

Commuting Local Hamiltonians Beyond 2D

John Bostanci 

Columbia University, New York, NY, USA

Yeongwoo Hwang 

Harvard University, Cambridge, MA, USA

Abstract

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. However, unlike the general local Hamiltonian problem, the exact complexity of the commuting local Hamiltonian problem (CLH) remains unknown. A number of works have shown that increasingly expressive families of commuting local Hamiltonians admit 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 for pairs of projectors and the Structure Lemma for C^* algebras. 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. Using our rounding technique, 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. [4]

2012 ACM Subject Classification Theory of computation → Quantum computation theory

Keywords and phrases Quantum complexity, commuting Hamiltonians, complexity theory, C^* algebras

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.25

Related Version *Extended Version*: <https://arxiv.org/abs/2210.09155>

Funding 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.

John Bostanci: JB is supported by Henry Yuen’s AFORS (award FA9550-21-1-036) and NSF CAREER (award CCF2144219).

Yeongwoo Hwang: 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 authors(s) and do not necessarily reflect the views of the National Science Foundation.

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.



© John Bostanci and Yeongwoo Hwang;

licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 25; pp. 25:1–25:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The local Hamiltonian problem [19] is one of the most important problems in quantum complexity theory, and the central problem of the field 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 [5]. 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; for example many advances in Gibbs state sampling come from analyzing the Gibbs states of commuting local Hamiltonians [17, 6, 11, 14], and work towards the quantum PCP conjecture has largely focused on quantum LDPC codes, and related code Hamiltonians [5, 8].

As it shares properties of classical constraint satisfaction problems and the local Hamiltonian problem, 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 the commuting local Hamiltonian problem has been studied explicitly by a number of authors, starting with the work of [7], who showed 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 [2], on locally expanding graphs [1], 2D over qubits [3], and on a square lattice over qudits [15] 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 local, or **(b)** the 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 Hamiltonian term; namely that they are all rank-1 projectors.

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 [9] 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 originally used to reduce the locality of QMA-hard local Hamiltonians from 3 to 2 [18]. Since then, perturbation gadgets have been used extensively to show many other, often much more constrained, Hamiltonians are

also QMA-complete [10]. 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 [15] that restricting to projectors does not 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 SAT [9, 20] 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, [10] 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 [4] 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 [3], 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 [7] 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 [1] who extended the result of [7] 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. [21] considered the 2D square lattice and showed qubit Hamiltonians defined in the lattice can be placed in NP. A follow up work by [3] 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, [15] extended the ideas contained in [21] 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 [12], 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.

► **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 [15], 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 [3] 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 2** (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 3** (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; roughly, this corresponds to the Hamiltonian being defined on a polygonal complex, with terms placed on faces of the complex and qudits on edges.

New techniques for rounding commuting local Hamiltonians. Our guided reductions will be constructed using the building block of “rounding schemes” (Definition 11) in which a Hamiltonian H over a set of registers $R = R_1 \otimes \cdots \otimes R_n$ is “rounded” to a Hamiltonian



■ **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.

$\tilde{H}$ 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 [7, 15] can be viewed as implementing a rounding scheme). Our contribution, however, is to formalize this notion and introduce more general rounding schemes.

In [15], 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 4** (Rounding pairs of projectors (informal version of Theorem 14)). *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 non-zero blocks arising from the Jordan decomposition of the product of the two projectors. We can view this as an extension of the rounding scheme from [15] 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 5** (Characterization of commuting rank 1 projectors (informal version of Theorem 17)). *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 [3]. 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 of the full version of the paper, 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 [3] 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 [15] 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 [13] 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.

2 Technical overview

The first key technique used in this work is the Structure Lemma of [7], which has played a central role in nearly every work on the complexity of commuting Hamiltonians.


► **Lemma 6** (The Structure Lemma). *Let $\mathcal{A} \subseteq \mathcal{L}(\mathcal{R})$ be a C^* -algebra on $\mathcal{R}$. Then there exists a direct sum decomposition $\mathcal{R} = \bigoplus_i \mathcal{R}_i$ and a tensor product structure $\mathcal{R}_i = \mathcal{R}_i^1 \otimes \mathcal{R}_i^2$ such that*

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

Our second key technique is Jordan’s Lemma which has found widespread use throughout quantum complexity theory, yet, to our best knowledge, has not yet been explicitly applied to the study of CLHs.

► **Lemma 7** (Jordan’s Lemma [16]). *For any two Hermitian projectors Π_P and Π_Q on a register $\mathcal{R}$, there exists a orthogonal decomposition of $\mathcal{R} = \bigoplus_b \mathcal{R}_b$ (the Jordan decomposition with respect to Π_P, Π_Q) into one- and two-dimensional subspaces $\mathcal{B} = \{\mathcal{S}_b\}_b$ (Jordan subspaces, or blocks) where each $\mathcal{R}_b$ 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.*



■ **Figure 2** Puncturing holes via the Structure Lemma.

- 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 $|q_b^1\rangle$.

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 [7]. An easy consequence of this lemma is the following corollary,

► **Corollary 8** (Structure of two commuting operators). *Let $\mathcal{A}_h \subseteq \mathcal{L}(\mathcal{R})$ be the C^* -algebra induced by a Hermitian h and $\mathcal{A}_{h'}$ be the C^* -algebra on $\mathcal{R}$ induced by h' that commutes with h . Let $\mathcal{R}_i^j$ be the decomposition induced by Lemma 6 applied to $\mathcal{A}_h$. Then the following holds:*

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

$$\mathcal{A}_{h'} \subseteq \bigoplus_i \text{id}(\mathcal{R}_i^1) \otimes \mathcal{L}(\mathcal{R}_i^2). \quad (3)$$

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

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

2-local setting. In the 2-local setting of [7], every pair of terms h, h' either do not interact, or interact *only* on a single-qudit register $\mathcal{R}$. This means that for each register, a single decomposition $\mathcal{R} = \bigoplus_i \mathcal{R}_i$ can be constructed so that every operators acts invariantly on the corresponding projectors $\pi_i \triangleq (\pi_i)_{\mathcal{R}_i}$. Concretely, each term h is *block diagonal* with respect to the subspaces $\mathcal{R}_i$. 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:

$$\lambda_0(H) = 0 \iff \exists i \text{ s.t. } \text{Tr} \left[\prod_j (\pi_i - \pi_i h_j \pi_i) \right] > 0. \quad (4)$$

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 $(R_1, \dots, \pi_i R_i \pi_i, \dots, R_n)$, where $\pi_i R_i \pi_i \leq R_i$ is the sub-space generated by projecting vectors of 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 $R_i = \otimes_h R_i^{(h)}$ where each $\pi_i h \pi_i$ acts on non-trivially only on $R_i^{(h)}$. This effectively decouples all terms on 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 R , then commutation between adjacent terms is not entirely determined by the terms' structure on 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 (4), we get

$$\lambda_0(H) = 0 \iff \quad (5)$$

$$\exists i, j \text{ s.t. } \text{Tr} \left[\tilde{\pi}_j (\mathbb{I} - h_1) \tilde{\pi}_j \tilde{\pi}_j (\mathbb{I} - h_3) \tilde{\pi}_j \pi_i (\mathbb{I} - h_2) \pi_i \pi_i (\mathbb{I} - h_4) \pi_i \prod_{j>4} (\mathbb{I} - h_j) \right] > 0. \quad (6)$$

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


$$A \triangleq \tilde{\pi}_j (\mathbb{I} - h_1) \tilde{\pi}_j \tilde{\pi}_j (\mathbb{I} - h_3) \tilde{\pi}_j \quad (7)$$

$$B \triangleq \pi_i (\mathbb{I} - h_2) \pi_i \pi_i (\mathbb{I} - h_4) \pi_i \prod_{j>4} (\mathbb{I} - h_j) \quad (8)$$

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 9** (Equivalence under projectors, implicit in [15]). *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 .*



■ **Figure 3** A degree 4 register R .

Then,

$$\lambda_0(H) = 0 \iff \quad (9)$$

$$\exists i \in [k_1], j \in [k_2] \text{ such that } \text{Tr} \left[\prod_{i \in S} (\pi_i^1 - \pi_i^1 h_1 \pi_i^1) \prod_{i \in T} (\pi_j^2 - \pi_j^2 h_2 \pi_j^2) \overline{H}_{rest} \right], \quad (10)$$

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.

Equation (5) is recovered by taking $S = \{1, 2\}$ and $T = \{3, 4\}$. This lemma therefore yields a statement equivalent to Equation (4). 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 [21], 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:

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


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 R and thus this register (which [21] calls a *removable* qudit) can be traced out from the system. [21] 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, [15] 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, [15] 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 a improved rounding scheme with the idea of “triangulations” from [3].

Guided reduction via triangulation. The key insight from [3] 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



■ **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 [3] the center terms are simply removed and corrected later.

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. [3] 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, [3] 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 [3], 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 [3] 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 [7]. There a simple consequence of the Structure Lemma is the following lemma.

► **Lemma 10.** *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 10 to remove the obstructions.

In comparison to previous techniques,

1. The semi-separable self-reduction of [15] 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. [21] 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 [3] 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 4, we give the key technical tools we will need in this paper. In Section 4.1, 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.
- Sections 6–7 apply the rounding scheme to various geometrically local instances of the commuting local Hamiltonian problem.
 - In Section 6, we prove that $2D^*\text{-CLH}^{(1)}$ is in NP (Theorem 24)
 - In Section 7, we prove that $3D^*\text{-CLH}^{(1)}$ is in NP (Theorem 26).

3 Preliminaries

See full version of the paper for extended preliminaries. All proofs are also omitted in this version of the paper.


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 9 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:

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

► **Definition 11** (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 $\tilde{H}_\Pi$ such that H has a 0-energy ground state if and only if $\tilde{H}_\Pi$ does.



■ **Figure 5** Result of Corollary 12.

Although in the above definition, we write $\tilde{H}_\Pi$ 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, another straightforward rounding scheme arises when each operator acting on a register *locally* commutes when restricted to that register. Such a register is called *classical*.

► **Corollary 12.** 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. [7, 15]) 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 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 [15], the improved rounding scheme we supply applies in all the situations where the rounding scheme of [15] 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 14 works for general CLH instances.

► **Definition 13.** We say that a matrix P survives a projector π if $\pi P \pi \neq 0$.

► **Theorem 14 (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. See full version for details. ◀

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 15 (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

$$\begin{aligned} \Pi P \Pi &= P \text{ and} \\ \Pi Q \Pi &= 0. \end{aligned} \tag{11}$$

► **Definition 16 (Singular commutation for rank 1 projectors).** Two rank 1 projectors P_{AB} and Q_{BC} commute in a singular way if the following holds.

$$\begin{aligned} P &= |\psi\rangle\langle\psi|_B \otimes \tilde{P}_A \text{ and} \\ Q &= |\psi\rangle\langle\psi|_B \otimes \tilde{Q}_C. \end{aligned} \tag{12}$$

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 17 (Commutation of rank 1 projectors).** Let P and Q be two rank 1 projectors. Then P and Q either commute in a singular (Definition 16) or reducing way (Definition 15).

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

► **Definition 18** (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 19.** 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 6) 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 20** (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. See full version for details. ◀

Combining Lemma 20 with the Structure Lemma (Lemma 6), we get the following corollary about the way that rank 1 projectors commute with neighboring terms.

► **Corollary 21.** *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 6 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 20. 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. See full version for details. ◀

► **Remark 22** (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 17 in the degree-4, rank-1 case. Consider a qudit register $\mathbb{R}$, acted on non-trivially by P_1, P_2 and Q_1, Q_2 , where the pairs P_1 and P_2 , and Q_1 and Q_2 only interact on $\mathbb{R}$ (as in Figure 6). By Theorem 17, 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 14. 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, P_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 23** (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

► **Theorem 24.** *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 [7], we get the following simple corollary.

► **Corollary 25.** *The local Hamiltonian problem for 2D*-CLH⁽¹⁾ is in NP.*



■ **Figure 6** An original Hamiltonian, (a), and the possible resulting Hamiltonians after applying the puncturing and rounding scheme.

The proof of Theorem 24 is inspired by [3]. However, our final result is incomparable as we handle the non-starred version of the problem for any local dimension, but [3] is able to handle Hamiltonian terms of any rank. As in [3], 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 14. See the full version of the proof.


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 assumptions,

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

► **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 8) from the simpler case.*



(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 7** A local view at a cubulation.



■ **Figure 8** 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.

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 9a) can exit through a distinct face.


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

► **Theorem 26** (Guided reduction for $3D^*\text{-CLH}^{(1)}$). *Let H be an instance of $3D^*\text{-CLH}^{(1)}$ with degree g on d -dimensional registers. Suppose we also have a bound k on the locality of any term. Then,*


- *If $g \leq \frac{1}{e}d^k$, then by the quantum Lovász Local Lemma the instance H is trivially satisfiable.*
- *Otherwise, assume $g > \frac{1}{e}d^k \geq 6$. If H satisfies Assumption 1 and Assumption 2, then there is a guided reduction from H to a 2-local commuting local Hamiltonian.*

► **Corollary 27** (Local Hamiltonian problem on $3D^*\text{-CLH}^{(1)}$ in NP). *For non-trivial $3D^*\text{-CLH}^{(1)}$ instances meeting the requirements of Theorem 26, 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 26 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.



■ **Figure 9** The process of constructing tubes of slices for a vertex of the cubulation.



■ **Figure 10** 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.

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 10. 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 Section 6 to create “tunnels” so that the edges can pass through holes in the Hamiltonian. This still is not quite sufficient, as there can be blockages in the generated tunnels. 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. See the full version for the proof.

References

- 1 Dorit Aharonov and Lior Eldar. On the complexity of commuting local hamiltonians, and tight conditions for topological order in such systems. In Rafail Ostrovsky, editor, *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 334–343. IEEE Computer Society, 2011. doi:10.1109/FOCS.2011.58.
- 2 Dorit Aharonov and Lior Eldar. The commuting local hamiltonian problem on locally expanding graphs is approximable in NP. *Quantum Inf. Process.*, 14(1):83–101, 2015. doi:10.1007/S11128-014-0877-9.
- 3 Dorit Aharonov, Oded Kenneth, and Itamar Vigdorovich. On the complexity of two dimensional commuting local hamiltonians. In Stacey Jeffery, editor, *13th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2018, Sydney, Australia, July 16-18, 2018*, volume 111 of *LIPIcs*, pages 2:1–2:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.TQC.2018.2.
- 4 Andris Ambainis, Julia Kempe, and Or Sattath. A quantum lovász local lemma. *J. ACM*, 59(5):24:1–24:24, 2012. doi:10.1145/2371656.2371659.
- 5 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*, pages 1090–1096, 2023. doi:10.1145/3564246.3585114.
- 6 Ivan Bardet, Ángela Capel, Li Gao, Angelo Lucia, David Pérez-García, and Cambyse Rouzé. Rapid thermalization of spin chain commuting hamiltonians. *Physical Review Letters*, 130(6):060401, 2023.
- 7 Sergey Bravyi and Mikhail N. Vyalyi. Commutative version of the local hamiltonian problem and common eigenspace problem. *Quantum Inf. Comput.*, 5(3):187–215, 2005. doi:10.26421/QIC5.3-2.
- 8 Nolan J. Coble, Matthew Coudron, Jon Nelson, and Seyed Sajjad Nezhadi. Local hamiltonians with no low-energy stabilizer states. In Omar Fawzi and Michael Walter, editors, *18th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2023, Aveiro, Portugal, July 24-28, 2023*, volume 266 of *LIPIcs*, pages 14:1–14:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.TQC.2023.14.
- 9 Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing, STOC '71*, pages 151–158, New York, NY, USA, 1971. Association for Computing Machinery. doi:10.1145/800157.805047.
- 10 Toby S. Cubitt and Ashley Montanaro. Complexity classification of local hamiltonian problems. *SIAM J. Comput.*, 45(2):268–316, 2016. doi:10.1137/140998287.
- 11 Zhiyan Ding, Bowen Li, Lin Lin, and Ruizhe Zhang. Polynomial-time preparation of low-temperature gibbs states for 2d toric code, 2024. arXiv:2410.01206.
- 12 David Gosset, Jenish C. Mehta, and Thomas Vidick. QCMA hardness of ground space connectivity for commuting hamiltonians. *Quantum*, 1:16, 2017. doi:10.22331/Q-2017-07-14-16.
- 13 David Gosset, Jenish C Mehta, and Thomas Vidick. Qcma hardness of ground space connectivity for commuting hamiltonians. *Quantum*, 1:16, 2017. doi:10.22331/Q-2017-07-14-16.
- 14 Yeongwoo Hwang and Jiaqing Jiang. Gibbs state preparation for commuting hamiltonian: Mapping to classical gibbs sampling, 2024. doi:10.48550/arXiv.2410.04909.
- 15 Sandy Irani and Jiaqing Jiang. Commuting local hamiltonian problem on 2d beyond qubits. *Communications in Mathematical Physics*, 406(11):273, 2025.
- 16 Camille Jordan. Essai sur la géométrie à n dimensions. *Bulletin de la Société Mathématique de France*, 3:103–174, 1875. doi:10.24033/bmsf.90.
- 17 Michael J Kastoryano and Fernando GSL Brandao. Quantum gibbs samplers: The commuting case. *Communications in Mathematical Physics*, 344:915–957, 2016.
- 18 Julia Kempe, Alexei Y. Kitaev, and Oded Regev. The complexity of the local hamiltonian problem. *SIAM J. Comput.*, 35(5):1070–1097, 2006. doi:10.1137/S0097539704445226.

- 19 Alexei Y. Kitaev, A. H. Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate studies in mathematics*. American Mathematical Society, 2002. URL: <https://bookstore.ams.org/gsm-47/>.
- 20 Leonid A Levin. Universal sequential search problems. *Problems of information transmission*, 9(3):265–266, 1973.
- 21 Norbert Schuch. Complexity of commuting hamiltonians on a square lattice of qubits. *Quantum Inf. Comput.*, 11(11-12):901–912, 2011. doi:10.26421/QIC11.11-12-1.