Definition 23 (Subspace restrictions of an algebra). *Given an algebra \mathcal{A} and subspace Π that commutes with all operators of \mathcal{A} , we define the restriction of \mathcal{A} onto Π , $(\mathcal{A})|_{\Pi}$, to be the algebra generated by $\Pi h \Pi$, $h \in \mathcal{A}$. We sometimes abuse notation and write $(\mathcal{A})|_{\mathbb{R}}$ to mean that Π is the projection onto \mathbb{R} .*

Remark 3. Note that if a projector commutes with \mathcal{A} , every element of the algebra on a register can be written as a sum of two operators, one in Π and one in $\text{id} - \Pi$.

The condition that Π commutes with \mathcal{A} implies that the restriction onto Π yields another algebra (i.e. the resulting set is closed under sums and products). In the next section, we show a special case to the Structure Lemma (Lemma 3.5) in the case when one of the operators is rank 1. Intuitively, rank 1 operators “fully span” their support, meaning that any term commuting with a rank 1 projector must be trivial in their overlap.

Lemma 5.2 (Commutation of rank 1 projectors). *Let P_{AB} be a rank 1 projector and Q_{BC} be another projector that commutes with P . Let \mathcal{A}_P and \mathcal{A}_Q be the induced algebras by P and Q on \mathbb{B} respectively. Then there exists a decomposition of $\mathbb{B} = \mathbb{B}_P \oplus \mathbb{B}_Q$ such that*

1. $\mathcal{A}_P = \mathcal{L}(\mathbb{B}_P)$, i.e. P acts as the full algebra on \mathbb{B}_P and as 0 on \mathbb{B}_Q , and
2. $(\mathcal{A}_Q)|_{\mathbb{B}_P}$ is trivial, and Q acts non-trivially only on \mathbb{B}_Q .

Proof. Let $P = |\psi\rangle\langle\psi|$ for some state $|\psi\rangle$. We can write the Schmidt decomposition of $|\psi\rangle$ as

$$|\psi\rangle = \sum_i \alpha_i |\psi_i^1\rangle_{\mathbb{A}} |\psi_i^2\rangle_{\mathbb{B}}. \quad (26)$$

Then we can write P as follows

$$P = \sum_{i,j} \alpha_i \alpha_j^\dagger |\psi_i^1\rangle\langle\psi_j^1| \otimes |\psi_i^2\rangle\langle\psi_j^2|. \quad (27)$$

Since the local algebra is independent of the choice of decomposition (Lemma 3.3), the algebra \mathcal{A}_P is equal to $\mathcal{L}(\text{span}(\{|\psi_i^2\rangle\}_i))$. Defining $\mathbb{B}_P := \text{span}(\{|\psi_i^2\rangle\}_i)$ shows the first part of the claim.

Next we characterize $(\mathcal{A}_Q)|_{\mathbb{B}_P}$. By assumption, \mathcal{A}_Q commutes with \mathcal{A}_P . Since the projector onto \mathbb{B}_P , $\Pi_{\mathbb{B}_P}$ is in \mathcal{A}_P , \mathcal{A}_Q commutes with $\Pi_{\mathbb{B}_P}$. This implies that any element of \mathcal{A}_Q is block diagonal with respect to the subspaces \mathbb{B}_P and \mathbb{B}_P^\perp , and we can consider the restricted algebra $(\mathcal{A}_Q)|_{\mathbb{B}_P}$. Now, it remains to be seen that the restriction of \mathcal{A}_Q onto \mathbb{B}_P commutes with \mathcal{A}_P if \mathcal{A}_Q commutes with \mathcal{A}_P . As the only algebra that commutes with the full algebra on a subspace is the trivial algebra, this will complete the theorem.

Fix an element h of $(\mathcal{A}_Q)|_{\mathbb{B}_P}$. By definition $h = \Pi_{\mathbb{B}_P} h' \Pi_{\mathbb{B}_P}$ where $h' \in \mathcal{A}_Q$. Additionally, every element of \mathcal{A}_P is in the +1-eigenspace of $\Pi_{\mathbb{B}_P}$ (and therefore commutes with the projector). Then we have the following for all operators $g \in \mathcal{A}_P$:

$$\begin{aligned} [h, g] &= hg - gh \\ &= \Pi_{\mathbb{B}_P} h' \Pi_{\mathbb{B}_P} g - g \Pi_{\mathbb{B}_P} h' \Pi_{\mathbb{B}_P} \\ &= \Pi_{\mathbb{B}_P} h' g \Pi_{\mathbb{B}_P} - \Pi_{\mathbb{B}_P} g h' \Pi_{\mathbb{B}_P} \\ &= \Pi_{\mathbb{B}_P} [h', g] \Pi_{\mathbb{B}_P} \\ &= 0. \end{aligned} \quad (28)$$

Here we use the fact that g commutes with the projector, and h' commutes with g by the assumption that \mathcal{A}_Q commutes with \mathcal{A}_P . Therefore, the theorem is proved. \square

Combining Lemma 5.2 with the Structure Lemma (Lemma 3.5), we get the following corollary about the way that rank 1 projectors commute with neighboring terms.

Corollary 5.3. *Let P_{AB} be a rank 1 projector and Q be another projector that commutes with it and overlaps on \mathbb{B} . Then the direct sum decomposition from Lemma 3.5 consists of two subspaces, $\mathbb{R}_1 \oplus \mathbb{R}_2$ with $\mathbb{R}_1 = (\mathbb{B}_P \otimes \mathbb{B}_Q)$, where \mathbb{B}_P is the subspace implied by Lemma 5.2. Moreover, the following is true.*

1. Within \mathbb{R}_1 , Q looks like $\text{id}_{\mathbb{B}_P} \otimes \tilde{Q}$ where \tilde{Q} is supported on \mathbb{B}_Q .
2. Within \mathbb{R}_2 , P is 0.

Proof. We first apply the Structure Lemma, which implies that we can decompose the space \mathbb{B} into $\mathbb{R}_1 \oplus \mathbb{R}_2$. By Lemma 5.2 item (1), P_{AB} acts non-trivially only one of \mathbb{R}_1 or \mathbb{R}_2 ; assume without loss of generality it is \mathbb{R}_1 . Next, Lemma 5.2 further implies that the algebra generated by P_{AB} on \mathbb{B} is the full algebra on a subspace $\mathbb{B}_P \subseteq \mathbb{R}_1$. Take \mathbb{B}_Q to be the subspace such that $\mathbb{R}_1 = \mathbb{B}_P \otimes \mathbb{B}_Q$; such a subspace exists by the Structure Lemma. Furthermore, the algebra of Q on \mathbb{B}_P is the identity, and thus $Q|_{\mathbb{R}_1} = \text{id} \otimes \tilde{Q}$. Finally, since P_{AB} acts non-trivially on \mathbb{R}_1 , it must act as 0 on \mathbb{R}_2 (as otherwise it would have rank > 1). \square

Proof of Theorem 5.1. We apply Corollary 5.3 to the case of two rank-1 projectors. If \tilde{Q} in the above corollary is non-zero, then $\text{rank}(Q) \geq \dim(\mathbb{B}_P)$; this implies $\dim(\mathbb{B}_P) = 1$ and $\mathbb{B}_P = \text{span}(|\psi\rangle\langle\psi|)$ for some state $|\psi\rangle$. We conclude that $Q = |\psi\rangle\langle\psi| \otimes \tilde{Q}$ and similarly for P , as it is the full algebra on \mathbb{B}_P . This is exactly the singular case (Definition 22).

Otherwise, if $\tilde{Q} = 0$ then we may set Π to be the projection onto R_1 to match the reducing case (Definition 21). \square

Remark 4 (On rank larger than 1). One may wonder whether these techniques extend to the rank > 1 setting. When attempting such a generalization, it turns out that the *singular* case becomes challenging. However, if we were able to assume that Q is of a lower rank than $\dim(\mathbb{B}_P)$, we see from the above that the singular case can not happen. Given this (admittedly specific) constraint, we can generalize our proof to higher ranks.

Connection to rounding schemes As a quick example, let's see how to apply Theorem 5.1 in the degree-4, rank-1 case. Consider a qudit register R , acted on non-trivially by P_1, P_2 and Q_1, Q_2 , where the pairs (P_1, P_2) and (Q_1, Q_2) only interact on R (as in Figure 6). By Theorem 5.1, each pair commute in a *reducing* or *reducing* way. We notice that if both pairs commute in a reducing way, then the corresponding projectors π_P and π_Q satisfy the requirements of Theorem 4.7. This yields a new Hamiltonian where,

1. One of P_1, P_2 is removed, WLOG let it be P_2 .
2. One of Q_1, Q_2 is removed, WLOG let it be Q_2 .
3. The remaining terms P_1, Q_1 are combined into a single term h_{merge} which commutes with all other Hamiltonian terms.

This is depicted in Figure 6c. Thus, we get the following guided reduction which punctures a single hole in the Hamiltonian.

Corollary 5.4 (Guided reduction for rank-1 2D-CLH-Grid). *Suppose for a 2D-CLH-Grid instance H there exists a qudit register R on which the diagonal terms commute in a reducing way. Then, there exists a guided reduction to a Hamiltonian \tilde{H} where all terms acting trivially on R are unchanged, and there exists only a single term h_{merge} acting non-trivially on R , spanning the supports of two terms which were adjacent and acted non-trivially on R in H .*

In the next section, we'll scale this reduction up by considering the other *singular* commuting terms as well.

6 2D*-CLH⁽¹⁾ is in NP

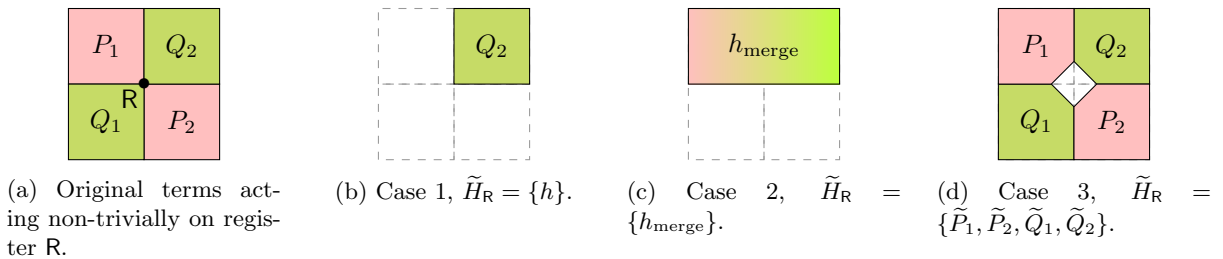


Figure 6: The original Hamiltonian and the possible resulting Hamiltonians from the rounding scheme of Lemma 6.3.

To begin, we can show that rank-1 2D-CLH instances are in NP for all local dimensions d . As noted before, our proof is through a *guided reduction* (Definition 1).

Theorem 6.1. *There is a guided reduction from rank-1 2D*-CLH⁽¹⁾ to 2-local CLH.*

Then, since any 2-local CLH is in NP by [BV04], we get the following simple corollary.

Corollary 6.2. *The local Hamiltonian problem for 2D*-CLH⁽¹⁾ is in NP.*

The proof of Theorem 6.1 is inspired by [AKV18]. However, our final result is incomparable as we handle the non-starred version of the problem for any local dimension (see Definition 13 versus Definition 14), but [AKV18] is able to handle Hamiltonian terms of any rank. As in [AKV18], we will show that the rank-1 condition allows us to create many holes in the 2D grid, so that after grouping qudit registers of dimension d into new registers of dimension at most $\mathcal{O}(d)$, the resulting Hamiltonian becomes 2-local. However, rather than creating holes by reduction to the Toric code, we use the rounding scheme developed in Theorem 4.7.

Rounding will allow us to prove following key lemma, which describes how to simplify a set of terms around a single register. For now we give a version of the lemma that only works for degree-4 registers; however, in Section 6.1 we'll show that a simple corollary of this lemma is a similar result for registers with higher degree.

Lemma 6.3 (Puncturing a 4-local register). *Let H be a Hamiltonian, and let R be a register acted on non-trivially $H_R = \{P_1, P_2, Q_1, Q_2\}$. So, $H = \sum_{h \in H_R} h + \sum_{h \in H_{rest}} h$, where H_{rest} are terms acting trivially on R . Furthermore, suppose that P_1 and P_2 intersect only on R , and same for Q_1, Q_2 . Then for any “way” the pair P_1 and P_2 and pair Q_1 and Q_2 commute (as in Theorem 5.1), there a rounding scheme yielding a set of at most 3 Hamiltonian terms \tilde{H}_R such that the Hamiltonian*

$$\tilde{H} = \sum_{h \in \tilde{H}_R} h + \sum_{h \in H_{rest}} h \quad \text{and} \quad \lambda_0(\tilde{H}) = 0 \iff \lambda_0(H) = 0 \quad (29)$$

Furthermore, the set \tilde{H}_R falls into one of three cases. See Figure 6 for matching figures.

1. $\tilde{H}_R = \{h\}$, with $h \in H_R$.
2. $\tilde{H}_R = \{h_{merge}\}$ with h_{merge} having support equal to a union of P_i, Q_j for some $i, j \in [2]$.
3. $\tilde{H}_R = \{\tilde{P}_1, \tilde{P}_2, \tilde{Q}_1, \tilde{Q}_2\}$ where R is removed from the support of each operator.

Corollary 6.4 (Degree $k > 4$ puncturing). *Lemma 6.3 holds whenever register R has degree $k > 4$. In this case, either R can be removed from the Hamiltonian, or all but 3 terms can be made to act trivially on R introducing a hole in the 2D surface.*

With this corollary, Theorem 6.1 follows identically as in [AKV18, Section 7].

Proof of Theorem 6.1. Fix some triangulation \mathcal{T} of H . For each triangle $T \in \mathcal{T}$, identify a qudit register R such that each neighbor of R is contained entirely in T . For each R , we apply Corollary 6.4, and the prover provides the projectors required by the rounding scheme. In each case of Lemma 6.3, we see that we puncture a hole adjacent to R . Thus, as in [AKV18], we place the corners of the co-triangulation at these holes, and draw paths from each hole to the centers of the edges of the surrounding triangles to construct the grouped register R_T (representing a higher-dimensional qudit). Given this grouping, terms either cross exactly one edge (and thus are 2-local) or act only on R_T (and are 1-local). This yields a guided reduction from 2D-CLH-Grid to 2-local CLH. \square

In the remainder of this section, we'll prove Lemma 6.3, deferring the proof of Corollary 6.4 to Section 6.1.

Proof of Lemma 6.3. As in Figure 6, consider the pairs P_1, P_2 and Q_1, Q_2 , each of which intersect only on register R . Then, the possible cases from Corollary 5.3 are,

- **(Case A)** Both decompositions are reducing.
- **(Case B)** Exactly one decomposition is reducing.
- **(Case C)** Neither decompositions are reducing. This implies $(P_1)|_R = (P_2)|_R = |\psi\rangle\langle\psi|$ and $(R_1)|_R = (R_2)|_R = |\varphi\rangle\langle\varphi|$.

In each of these cases, we'll show that we obtain one of the three cases from the lemma statement. Since P_1, P_2 commute and intersect only on R , the Structure Lemma induces a direct sum decomposition of $\mathcal{H}_R = \oplus_i (\mathcal{H}_R)_i$. We denote the projector onto one of these subspaces provided by the prover as π_P . Similarly, we obtain a projector π_Q , from the decomposition induced by Q_1 and Q_2 .

Case A (Both reducing). In the case when both are reducing, we claim that we can apply Theorem 4.7. In particular, by the definition of reducing, one of two Q terms does not survive π_Q , leaving only a single term that does not commute with π_P , and similarly only a single P term survives π_P , leaving a single P term which does not commute with π_Q . We let \tilde{H}_R be the Hamiltonian produced by Theorem 4.7. The theorem guarantees that \tilde{H}_R satisfies Item 2.

Case B (One reducing). Without loss of generality, say P_1, P_2 commute in a singular manner (so $P_1 = \tilde{P}_1 \otimes |\psi\rangle\langle\psi|$ and $P_2 = \tilde{P}_2 \otimes |\psi\rangle\langle\psi|$) and Q_1, Q_2 commute in a reducing manner. If π_P is onto $\mathcal{H}_{\text{rest}}$, then we again left with a single term, leaving us in Item 1. If π_Q is onto $\mathcal{H}_{\text{rest}}$, then all remaining terms act as $|\psi\rangle\langle\psi|$ on R , and R is a classical register. Thus, by Observation 3.7 we can obtain an equivalent Hamiltonian where each term acts trivially on R and we obtain Item 3). Otherwise, we assume $\pi_P = |\psi\rangle\langle\psi|$ and π_Q preserves Q_1 . Thus (by Lemma 2.1), we it suffices to analyze

$$\text{Tr} \left[(\psi - \tilde{P}_1 \otimes \psi)(\psi - \tilde{P}_2 \otimes \psi)(\pi_Q - Q_1) \prod_{h \in S} (\mathbb{I} - h) \right]. \quad (30)$$

Now, pull out a ψ from either side of the two P -type terms, and conjugate $\pi_Q - Q_1 \rightarrow \alpha \cdot \psi - \psi Q_1 \psi$ (where $\alpha = \text{Tr}[\pi_Q \psi]$). Since ψ and $\psi Q_1 \psi$ are simultaneously diagonalizable, we can use the rounding technique from [LJ23] to round $\alpha \psi - \psi Q_1 \psi \rightarrow \psi - \mathcal{R}(\psi Q_1 \psi)$ ³ which yields a projector commuting with P_1, P_2 , as well as all remaining terms $h \in S$. Thus, the register R again is classical, and we obtain Item 3.

Case C (Both singular). First, if one of π_P or π_Q project onto $\mathcal{H}_{\text{rest}}$, then the q becomes classical. Otherwise we can write,

$$\begin{aligned} P_1 &= |\psi\rangle\langle\psi| \otimes \tilde{P}_1 & P_2 &= |\psi\rangle\langle\psi| \otimes \tilde{P}_2 \\ Q_1 &= |\varphi\rangle\langle\varphi| \otimes \tilde{Q}_1 & Q_2 &= |\varphi\rangle\langle\varphi| \otimes \tilde{Q}_2 \end{aligned} \quad (31)$$

with $\pi_P = |\psi\rangle\langle\psi|$ and $\pi_Q = |\varphi\rangle\langle\varphi|$. By commutation of P_1, P_2 and Q_1, Q_2 , we have that $[\tilde{P}_1, \tilde{P}_2] = [\tilde{Q}_1, \tilde{Q}_2] = 0$. Similarly, since each $h \in S$ acts trivially on R , their commutation with P_1, P_2, Q_1, Q_2 is equivalent to their commutation with the tilde-d terms. Once again we can apply π_P and π_Q to the P and Q -type terms respectively via Lemma 2.1, and obtain,

$$\begin{aligned} & \text{Tr} \left[(\psi - \psi \otimes \tilde{P}_1)(\psi - \psi \otimes \tilde{P}_2)(\varphi - \varphi \otimes \tilde{Q}_1)(\varphi - \varphi \otimes \tilde{Q}_2) \prod_{h \in S} (\text{id} - h) \right] \\ &= \text{Tr} \left[\varphi_R \psi_R (\text{id} - \text{id} \otimes \tilde{P}_1) \left(\text{id}_R \otimes (\text{id} - \tilde{P}_2) \right) \psi_R \varphi_R (\text{id} - \text{id} \otimes \tilde{Q}_1) \left(\text{id}_R \otimes (\text{id} - \tilde{Q}_2) \right) \prod_{h \in S} (\mathbb{I} - h) \right] \\ &= \langle \psi | \varphi \rangle^2 \text{Tr} \left[\langle \psi |_R (\text{id} - \text{id} \otimes \tilde{P}_1) | \psi \rangle_R \cdot (\text{id} - \tilde{P}_2) \langle \varphi |_R (\text{id} - \text{id} \otimes \tilde{Q}_1) | \varphi \rangle_R \cdot (\text{id} - \tilde{Q}_2) \prod_{h \in S} (\text{id} - h) \right] \\ &= \langle \psi | \varphi \rangle^2 \text{Tr} \left[(\text{id} - \tilde{P}_1)(\text{id} - \tilde{P}_2)(\text{id} - \tilde{Q}_1)(\text{id} - \tilde{Q}_2) \prod_{h \in S} (\text{id} - h) \right] \end{aligned} \quad (32)$$

where in the second line, we have pulled out the ψ, φ projectors, and used the cyclic property of the trace to pull one φ to the LHS. To obtain the third line, we take the trace with respect to R . If the original

³More precisely, we note that $\alpha \cdot \psi \succeq \psi Q_1 \psi$. Then, pull the positive constant α out of the trace. Then, within each non-zero eigenspace of $\psi Q_1 \psi$, ψ has eigenvalue 1. Finally, the rounding map \mathcal{R} takes each non-zero eigenvalue to 0, which by [LJ23] yields an equivalent trace expression.

Hamiltonian had a non-empty ground space, then there is a choice of ψ, φ such that $\langle \psi | \varphi \rangle^2 > 0$. Besides the constant factor, checking whether the trace is greater than 0 is equivalent to the condition that

$$\lambda_0(\tilde{H}) = 0 \quad \text{where} \quad \tilde{H} = \tilde{P}_1 + \tilde{P}_2 + \tilde{Q}_1 + \tilde{Q}_2 + \sum_{h \in S} h, \quad (33)$$

and \tilde{H} is a Hamiltonian which does not act on R . Moreover, this Hamiltonian is commuting, by our observation earlier. Thus, we have rounded to a CLH instance satisfying Item 3. \square

6.1 Extending to degree $k > 4$

We now show how Corollary 6.4 follows from Lemma 6.3.

Corollary 6.4 (Degree $k > 4$ puncturing). *Lemma 6.3 holds whenever register R has degree $k > 4$. In this case, either R can be removed from the Hamiltonian, or all but 3 terms can be made to act trivially on R introducing a hole in the 2D surface.*

Recall that the degree of a register is the number of terms acting non-trivially on it. We can extend the above proof to work for registers with degree larger than 4. Since we are working with planar, 2D Hamiltonians, we may index the terms for a single register R as H_1, \dots, H_k , such that for each H_i , H_i shares two registers with H_{i-1} and H_{i+1} , and only the register R for H_j , $|i - j| > 1$ ⁴. Define two sets, \mathcal{P} and \mathcal{Q} such that,

$$\begin{aligned} \mathcal{Q} &= \{H_{2i-1} : i \in \lfloor k/2 \rfloor\} \\ \mathcal{P} &= \{H_{2i} : i \in \lfloor k/2 \rfloor\} \end{aligned} \quad (34)$$

i.e. \mathcal{P} contains the even indexed H_i 's, and \mathcal{Q} contains the odds. If k is even then each pair $P_1, P_2 \in \mathcal{P}$ and $Q_1, Q_2 \in \mathcal{Q}$ share only a single register. On the other hand, if k is odd, then Q_1 and Q_k intersect on 2 registers. In this case, we redefine $Q_1 \triangleq Q_1 + Q_k$, yielding a single rank 2 term. Since Corollary 5.3 works as long as all but a single term is rank-1, we may assume that k is even.

Since each $P_i, P_j \in \mathcal{P}$ only intersect on a single register, we can analyze their local algebra on R using the Structure Lemma. By iteratively applying Lemma 5.2, we see that each P_i has an induced algebra \mathcal{A}_{P_i} which is full on some subspace. Moreover, these subspaces are orthogonal unless they are 1-dimensional. Group the operators $P_i \in \mathcal{P}$ into sets $S_1^{\mathcal{P}}, \dots, S_\ell^{\mathcal{P}}$ such that every $S_i^{\mathcal{P}}$ has a corresponding register R_i with $\mathcal{A}_{P_i} = \mathcal{L}(R_i)$ for each $P_i \in S_i^{\mathcal{P}}$. These subspaces further satisfy the following properties

1. The R_i 's partition the Hilbert space of register R , i.e. $R = \bigoplus_{i=1}^{\ell} R_i$.
2. If $|S_i^{\mathcal{P}}| > 1$, then $\dim(R_i) = 1$.

The second condition follows by applying Corollary 5.3 to the elements that are the full algebra on the same subspace. Similarly, we may define sets $S_i^{\mathcal{Q}}$. Here we point out that if we have a rank-2 term from k being odd, the second condition still holds. In particular, if there is another rank 1 projector in $S_i^{\mathcal{P}}$, then both terms in the sum $Q_1 + Q_k$ must individually commute with it, and thus the dimension of $R_i^{\mathcal{Q}}$ must be 1 by Lemma 5.2.

The condition that our Hamiltonian as a 0-energy ground space is equivalent to the condition that,

$$\text{Tr} \left[\prod_{P \in \mathcal{P}} (\mathbb{I} - P) \prod_{Q \in \mathcal{Q}} (\mathbb{I} - Q) \prod_{h \in S} (\mathbb{I} - h) \right] > 0. \quad (35)$$

Write $\Pi_a^{\mathcal{P}}$ and $\Pi_b^{\mathcal{Q}}$ as the projectors onto $\mathcal{H}_a^{\mathcal{P}}$ and $\mathcal{H}_b^{\mathcal{Q}}$ respectively. By Lemma 2.1, we can further write the equivalent statement,

$$\exists a, b \text{ s.t. } \text{Tr} \left[\prod_{P \in \mathcal{P}} (\Pi_b^{\mathcal{P}} - \Pi_b^{\mathcal{P}} P \Pi_b^{\mathcal{P}}) \prod_{Q \in \mathcal{Q}} (\Pi_a^{\mathcal{Q}} - \Pi_a^{\mathcal{Q}} Q \Pi_a^{\mathcal{Q}}) \prod_{h \in S} (\mathbb{I} - h) \right] > 0 \quad (36)$$

⁴Arithmetic here is performed modulo k , so that $k + 1 \equiv 1$

This choice of a, b zeroes out all terms except those in $S_a^{\mathcal{P}}$ and $S_b^{\mathcal{Q}}$. At this point, we can case on the size of each set, resembling the cases **A**, **B**, and **C** considered in the previous section.

- $|S_a^{\mathcal{P}}| = |S_b^{\mathcal{Q}}| = 1$. In this case, we're left with a single Q -type and P -type term, and we can applying the analysis of **Case A**.
- $|S_a^{\mathcal{P}}| = 1$ and $|S_b^{\mathcal{Q}}| > 1$. This implies that each $Q \in S_b^{\mathcal{Q}}$ is of the form $\varphi \otimes \tilde{Q}$, for a fixed state φ . This corresponds to **Case B**, yielding a classical qudit. The case when $|S_a^{\mathcal{P}}| > 1$ and $|S_b^{\mathcal{Q}}| = 1$ is identical.
- $|S_a^{\mathcal{P}}|, |S_b^{\mathcal{Q}}| > 1$. Then each non-zero $P = \psi \otimes \tilde{P}$ and $Q = \varphi \otimes \tilde{Q}$. This corresponds to case **Case C**, and we can conjugate each Q term with ψ . This yields a classical qudit.

7 Commuting Hamiltonians in three dimensions

In this section, we show the rank-1 assumption allows us to place commuting Local Hamiltonians over a 3D complex in NP. We recall the following important assumption,

Assumption 2 (Generalized cubic lattice). *We maintain some connection with the simpler cubic lattice case. This connection is through two restrictions on the lattice.*

1. *If there are two Hamiltonian terms h_1, h_2 sharing a face, then there is no term which shares a face with both h_1 and h_2 . Intuitively, this prevents “Jenga-like” structures, which will make it hard to draw an analogue to the tunnel picture (Figure 9) from the simpler case.*
2. *Each Hamiltonian term should have at least 5 faces. This condition will ensure that when we punctures holes for the vertex, there are enough faces so that each boundary path (e.g. the lines in Figure 11a) can exit through a distinct face.*

Formally, we state our main result, and its immediate corollary, as follows.

Theorem 7.1 (Guided reduction for 3D*-CLH⁽¹⁾). *Let H be an instance of 3D*-CLH⁽¹⁾ with degree g (Definition 6) on d -dimensional registers. Suppose we also have a bound k on the locality of any term. Then,*

- *If $g \leq \frac{1}{\epsilon} d^k$, then by Lemma 3.2 the instance H is trivially satisfiable.*
- *Otherwise, assume $g > \frac{1}{\epsilon} d^k \geq 6$. If H satisfies Assumption 1 (large minimum degree) and Assumption 2, then there is a guided reduction from H to a 2-local commuting local Hamiltonian.*

Corollary 7.2 (Local Hamiltonian problem on 3D*-CLH⁽¹⁾ in NP). *For non-trivial 3D*-CLH⁽¹⁾ instances meeting the requirements of Theorem 7.1, the local Hamiltonian problem on these instances is in NP.*

Note that the *cubic lattice* implies a degree bound of $g = 4$ and locality $k = 8$, for which the first case of Theorem 7.1 applies (and thus the instance is trivially satisfiable). However, most of the ideas for the general case are extensions of ideas from the cubic lattice, so we will first describe our proof on this more restricted setting.

Our strategy in this section is to induce a “cubulation” in the 3D space, where a lattice of cubes is superimposed over the Hamiltonian terms, as in Figure 7. We refer to the set of larger cubes as \mathcal{C} . Given a cubulation \mathcal{C} , we can classify terms in the original Hamiltonian as into 3 “types:” those intersecting a face of the cubulation, intersecting an edge, or intersecting a vertex. Those intersected a face are automatically 2-local. In the 2D setting terms intersecting edges were also 2-local. In the 3D case, these are 4-local. Therefore, we will need to use our rounding scheme from Corollary 6.4 to create “tunnels” so that the edges can pass through holes in the Hamiltonian. This still is not quite sufficient, as there can blockages in the generated tunnels, see Figure 10b. To address this issue, we use the fact that the result of rounding schemes are also valid CLH instances and that these blockages are essentially 2-local interactions. Therefore, we will be able to apply the 2-local rounding scheme in the form of Corollary 4.3 to fix this case.




Figure 7: An example of a cubulation. The blue edges represent the cubulation, where the qubits (on edges) inside each blue cube are grouped into a single qudit. The green shading emphasizes some terms intersecting the cubulation.

7.1 3D cubic lattice

In the section, we consider the special case of $3D^*$ -CLH⁽¹⁾-Grid. As mentioned above, we superimpose a cubulation \mathcal{C} over the 3D space, as in Figure 7. Given a cubulation \mathcal{C} , we can classify the Hamiltonian terms into 3 types. We use Figure 8a as a reference.

1. Terms intersecting a *face* of the cubulation (e.g. contained in S_4).
2. Terms intersecting an *edge* of the cubulation (e.g. terms contained in S_1).
3. Terms intersecting a *vertex* of the cubulation (e.g. terms contained in S_0 and S_2).


Type 1 terms are easy; these are automatically 2-local as faces are shared by at most two adjacent $C, C' \in \mathcal{C}$. Type 2 and 3 are harder. For these remaining terms, the idea is to apply our rounding technique from the 2D case to a “slice” of terms, i.e. Figure 8b. The key observations is that from the point of a qudit register R on the center edge, the local geometry “looks like” the 2D grid. As a result, we are in a setting where we can apply Lemma 6.3, restated here for convenience.

Lemma 6.3 (Puncturing a 4-local register). *Let H be a Hamiltonian, and let R be a register acted on non-trivially $H_R = \{P_1, P_2, Q_1, Q_2\}$. So, $H = \sum_{h \in H_R} h + \sum_{h \in H_{rest}} h$, where H_{rest} are terms acting trivially on R . Furthermore, suppose that P_1 and P_2 intersect only on R , and same for Q_1, Q_2 . Then for any “way” the pair P_1 and P_2 and pair Q_1 and Q_2 commute (as in Theorem 5.1), there a rounding scheme yielding a set of at most 3 Hamiltonian terms \tilde{H}_R such that the Hamiltonian*

$$\tilde{H} = \sum_{h \in \tilde{H}_R} h + \sum_{h \in H_{rest}} h \quad \text{and} \quad \lambda_0(\tilde{H}) = 0 \iff \lambda_0(H) = 0 \quad (29)$$

Furthermore, the set \tilde{H}_R falls into one of three cases. See Figure 6 for matching figures.

1. $\tilde{H}_R = \{h\}$, with $h \in H_R$.
2. $\tilde{H}_R = \{h_{merge}\}$ with h_{merge} having support equal to a union of P_i, Q_j for some $i, j \in [2]$.
3. $\tilde{H}_R = \{\tilde{P}_1, \tilde{P}_2, \tilde{Q}_1, \tilde{Q}_2\}$ where R is removed from the support of each operator.



(a) The light-gray dashed lines indicate the individual Hamiltonian terms. The black boxes depict the Hamiltonian terms intersecting the boundaries of the “cubulation” (boundaries depicted by the red, green, and blue lines). The shaded in terms indicates a “slice” (denoted S_i) containing a term intersecting a boundary of the cubulation. Notice S_4 only intersects the plane defined by the green and blue lines.

(b) A single slice centered around a register R .

Figure 8: A local view at a cubulation.






Figure 9: The set of “slices” S_i around an edge of the cubulation (in blue). Each slice is centered around the register R_i . The Hamiltonian term containing the edge is colored.




(a) The slice for R_2 falls into Item 3, and the slices for R_1 and R_0 fall into Item 2. The orientation of the merged blocks permits the boundary to pass through.




(b) The orientation of the merged terms does not leave a hole for the boundary to pass through.


Figure 10: Given a set of terms with a cubulation edge (in bold black) passing through, we can apply Lemma 6.3 to attempt to create a hole. Depending on the orientation, this process may fail.



(a) A slice S_0 containing the vertex term, which yields the merged term h . The dotted blue, red, and green lines indicate the boundaries of the cubulation. The solid lines indicate the edges on which Lemma 6.3 will be applied.



(b) An example of the slices on which we'll apply Lemma 6.3.



(c) An example of possible “blockages.” The merged blue terms on the left and right block the blue boundary from pass through an empty region.

Figure 11: The process of constructing tubes of slices for a vertex of the cubulation.

Edge-type terms. Consider the set of terms S_1 ; as indicated in Figure 8a, we pick the slice so that it is perpendicular to the (blue) edge. These slices look like Figure 9. For each slice we will apply Lemma 6.3. As long as the reduced slices permit a (possibly curving) boundary to pass through empty areas of the geometry, we ensure that only 2-local terms remain, as the remaining terms pass through *faces* of the cubulation (and thus are 2-local). One can check that in most cases, two adjacent slices to create such a hole:

- Type 1 leaves a hole when adjacent to *any* other type.
- Type 3 removes the center register and thus always leaves a hole.

The trouble is when there are two adjacent Type 2 slices, and specifically, when the terms line up as in h_0 and h_1 in Figure 10b (it's easy to see that any other orientation also leaves a hole). However, we now use a crucial aspect of Lemma 6.3, which is that the merged terms commute with all terms outside of the slice. Under this type of blockage, the (green and blue) merged terms are the *only* terms acting non-trivially on their shared edge and thus can be viewed as 2-local operators acting on their intersection. Moreover, all other terms act trivially there. Thus, we can directly apply Corollary 4.3 which provides a rounding scheme creating a hole between h_1 and h_0 .

Vertex-type terms. As before, we'll apply Lemma 6.3 to a slice with the term containing the vertex (slice S_0 in Figure 11). The corner is then placed in the hole created by the lemma. The next step is to apply puncturing again to surround qudit registers to create holes for the cubulation boundaries. In Figure 11a, we draw the (dotted) boundaries, and describe the edges (registers) on which Lemma 6.3 will be applied to create holes. Two things to call out,

1. The edges are chosen so that slices intersect on at most one term. This prevents “merged” terms from Lemma 6.3 impacting the other edges. For instance, although the bottom green register intersects h (and possibly the merged term of the right blue boundary), its placement ensures the green register remains degree-4.
2. For the north/south green boundary, we do not pick the edge pointing straight down, as the merged term h blocks our path. Thus, we first go “forwards,” then eventually turn back down.

For slices which are adjacent to two removed terms from S_0 (e.g. the top green, and both red boundaries), no special care is needed; we can apply Corollary 4.3 to puncture holes if obstructions arise. However, there are also slices which are perpendicular to S_0 , see, e.g. Figure 11c. For these, it's not immediately clear the 2-local rounding scheme applies. Nonetheless, we can show that as long as a slice is adjacent to *at least one* removed term of a previous slice, we can make a continuous hole. This will allow us to puncture holes for the blue and bottom green edges. For reference, the names in the following lemma are chosen to line up with the right volume in Figure 11c.

Lemma 7.3 (General blockages). *Suppose there exists a slice S_1 which is adjacent (via a face) to a removed term h_{removed} . After applying Lemma 6.3 (or Corollary 6.4) to S_1 , we either immediately have punctured a continuous hole, or there is a further projection that can be applied to an edge adjacent to h_{removed} and S_1 which punctures a hole from h_{removed} through S_1 .*

Proof. Let f be a face of the removed term h_{removed} , from which we'll attempt to pass through S_1 . Let h be the (original) term of S_1 adjacent to f . Assume we have applied Corollary 6.4 on R_3 , the central register of S_1 . If h is removed via Corollary 6.4, then we are done. Otherwise, we make the two geometric observations:

- There are at most two edges on f (R_1, R_2) which are adjacent to R_3 . This is because any third edge would subdivide f .
- For any edge R on f , there are at most 2 terms on which h shares more than just the edge R . This is because any edge induces a slice, sandwiching h by two other terms.

Since h_{removed} is adjacent to h via f (and thus via R_1, R_2), and there exists two distinct terms h_1, h_2 which are adjacent to h *within* the slice S_1 , the second item implies that

- Across R_1 , h is adjacent to *only* h_{removed} and some $h_1 \in S_1$ on more than just R_1 .
- Across R_2 , h is adjacent to *only* h_{removed} and some $h_2 \in S_2 \setminus \{h_1\}$ on more than just R_2 .

But Corollary 6.4 merges h with at most one adjacent term in S_1 , implying either h_1 or h_2 are removed. Without loss of generality, assume h_2 is removed. This implies that every other term acting non-trivially on R_2 intersects h *only* on R_2 and thus across R_2 , all terms can be treated as 2-local. As before, we apply Corollary 4.3 to create a hole through this edge. \square

Beyond the obstruction depicted in Figure 11c, this also permits us to “turn,” as required by the green boundary in Figure 11a.




Figure 12: The green slice and “turning” green boundary. One can check that applying Lemma 7.3 with the above choice of h_{removed} and $S_1 \triangleq S$, we puncture a hole through S .

7.2 General 3D lattices


We now turn to the second case in Theorem 7.1. Although similar ideas from the cubic lattice case work here, there are several components which need to be generalized. First, the idea of a tube of slices is not well-defined. One idea is to choose some path $P = (e_1, e_2, \dots)$ and attempt to puncture a hole along this edge (e.g. by applying Lemma 6.3 to each edge e_i). This has the undesirable property that the resulting “cubulation” depends heavily on the individual Hamiltonian terms, unlike in the proof of [AE11], where the triangulation can be made independent of the Hamiltonian. To emulate this idea, we partition the 3D space in a cubic grid, large enough that within each element C of the grid \mathcal{C} , there is a term (volume) h_C such that all neighboring volumes are fully contained within C . We’ll use this center term h_C to define the co-cubulation. At a high level, this involves selecting a center of grid element C , and then drawing paths from the center to the geometric centers of each face of C . This is essentially the co-triangulation idea of [AE11]. We now show how to pick the centers and paths so that the Hamiltonian can be made 2-local. We recall the key assumption we make in this section for convenience:

Assumption 2 (Generalized cubic lattice). *We maintain some connection with the simpler cubic lattice case. This connection is through two restrictions on the lattice.*

1. *If there are two Hamiltonian terms h_1, h_2 sharing a face, then there is no term which shares a face with both h_1 and h_2 . Intuitively, this prevents “Jenga-like” structures, which will make it hard to draw an analogue to the tunnel picture (Figure 9) from the simpler case.*
2. *Each Hamiltonian term should have at least 5 faces. This condition will ensure that when we punctures holes for the vertex, there are enough faces so that each boundary path (e.g. the lines in Figure 11a) can exit through a distinct face.*

Puncturing edges As in Section 7.1, we’d like to define a notion of a “tube” of slices. In the cubic lattice case, such a tube is essentially defined by a sequence of faces, F_1, F_2, \dots , such that there is a term h sharing adjacent faces F_i, F_{i+1} . Then, the slice S_i will be defined by an edge (register) R_i connecting F_i and F_{i+1} . The terms of the slice are those volumes which act non-trivially on R_i . This generalizes the notion of tubes, and we can apply Corollary 6.4 to each R_i . There are several complications which arise in the general 3D case.

1. Two faces F_i, F_{i+1} share an edge.
2. A volume of S_i is not directly adjacent to a volume in S_{i+1} , as in Figure 13a
3. There are no adjacent removed terms between slices. This generalizes Figure 10b and is depicted in Figure 13c.



(a) Two adjacent terms sharing a face. In this example, two of the surrounding (pink) terms do not also share a face. However, by treating the terms in the blue region, together with the right pink term, as one larger term, we can recover the “cubic” structure of Section 7.1.

(b) A sequence of faces where adjacent faces share an edge.

(c) One might imagine that we could have completely disjoint removed terms, such that we cannot even apply Corollary 4.3. However, this is avoided by the number of terms guaranteed to be trivial by Corollary 6.4.

Figure 13: Complications in general 3D geometries.

Item 1 is easily dealt with; any two non-equal faces sharing a volume must have *some* edge separating them so we can apply Corollary 6.4 to this edge instead.

For Item 2, we may no longer have a direct path between the two slices, e.g. if the pink terms in Figure 13a are removed. However, in this case, any term in the blue region does not intersect h or h' on a face, as this would violate the first item of Assumption 2. Since h and h' both commute with any term in the blue region, there exists a projector on R_3 which decouples these two terms and induces a hole from face F_{k-1} to F_{k+1} .

For Item 3, Corollary 6.4 ensures that (the regions corresponding to) 3 terms are non-trivial. Since the local degree is at least 6, this implies that the 3 non-trivial terms can only create a blockage when they do not share any faces. But then we can directly apply Corollary 4.3 to create a hole.

Puncturing vertices Choose an edge R (i.e. a qudit register) of the term h_C , and consider the set of terms acting non-trivially on R . Since we are in the high degree setting, these form a slice S_0 with degree > 4 , and we can apply Corollary 6.4. This ensures that there are at least 2 face-adjacent terms which are now trivial in the resulting Hamiltonian and we can place the corner of the co-cubulation in this hole. We now need to route 6 paths through the surrounding faces. Suppose there are k removed terms, and each term has at least f faces. Since each term is adjacent to two other terms in the slice, 2 faces from each term are internal. Thus, the number of faces surrounding the hole is at least,

$$N_{\text{faces}} \geq (f - 2) \cdot k \quad (37)$$

By Item 2 of Assumption 2, each term has at least 5 faces, and k is at least 2 so $N_{\text{faces}} \geq 6$, which is exactly the number of distinct edges exiting the corner. We now can make an argument similar to Section 7.1. First, we pick the edges R_1, \dots, R_6 so that each pair $R_i \neq R_j$ are not adjacent via a single edge. It’s possible that R_i, R_j share a single face (and thus one volume), but any other volume would need to span two edges to intersect both both registers, which violates convexity. One can check that these edges can be chosen in this way, by imagining flattening out the N_{faces} exposed faces of S_0 and pick 6 vertices to serve as the origins of each edge on the resulting 2D surface.

Once this is done, as in the cubic case, any edge adjacent to two removed terms can be easily dealt with. For edges adjacent to only single term, one can check that the observations made in the proof of Lemma 7.3 still hold in the general 3D lattice case and therefore holes can be made through these terms as well.

References

- [ABN23] Anurag Anshu, Nikolas P Breuckmann, and Chinmay Nirkhe. “NLTS Hamiltonians from good quantum codes”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023, pp. 1090–1096.
- [AE11] Dorit Aharonov and Lior Eldar. “On the Complexity of Commuting Local Hamiltonians, and Tight Conditions for Topological Order in Such Systems”. In: *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*. 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science. Oct. 2011, pp. 334–343. DOI: [10.1109/FOCS.2011.58](https://doi.org/10.1109/FOCS.2011.58). URL: <https://ieeexplore.ieee.org/abstract/document/6108194> (visited on 11/01/2023).
- [AE15] Dorit Aharonov and Lior Eldar. “The Commuting Local Hamiltonian Problem on Locally Expanding Graphs Is Approximable in NP ”. In: *Quantum Information Processing* 14.1 (Jan. 1, 2015), pp. 83–101. ISSN: 1573-1332. DOI: [10.1007/s11128-014-0877-9](https://doi.org/10.1007/s11128-014-0877-9). URL: <https://doi.org/10.1007/s11128-014-0877-9> (visited on 10/25/2023).
- [AKS12] Andris Ambainis, Julia Kempe, and Or Sattath. “A Quantum Lovasz Local Lemma”. In: *Journal of the ACM* 59.5 (Oct. 2012), pp. 1–24. ISSN: 0004-5411, 1557-735X. DOI: [10.1145/2371656.2371659](https://doi.org/10.1145/2371656.2371659). arXiv: [0911.1696 \[quant-ph\]](https://arxiv.org/abs/0911.1696). URL: <http://arxiv.org/abs/0911.1696> (visited on 07/10/2024).
- [AKV18] Dorit Aharonov, Oded Kenneth, and Itamar Vigdorovich. *On the Complexity of Two Dimensional Commuting Local Hamiltonians*. Mar. 29, 2018. DOI: [10.48550/arXiv.1803.02213](https://arxiv.org/abs/1803.02213). arXiv: [1803.02213 \[quant-ph\]](https://arxiv.org/abs/1803.02213). URL: <http://arxiv.org/abs/1803.02213> (visited on 10/02/2023). Pre-published.
- [Bar+23] Ivan Bardet et al. “Rapid thermalization of spin chain commuting Hamiltonians”. In: *Physical Review Letters* 130.6 (2023), p. 060401.
- [BV04] S. Bravyi and M. Vyalyi. *Commutative Version of the K-Local Hamiltonian Problem and Common Eigenspace Problem*. Dec. 16, 2004. DOI: [10.48550/arXiv.quant-ph/0308021](https://arxiv.org/abs/quant-ph/0308021). arXiv: [quant-ph/0308021](https://arxiv.org/abs/quant-ph/0308021). URL: <http://arxiv.org/abs/quant-ph/0308021> (visited on 10/02/2023). Pre-published.
- [CM16] Toby Cubitt and Ashley Montanaro. *Complexity Classification of Local Hamiltonian Problems*. Mar. 27, 2016. DOI: [10.48550/arXiv.1311.3161](https://arxiv.org/abs/1311.3161). arXiv: [1311.3161 \[quant-ph\]](https://arxiv.org/abs/1311.3161). URL: <http://arxiv.org/abs/1311.3161> (visited on 12/11/2023). Pre-published.
- [Cob+23] Nolan J. Coble et al. “Local Hamiltonians with No Low-Energy Stabilizer States”. en. In: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. DOI: [10.4230/LIPICS.TQC.2023.14](https://drops.dagstuhl.de/entities/document/10.4230/LIPICS.TQC.2023.14). URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPICS.TQC.2023.14>.
- [Coo71] Stephen A. Cook. “The complexity of theorem-proving procedures”. In: *Proceedings of the Third Annual ACM Symposium on Theory of Computing*. STOC ’71. Shaker Heights, Ohio, USA: Association for Computing Machinery, 1971, pp. 151–158. ISBN: 9781450374644. DOI: [10.1145/800157.805047](https://doi.org/10.1145/800157.805047). URL: <https://doi.org/10.1145/800157.805047>.
- [Din+24] Zhiyan Ding et al. *Polynomial-Time Preparation of Low-Temperature Gibbs States for 2D Toric Code*. 2024. arXiv: [2410.01206 \[quant-ph\]](https://arxiv.org/abs/2410.01206). URL: <https://arxiv.org/abs/2410.01206>.
- [Gha+15] Sevag Gharibian et al. “Quantum Hamiltonian Complexity”. In: *Foundations and Trends® in Theoretical Computer Science* 10.3 (Oct. 12, 2015), pp. 159–282. ISSN: 1551-305X, 1551-3068. DOI: [10.1561/0400000066](https://www.nowpublishers.com/article/Details/TCS-066). URL: <https://www.nowpublishers.com/article/Details/TCS-066> (visited on 10/25/2023).
- [GMV16] David Gosset, Jenish C. Mehta, and Thomas Vidick. *QCMA Hardness of Ground Space Connectivity for Commuting Hamiltonians*. arXiv.org. Oct. 12, 2016. DOI: [10.22331/q-2017-07-14-16](https://arxiv.org/abs/1610.03582v2). URL: <https://arxiv.org/abs/1610.03582v2> (visited on 01/11/2024).
- [GMV17] David Gosset, Jenish C Mehta, and Thomas Vidick. “QCMA hardness of ground space connectivity for commuting Hamiltonians”. In: *Quantum* 1 (2017), p. 16.

- [HJ24] Yeongwoo Hwang and Jiaqing Jiang. *Gibbs state preparation for commuting Hamiltonian: Mapping to classical Gibbs sampling*. 2024. arXiv: [2410.04909](https://arxiv.org/abs/2410.04909) [quant-ph]. URL: <https://arxiv.org/abs/2410.04909>.
- [IJ23] Sandy Irani and Jiaqing Jiang. *Commuting Local Hamiltonian Problem on 2D beyond Qubits*. Version 1. Sept. 9, 2023. DOI: [10.48550/arXiv.2309.04910](https://doi.org/10.48550/arXiv.2309.04910). arXiv: [2309.04910](https://arxiv.org/abs/2309.04910) [quant-ph]. URL: <http://arxiv.org/abs/2309.04910> (visited on 10/26/2023). Pre-published.
- [Jia23] Jiaqing Jiang. *Local Hamiltonian Problem with Succinct Ground State Is MA-Complete*. Sept. 18, 2023. DOI: [10.48550/arXiv.2309.10155](https://doi.org/10.48550/arXiv.2309.10155). arXiv: [2309.10155](https://arxiv.org/abs/2309.10155) [quant-ph]. URL: <http://arxiv.org/abs/2309.10155> (visited on 10/24/2023). Pre-published.
- [Jor75] Camille Jordan. “Essai sur la géométrie à n dimensions”. In: *Bulletin de la Société Mathématique de France* 3 (1875), pp. 103–174. DOI: [10.24033/bmsf.90](https://doi.org/10.24033/bmsf.90).
- [KB16] Michael J Kastoryano and Fernando GSL Brandao. “Quantum Gibbs samplers: The commuting case”. In: *Communications in Mathematical Physics* 344 (2016), pp. 915–957.
- [KKR05] Julia Kempe, Alexei Kitaev, and Oded Regev. *The Complexity of the Local Hamiltonian Problem*. Oct. 2, 2005. DOI: [10.48550/arXiv.quant-ph/0406180](https://doi.org/10.48550/arXiv.quant-ph/0406180). arXiv: [quant-ph/0406180](https://arxiv.org/abs/quant-ph/0406180). URL: <http://arxiv.org/abs/quant-ph/0406180> (visited on 04/09/2024). Pre-published.
- [KSV02] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalii. *Classical and Quantum Computation*. USA: American Mathematical Society, June 2002. 257 pp. ISBN: 978-0-8218-3229-5.
- [Lev73] Leonid A Levin. “Universal sequential search problems”. In: *Problems of information transmission* 9.3 (1973), pp. 265–266.
- [Sch11] Norbert Schuch. “Complexity of Commuting Hamiltonians on a Square Lattice of Qubits”. In: *Quantum Information & Computation* 11.11-12 (Nov. 1, 2011), pp. 901–912. ISSN: 1533-7146.

A Verifiers for commuting local Hamiltonians

Much of the work surrounding commuting local Hamiltonians has been directed towards showing ever increasing classes of them are classically verifiable. However, showing that commuting local Hamiltonians contain problems that are not in NP has proven much harder; to this day it is not even known if commuting local Hamiltonians can capture the complexity of problems in MA. In this section, we do not make any formal progress towards showing the hardness of commuting local Hamiltonians, but try to provide a conceptual framework that might be useful for studying commuting local Hamiltonians in the future. We provide a class of quantum verifiers for whom the problems that they can verify exactly captures the complexity of commuting local Hamiltonians, although the definition of the verifier is mostly specifically designed to capture commuting local Hamiltonians. Consider the following complexity class

Definition 24 (QIMA_k). *An k -local instant verifier V is a time-uniform polynomial-time algorithm on two registers AB that consists of $m = \text{poly}(n)$ many $k + 1$ -local B -controlled gates G_1, \dots, G_m that commute with each other. V acts on a witness $|\psi\rangle_{\text{A}}$, and an auxiliary register B initialized in the state $|+\rangle^{\otimes n+}$, and V measures $|+\rangle\langle+|^{\otimes n+} \otimes \text{id}$ for some l to accept.*

A promise problem L belongs to QIMA_k if there exists a k -local instant verifier such that for all $x \in L$, there exists a witness $|\psi_x\rangle$ such that $V_x(|\psi_x\rangle)$ accepts with probability 1, and for all $x \notin L$, for all $|\psi\rangle$, $V_x(|\psi\rangle)$ accepts with probability at most $1/2$.

Lemma A.1. *The k -local commuting local projectors problem is in QIMA_k .*

Proof. Consider the instant verifier that performs the following unitary

$$U_i = \text{id} - 2\Pi_i,$$

controlled on the i 'th $|+\rangle$ ancilla register. Then the verifier measures all m ancilla registers. If $|\psi\rangle$ is a joint 0-energy state of all verifiers, then for all i , $U_i|\psi\rangle = |\psi\rangle$, and otherwise for every state $|\psi\rangle$ there exists an i such that $\text{Re}(\langle\psi|U_i|\psi\rangle) < c$ for some c inverse-polynomial away from 1. Thus, if the commuting projectors

have a 0 energy ground state, the ground state is accepted by the verifier with probability 1, and otherwise every state is rejected by the verifier with probability at least inverse polynomial probability. Applying parallel repetition allows us to amplify the advantage without changing the locality of the unitaries. \square

Lemma A.2. *For all k , the k -local commuting local projectors problem is hard for QIMA_k .*

Proof. Let G_1, \dots, G_m be the gates of a k -local instant verifier V . Then consider the following commuting local Hamiltonian:

$$h_i = \text{id} - \frac{G_i + G_i^\dagger}{2}.$$

We note that every h_i commutes, and if $|\psi\rangle$ (the witness) is a $+1$ -eigenstate of G_i , it is a 0-energy eigenstate of h_i . Note that not all of the h_i are projectors, however using the reduction from [IJ23], the prover can send a (classical) NP witness containing λ_i . Finally note that every h_i is Hermitian and commutes with every other h_i if the G_i commute with each other. \square

We can define the complexity class QIMA_{\log} to be the complexity class defined in a similar fashion where k is allowed to be a function of the input size, and is $\log(n)$. Unfortunately very little is known about reductions between commuting local Hamiltonian problems, for example it is not known different localities are equally hard, or if reducing the rank of local projectors changes the difficulty of the problem. We think these are interesting open questions for future work, and we hope that by defining some complexity class that captures the difficulty of commuting local Hamiltonians, there will be more of a formal framework to study the hardness of these problems.

B Lemmas for commuting local Hamiltonians

Lemma B.1 (Rounding preserves commutation). *Suppose A and B are two commuting, $N \times N$ Hermitian matrices. Suppose B has eigendecomposition $B = \sum_i \lambda_i^A \Pi_i$, with $\lambda_i^A \neq \lambda_j^A$ for $i \neq j$. Then, the rounded matrix $\text{round}(B) = \sum_{i: \lambda_i^A > 0} \Pi_i$ commutes with A .*

Proof. Since A and B commute, A acts invariantly on each eigenspace Π_i of B and thus commutes with Π_i individual. This easily implies that A commutes with $\text{round}(B)$. \square

Lemma 2.1 (Equivalence under projectors, implicit in [IJ23]). *Suppose we have a CLH instance $H = \sum_{i \in [m]} h_i$ and two POVM measurements $\Pi_1 = \{\pi_1^1, \dots, \pi_{k_1}^1\}$ and $\Pi_2 = \{\pi_1^2, \dots, \pi_{k_2}^2\}$ and disjoint sets of Hamiltonian terms $S, T \subseteq [m]$ such that,*

- All terms h_i with $i \in [m] \setminus T$ commute with Π_1 .
- All terms h_i with $i \in [m] \setminus S$ commute with Π_2 .
- All remaining terms h_i with $i \in [m] \setminus (T \cup S)$ commute with both Π_1 and Π_2 .

Then,

$$\lambda_0(H) = 0 \iff \exists i \in [k_1], j \in [k_2] \text{ such that } \text{Tr} \left[\prod_{i \in S} (\pi_i^1 - \pi_i^1 h_i \pi_i^1) \prod_{i \in T} (\pi_j^2 - \pi_j^2 h_i \pi_j^2) \overline{H}_{rest} \right], \quad (5)$$

where $\overline{H}_{rest} \triangleq \prod_{i \in [m] \setminus (S \cup T)} (\text{id} - h_i)$. Note that the ordering of the terms on the RHS matters.

Proof. Since H is commuting,

$$\lambda_0(H) = 0 \iff \text{Tr} \left[\prod_i (\text{id} - h_i) \right] > 0.$$

This fact can be easily seen by considering a simultaneously diagonalizing basis for all h_i 's. For the RHS, we have

$$\mathrm{Tr} \left[\prod_i (\mathrm{id} - h_i) \right] = \sum_{i,j} \mathrm{Tr} \left[\pi_i^1 (\mathrm{id} - h_1) \pi_i^1 \pi_j^2 (\mathrm{id} - h_2) \pi_j^2 \overline{H}_{\mathrm{rest}} \right], \quad (38)$$

where we used that π_i^1 commutes with h_1 and thus $\pi_i^1 h_1 \pi_k^1 = 0$ for $i \neq k$ and similarly for h_2 . Additionally, we write $\overline{H}_{\mathrm{rest}} = \prod_{i>2} (\mathrm{id} - h_i)$. Suppose each summand on the RHS is non-negative. If the LHS is 0, this immediately means all terms on the right are 0. On the other hand, if the LHS is positive, then there must be some i, j such that

$$\mathrm{Tr} \left[\pi_i^1 (\mathrm{id} - h_1) \pi_i^1 \pi_j^2 (\mathrm{id} - h_2) \pi_j^2 \overline{H}_{\mathrm{rest}} \right] > 0.$$

This proves the claim. It remains to show each summand is positive.

Consider arbitrary i, j and write $\pi_1 = \pi_i^1$ and $\pi_2 = \pi_j^2$. Since π_2 commutes with h_2 , we know that these operators are diagonal in the same basis. In this basis, write $\pi_2 = \sum_i |i\rangle\langle i|$ and $h_2 = \sum_i \lambda_i |i\rangle\langle i|$. Thus,

$$\pi_2 - \pi_2 h_2 \pi_2 = \sum_i (1 - \lambda_i) |i\rangle\langle i|$$

But since $\|h_2\|_\infty \leq 1$, each coefficient is non-negative and this matrix is PSD. Next, by assumption h_i with $i > 2$ commute with both π_2 and h_2 ; thus, $[\pi_2 - \pi_2 h_2 \pi_2, \overline{H}_{\mathrm{rest}}] = 0$. The product of two commuting PSD operators is also PSD and thus $(\pi_2 - \pi_2 h_2 \pi_2) \overline{H}_{\mathrm{rest}} \succeq 0$. Finally, we have by the same argument that $\pi_1 - \pi_1 h_1 \pi_1 \succeq 0$. Thus

$$(\pi_1 - \pi_1 h_1 \pi_1) ((\pi_2 - \pi_2 h_2 \pi_2) \overline{H}_{\mathrm{rest}})$$

is the product of two PSD operators and this implies the trace is non-negative. \square