Complexity Classification of Product State Problems for Local Hamiltonians

John Kallaugher
Sandia National Laboratories, Albuquerque, NM, USA

Ojas Parekh
Sandia National Laboratories, Albuquerque, NM, USA

Kevin Thompson
Sandia National Laboratories, Albuquerque, NM, USA

Yipu Wang
Sandia National Laboratories, Albuquerque, NM, USA

Justin Yirka
Sandia National Laboratories, Albuquerque, NM, USA
The University of Texas at Austin, TX, USA

Abstract

Product states, unentangled tensor products of single qubits, are a ubiquitous ansatz in quantum computation, including for state-of-the-art Hamiltonian approximation algorithms. A natural question is whether we should expect to efficiently solve product state problems on any interesting families of Hamiltonians.

We completely classify the complexity of finding minimum-energy product states for Hamiltonians defined by any fixed set of allowed 2-qubit interactions. Our results follow a line of work classifying the complexity of solving Hamiltonian problems and classical constraint satisfaction problems based on the allowed constraints. We prove that estimating the minimum energy of a product state is in P if and only if all allowed interactions are 1-local, and NP-complete otherwise. Equivalently, any family of non-trivial two-body interactions generates Hamiltonians with NP-complete product-state problems. Our hardness constructions only require coupling strengths of constant magnitude.

A crucial component of our proofs is a collection of hardness results for a new variant of the Vector Max-Cut problem, which should be of independent interest. Our definition involves sums of distances rather than squared distances and allows linear stretches.

We similarly give a proof that the original Vector Max-Cut problem is NP-complete in 3 dimensions. This implies hardness of optimizing product states for Quantum Max-Cut (the quantum Heisenberg model) is NP-complete, even when every term is guaranteed to have positive unit weight.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases quantum complexity, quantum algorithms, local hamiltonians

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.63

Related Version *Revised Version*: https://arxiv.org/abs/2401.06725v2

Funding This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Accelerated Research in Quantum Computing, Fundamental Algorithmic Research for Quantum Computing. O.P. was also supported by U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator. JY was partially supported by Scott Aaronson's Simons Investigator Award.

Acknowledgements Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. This written work is authored by an employee of

© John Kallaugher, Ojas Parekh, Kevin Thompson, Yipu Wang, and Justin Yirka;
licensed under Creative Commons License CC-BY 4.0

NTESS. The employee, not NTESS, owns the right, title and interest in and to the written work and is responsible for its contents. Any subjective views or opinions that might be expressed in the written work do not necessarily represent the views of the U.S. Government. The publisher acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this written work or allow others to do so, for U.S. Government purposes. The DOE will provide public access to results of federally sponsored research in accordance with the DOE Public Access Plan.

1 Introduction

Product states, unentangled tensor products of single-qubit states, have served as an effective focus for better understanding quantum phenomena. Because general quantum states cannot be described efficiently, approximation algorithms must be restricted to output some subset of states, an ansatz. Mean-field approaches are common as first steps in statistical mechanics, and recent approximation algorithms for extremal energy states of local Hamiltonians have relied on proving that product states provide good approximations in particular regimes [17, 4, 5, 18, 31, 32]. In fact, because in some natural regimes the ground states are rigorously well-approximated by product states [4], optimal approximation algorithms for local Hamiltonians on arbitrary interaction graphs must be capable of finding good product states. Understanding product state optimization is essential for understanding the complexity of Hamiltonian approximation generally.

Product states are a natural intermediate between classical and quantum states, allowing for superposition but not entanglement. Unlike general quantum states, they have succinct classical descriptions: a single-qubit pure state can be specified by two complex numbers, and an n-qubit pure product state can be specified by $2n$ complex numbers. One could consider "more quantum" intermediates, in the form of reduced states of two or more qubits. However, verifying the consistency of a set of quantum marginals is a QMA-complete problem, even for 2-qubit reduced states [27, 10]. Therefore, product states are uniquely useful when optimizing directly over state vectors.

We study the following question: for a family of Hamiltonians defined by a given set of allowed interactions, what is the complexity of computing the extremal energy over product states? Additionally, how does the complexity of optimizing over product states relate to that of optimizing over general states? For example, for a QMA-hard local Hamiltonian, must finding the optimal product state in turn be NP-hard?[1]

This question follows a long line of work classifying the complexity of constraint satisfaction problems (CSPs) based on the sets of allowed constraints, clauses, or interactions between variables. In particular, the dichotomy theorem of Schaefer [35] showed that for any set of allowed Boolean constraints, the family of CSPs is either efficiently decidable or is NP-complete. In the context of quantum problems, Cubitt and Montanaro [14] introduced a similar classification of ground state energy problems for 2-local Hamiltonians, showing that for any fixed set of allowed 2-qubit interactions, the 2-LH problem is either in P or NP-, StoqMA-, or QMA-complete (the StoqMA case relies on the concurrent work of Bravyi and Hastings [6]). We briefly survey some of this line of classical and quantum work in Related Work below.

[1] The product state problem is always in NP since product states have succinct classical descriptions with which we can compute the expected energy contribution from each local Hamiltonian term in time polynomial in the size of the term.

While the complexity of finding the extremal states (e.g. the ground state) of 2-local Hamiltonians is well understood, the complexity of finding optimal product state solutions has been only sparsely studied [22]. The only NP-hardness results for such problems are based on mapping a classical problem to a diagonal 0-1 valued Hamiltonian [39].

An additional motivation for our study is the hope of developing new methods for identifying families of local Hamiltonians for which problems involving *general* ground states are *not* hard. While a complete complexity classification for the general 2-LH problem is known, more refined attempts at classification which take into account restrictions on the sign of the weights or geometry of the system are currently incomplete [33]. Developing algorithms for product states is a "mostly classical" problem that is easier to analyze, and progress involving product states may inform our expectations regarding general states.

In this work, we completely classify the complexity of finding optimal product states for families of 2-local Hamiltonians. In fact, we find the complexity of the product state problem is fully determined by the complexity of the general local Hamiltonian problem: if the general problem is NP-hard, the product state problem is NP-complete, and otherwise it is in P. To arrive at our results, we study a variant of the Vector Max-Cut which should be of independent interest especially to the optimization community. As a corollary to our classification theorem, we give the first published proof that estimating optimal product state energies in the Quantum Max-Cut model is NP-complete, and we show hardness holds even for unweighted Hamiltonians.[2]

1.1 Our Contributions

Formal definitions are given in Section 2. A k-local Hamiltonian is a sum of Hamiltonian terms each of which only acts non-trivially on at most k qubits, analogous to k-variable Boolean clauses. k-LH denotes the problem of estimating the ground state energy (the minimum eigenvalue across all states) of a k-local Hamiltonian to inverse-polynomial additive precision. Given a set of local terms $\mathcal{S}$, $\mathcal{S}$-LH is k-LH restricted to Hamiltonians such that every term belongs to $\mathcal{S}$. Finally, prodLH and $\mathcal{S}$-prodLH are the restrictions of these problems to product states, i.e. to minimize $\langle\phi|H|\phi\rangle$ where H is the Hamiltonian and $|\phi\rangle$ ranges over tensor products of single-qubit states.[3]

2-local $\mathcal{S}$-prodLH

The classification of the general ground state energy problem by Cubitt and Montanaro [14] completely classifies $\mathcal{S}$-LH for any fixed set $\mathcal{S}$ of 2-qubit terms, showing it is either in P or it is one of NP-, StoqMA-, or QMA-complete.

In the same vein, we give a complete classification of product state complexity for families of 2-local Hamiltonians as a function of the set of allowed 2-qubit interactions. For any given set $\mathcal{S}$ of 2-qubit terms, we prove the problem $\mathcal{S}$-prodLH is either in P or is NP-complete. To the best of our knowledge, ours is the first systematic inquiry into the complexity of product state problems.

▶ **Theorem 1.** *For any fixed set of 2-qubit Hamiltonian terms $\mathcal{S}$, if every matrix in $\mathcal{S}$ is 1-local then $\mathcal{S}$-prodLH is in P, and otherwise $\mathcal{S}$-prodLH is NP-complete.*

[2] A more complex unpublished proof based on large graph cycles was known earlier by Wright [40].
[3] An earlier version of this work considered *exact* versions of product state and graph problems. We have now improved our hardness results to hold up to an inverse-polynomial additive gap.

Additionally, our hardness constructions only require coupling strengths (weights) of at most constant magnitude. This is preferable in practice and contrasts with most known QMA-hardness constructions.

The sets for which [14] shows $\mathcal{S}$-LH is in P are the same for which we show $\mathcal{S}$-PRODLH is in P, i.e. those containing only 1-local terms. This immediately implies that a family of 2-local Hamiltonians has efficiently-computable minimum product state energy if and only if it has efficiently-computable ground state energy.

▶ **Corollary 2.** *For any fixed set of 2-qubit Hamiltonian terms $\mathcal{S}$, the problem $\mathcal{S}$-LH is in* P *if and only if $\mathcal{S}$-PRODLH is in* P*.*

▶ **Corollary 3.** *For any fixed set of 2-qubit Hamiltonian terms $\mathcal{S}$, the problem $\mathcal{S}$-LH is* NP*-hard if and only if $\mathcal{S}$-PRODLH is* NP*-complete.*

Our results imply that hardness of product state approximations is not restricted to Hamiltonians for which product states well approximate ground states: for any QMA-hard family of terms, our result implies that (assuming QMA $\neq$ NP) we can construct a family of local Hamiltonians that are NP-hard to product state approximate *and* for which the product states do not well approximate the ground state by constructing Hamiltonians on two systems, each with one of these properties, and taking a disjoint union thereof. This implies that algorithms using product states to approximate the ground states of QMA-hard Hamiltonians face a "double penalty": hardness of approximating product states which themselves imperfectly approximate ground states.

The Stretched Linear Vector Max-Cut problem

Our hardness constructions for product state problems embed an objective function which we prove is NP-complete. This objective function generalizes the classical MAX-CUT problem. Given work on other variations of MAX-CUT, we expect this problem and our reductions should be of independent interest, especially to the optimization and approximation communities.

In MAX-CUT, one is given a graph $G = (V, E)$ and asked to assign to each vertex v a label $\hat{v} = \pm 1$ so as to achieve the maximum number of oppositely labeled adjacent vertices:

$$\mathrm{MC}(G) = \frac{1}{2} \max_{\hat{\imath}=\pm 1} \sum_{ij \in E} (1 - \hat{\imath}\hat{\jmath}) = \frac{1}{2} \max_{\hat{\imath}=\pm 1} \sum_{ij \in E} |\hat{\imath} - \hat{\jmath}|. \tag{1}$$

A problem referred to as VECTOR MAX-CUT, RANK-k-MAX-CUT, or MAX-CUT$_k$ (MC$_k$) has been studied [8, 7, 22] which generalizes MAX-CUT to assigning k-dimensional unit vectors so as to maximize the angles between adjacent vertex labels, or equivalently to maximize the squared distances between adjacent vertex labels:

$$\mathrm{MC}_k(G) = \frac{1}{2} \max_{\hat{\imath} \in S^{k-1}} \sum_{ij \in E} (1 - \hat{\imath} \cdot \hat{\jmath}) = \frac{1}{4} \max_{\hat{\imath} \in S^{k-1}} \sum_{ij \in E} \|\hat{\imath} - \hat{\jmath}\|^2. \tag{2}$$

Our new problem can be seen as a stretched and linear version of MC$_k$. The goal in W-LINEAR-MAX-CUT (MC$^{\mathrm{L}}_W$) is to assign unit vectors so as to maximize the distance between adjacent labels:

$$\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}(G) = \frac{1}{2} \max_{\hat{\imath} \in S^{d-1}} \sum_{ij \in E} \|W\hat{\imath} - W\hat{\jmath}\| = \frac{1}{2} \max_{\hat{\imath} \in S^{d-1}} \sum_{ij \in E} \sqrt{\|W\hat{\imath}\| + \|W\hat{\jmath}\| - 2(W\hat{\imath})^\top (W\hat{\jmath})}, \tag{3}$$

where W is a fixed $d \times d$ diagonal matrix. Comparing MC_k and $\mathrm{MC}^{\mathrm{L}}_W$, our problem sums over un-squared distances and incorporates a linear stretch given by W. We consider the decision version of this problem, in which the objective is to test whether the optimal solution is at least b or no more than a, for $b - a \geq 1/\operatorname{poly}(n)$. Note that, unlike $\mathcal{S}$-LH, this is an unweighted problem – one could naturally define a weighted version but our hardness results will not require this.

Geometrically, MC_k corresponds to embedding a graph into the surface of a unit sphere with the objective of maximizing the sum of the squared lengths of every edge. Likewise, our problem $\mathrm{MC}^{\mathrm{L}}_W$ corresponds to embedding a graph into the surface of a d-dimensional ellipsoid, with radii defined by the entries of W, with the objective of maximizing the sum of the (non-squared) edge lengths.

Despite being generalizations of the NP-complete Max-Cut problem, hardness of neither MC_k nor $\mathrm{MC}^{\mathrm{L}}_W$ is trivial. The Goemans-Williamson approximation algorithm for Max-Cut on an n-vertex graph begins with efficiently computing the solution to MC_n via an SDP. In fact, deciding MC_k is known to be in P for any $k = \Omega(\sqrt{|V|})$, [28, Theorem 8.4] or [3, (2.2)]. And while it has been conjectured by Lovász that MC_k is NP-complete for all constants $k \geq 1$, in [29, p. 236] and earlier, no proof has been given for any $k > 1$.

Our main theorem concerning W-linear-Max-Cut, which is used to prove Theorem 1, is the following.

▶ **Theorem 4.** *For any fixed non-negative $W = \operatorname{diag}(\alpha, \beta, \gamma)$ with at least one of α, β, γ nonzero, W-linear-Max-Cut is NP-complete.*

Quantum Max-Cut Product States and MC_3

As a corollary of our classification theorem, we give the first published proof of the fact that product state optimization in the Quantum Max-Cut (QMC) model is NP-hard. This model, also known as the anti-ferromagnetic Heisenberg model, is equivalent to $\mathcal{S}$-LH with $\mathcal{S} = \{XX + YY + ZZ\}$. We note that a sketch of a different proof for this specific problem was previously known but unpublished [40]. That proof was based on large graph cycles, and our gadgets are simpler to analyze.

However, the proof of Theorem 1 utilizes Hamiltonian gadgets involving negative weights (unlike the aforementioned proof of [40]). This leaves open whether prodQMC remains NP-hard on unweighted graphs. In Section 5, we give a direct proof of hardness using the fact that the unweighted product-state version of QMC is equivalent to MC_3 (Equation (2)). Our work then is also the first published proof that MC_k is NP-complete for some $k > 1$ (in our case, $k = 3$), partially resolving a conjecture of Lovász [29, p. 236]. Note that as with $\mathrm{MC}^{\mathrm{L}}_W$, we consider the decision version in which the goal is to determine whether the value is above b or below a, for b and a with inverse-poly separation.

▶ **Theorem 5.** MC_3 *is NP-complete.*

▶ **Corollary 6.** Quantum Max-Cut *restricted to product states,* prodQMC*, is NP-complete, even when all terms are restricted to have positive unit weight.*

1.2 Proof Overview

2-local $\mathcal{S}$-prodLH

As product states have classical descriptions and their energies can be calculated in polynomial time, $\mathcal{S}$-prodLH is automatically in NP, so we focus on how we show hardness. Our approach is in two parts. We show how to reduce $\mathrm{MC}^{\mathrm{L}}_W$ to $\mathcal{S}$-prodLH, and later we show that $\mathrm{MC}^{\mathrm{L}}_W$

NP-complete (Theorem 4). More precisely, the first part of our approach is to show that for any $\mathcal{S}$ containing a strictly 2-local term, there exists a corresponding weight matrix W meeting the conditions of Theorem 4 so that $\mathrm{MC}^{\mathrm{L}}_W$ is NP-hard. For any instance of $\mathrm{MC}^{\mathrm{L}}_W$ with this fixed W, we show how to construct a Hamiltonian from $\mathcal{S}$ such that the minimum product state energy encodes the $\mathrm{MC}^{\mathrm{L}}_W$ value, yielding Theorem 1.

We interpret problems over product states as optimization problems on the collection of Bloch vectors for each single-qubit state. For example, consider $\mathcal{S}$-PRODLH with the specific set $\mathcal{S} = \{XX + YY + ZZ\}$ (the QMC model). In this case, by writing each qubit v in the Bloch vector representation

$$|\phi_v\rangle\langle\phi_v| = \frac{1}{2}\left(I + v_1 X + v_2 Y + v_3 Z\right),$$

the energy contributed by an interaction between qubits u and v is

$$\mathrm{tr}\left(\left(X^u X^v + Y^u Y^v + Z^u Z^v\right)|\phi_u\rangle\langle\phi_u|\,|\phi_v\rangle\langle\phi_v|\right) = u_1 v_1 + u_2 v_2 + u_3 v_3. \tag{4}$$

So, given a Hamiltonian which is the sum of XYZ interactions between pairs of qubits, the problem of estimating the extremal product state energies is equivalent to optimizing the objective function

$$\sum_{uv} w_{uv}\, u \cdot v$$

over 3-dimensional unit vectors, where each edge uv corresponds to the Hamiltonian's (weighted) interaction graph. Up to constant shifts and scaling, this is equivalent to MC_3, introduced in Equation (2).

More generally, because the Pauli matrices are a basis for all Hermitian matrices, any 2-qubit interaction can be written as $H = \sum_{i,j=1}^{3} M_{ij}\sigma_i\sigma_j + \sum_{k=1}^{3}(*)c_k\sigma_k I + w_k I\sigma_k$, where the σ_i are the Pauli matrices, for some M, c, w. Then, the energy of a given product state is calculated similarly to Equation (4). However, the resulting expression potentially contains many terms.

Our approach is to take an arbitrary 2-qubit term and insert it into gadgets which simplify the energy calculations. First, in the proof of Theorem 1, we borrow a trick of [14] and symmetrize the terms. For any 2-qubit interaction H^{ab} on qubits a, b, the combined interaction $H^{ab} + H^{ba}$ is symmetric, invariant under swapping the qubits. A similar trick handles the case of anti-symmetric terms.

Second, in the proofs of Lemmas 14 and 15, we show how to use a symmetric or anti-symmetric term to embed the $\mathrm{MC}^{\mathrm{L}}_W$ value into the minimum energy of a gadget. We begin by removing the 1-local terms, such as $\sigma_1 I$ or $I\sigma_2$, again taking inspiration from gadgets used by [14]. For two qubits u, v corresponding to two vertices in a $\mathrm{MC}^{\mathrm{L}}_W$ instance, the gadget adds two ancilla qubits and weights each interaction within the gadget to effectively cancel out the 1-local terms. When the two ancilla qubits vary freely, we find the minimum energy contributed by the entire four-qubit gadget is determined by the distance between the states of u and v. Although each individual edge contributes energy proportional to the squared distance between their states, the overall gadget contributes energy proportional to just the distance of the two "vertex qubits", $-\|Mu - Mv\|$. With some massaging, we can treat M as a non-negative diagonal matrix which meets the conditions of Theorem 4.

Therefore, as desired, we have that the $\mathrm{MC}^{\mathrm{L}}_W$ value for an NP-complete instance of $\mathrm{MC}^{\mathrm{L}}_W$ can be embedded into the minimum product state energy of an $\mathcal{S}$-PRODLH instance.

To prove that $\mathrm{MC}^{\mathrm{L}}_W$ is NP-hard for any fixed diagonal non-negative W with at least one nonzero entry, we divide into three cases, in Lemmas 21–23. Our first proof is a reduction from the standard Max-Cut problem, while the other two are reductions from 3-Coloring.

When there is a unique largest entry of W, we reduce from Max-Cut by taking any input graph G and forming G' by adding large star gadgets around each of the vertices of G, each using many ancilla vertices. Because the ancilla vertices in each gadget have just one neighbor (the original vertex at the center of the gadget), their optimal vector labels given any choice of labels for the center vertices are the negation of the center vertex labels. This means they heavily penalize assigning the center vertices any labels that are not along the highest-weight axis. Therefore, when the maximum entry of W is unique, the optimal $\mathrm{MC}^{\mathrm{L}}_W$ assignment to G' will have its vector labels almost entirely along the highest-weight axis. The assignment can trivially earn the maximum possible value on the star gadgets, and the amount additional amount it can earn on the original edges of G corresponds to the Max-Cut value of G.

When all of the entries of W are equal, we reduce from 3-Coloring. Given a graph G, we construct G' by replacing each edge with a 4-clique gadget, made by adding one ancilla vertex per gadget, along with a single ancilla vertex shared by every gadget. We show that G is 3-colorable iff G' has a sufficiently large $\mathrm{MC}^{\mathrm{L}}_W$ value. Specifically, we show this holds iff there is a vector assignment that simultaneously achieves (nearly) the maximum value on all of the clique gadgets. Achieving the maximum objective value on a clique corresponds to maximizing the total distance between each pair of vectors, and this enforces a predictable arrangement.

When weights are equal, assigning these vectors can be viewed as inscribing vectors in the unit sphere, and it is known that maximal perimeter polyhedra inscribed in the sphere must be regular. So for a 4-clique, the vector labels must form a regular tetrahedron. We carefully argue that for regular tetrahedra, fixing two of the vertices (approximately) fixes the other two vertices up to swapping (several of these geometric facts are proved in Section A). This means that the clique gadget corresponding to an edge uv in G shares two vertices with any gadget corresponding to an edge vw incident to the first edge: v and the "global" ancilla shared by all gadgets. This means that, once we fix the vector assigned to the global ancilla, the choice of a vector for v restricts the labels of both u and w to be chosen from a set of two vectors. So simultaneously optimizing every clique gadget is possible iff, for each connected component of G', we can 3-color that component with three vectors (corresponding to, for any v in the component, the vector assigned to v, the vector assigned to the global ancilla vertex, and the two vectors that can share a maximal tetrahedron with those two).

Finally, when the two largest entries of W are equal but distinct from the third, we combine the two previous approaches. Inserting star gadgets effectively reduces the problem from three dimensions to two, by penalizing vector assignments not in the 2d space corresponding to the two largest entries of W. We then add 3-clique gadgets, with one ancilla for each edge in G, and optimizing these over two dimensions corresponds to inscribing maximal perimeter triangles in the unit circle. Now, assigning a vector to one vertex fixes the optimal vectors assigned to the other two vertices (again up to swapping them) and so there is again a one-to-one correspondence between vector assignments simultaneously optimizing every clique gadget and 3-colorings of the connected components of G.

Three-Dimensional Vector Max-Cut

Our proof that MC_k with positive unit weights is NP-complete runs on very similar lines to our 3-Coloring reduction for $\mathrm{MC}^{\mathrm{L}}_W$ with all weights equal. However, there is one additional complication: that proof depended on the fact that the sum of the side lengths of the tetrahedron is uniquely maximized by choosing it to be regular. This is not the case when we instead consider the squared side lengths, for instance assigning half the vectors to one pole of the sphere and half to the other would achieve the same bound. So we use a different gadget: instead of replacing every edge with a 4-clique, we replace every edge with a 4-clique that in turn has its edges replaced with triangles. It turns out that this gadget *does* have a unique optimal MC_k assignment, which in particular assigns the vectors in the 4-clique to a regular tetrahedron, allowing us to proceed along the same lines as the aforementioned proof.

1.3 Related Work

Brandão and Harrow [4] give simple conditions under which 2-local Hamiltonians have product states achieving near-optimal energy, such as systems with high-degree interaction graphs. This suggests that unless NP = QMA, such Hamiltonians cannot be QMA-complete. Since then, there has been a line of work on the relationship between product states and general ground states in other Hamiltonians, in the more general case when the two problems are not equivalent. See e.g. [5, 17, 20].

Product states have especially been studied in the context of the Quantum Max-Cut problem, introduced by Gharibian and Parekh [18]. Briët, de Oliveira Filho, and Vallentin [8] give hardness results conditional on the unique games conjecture for approximating the optimal product state in the QMC model. Related work by Hwang, Neeman, Parekh, Thompson, and Wright [22] also gives tight hardness results for the QMC problem under a plausible conjecture. Parekh and Thompson [32] give an optimal approximation algorithm for QMC when using product states.

See [22] for an exposition of the relationship between QMC restricted to product state solutions and the Vector Max-Cut problem. Studying vector solutions to Max-Cut has a long history [29], including the seminal Goemans-Williamson algorithm [19]. This study is usually with the goal of a solution to the original Max-Cut problem, which relates to approximation ratios of integer and semidefinite programs. Bounding these ratios has been referred to in terms of Grothendieck problems and inequalities: see [7, 2] for further context on this nomenclature. A tight NP-hardness result is known for the non-commutative Grothendieck problem [9] which also generalizes the "little" Grothendieck problem over orthogonal groups [2]. Iterative algorithms (heuristics) for solving MC_k are also well-studied in the literature (see [11] and citing references), since in practice solving MC_k is often faster than solving the corresponding SDP relaxation.

Cubitt and Montanaro [14] classified $\mathcal{S}$-LH for sets $\mathcal{S}$ of 2-qubit terms. This work relies in turn on the work of Bravyi and Hastings [6] to classify the StoqMA case. [14] also examined a variant where $\mathcal{S}$ is assumed to always contain arbitrary single-qubit terms. Follow-up work by Piddock and Montanaro [33] began investigating classifying the complexity of $\mathcal{S}$-LH under the additional restrictions of positive weights (anti-ferromagnetic model) and/or interactions restricted to a spatial geometry such as a 2D grid. [15, 34] continued along these lines, and introduced Hamiltonian simulations rather than computational reductions.

In classical computer science, Schaefer [35] gave a dichotomy theorem showing that given a fixed set of allowed constraints, the family of CSPs is either decidable in P or is NP-complete; but see Section 2 or [14] for some important assumptions that are made in the

quantum versus classical models. In fact, Schaefer's classification offers a more fine-grained classification with classes within P. Later, [12, 25] gave a similar dichotomy theorem for the complexity of Max-SAT and related optimization problems, where the question is not just whether all clauses are simultaneously satisfiable but how many are simultaneously satisfiable. Applying weights to constraints becomes relevant with Max-SAT, and is covered in their work. This work is especially relevant given k-LH is more analogous to Max-SAT than to SAT. Continuing, Jonsson [23] classified these problems when both positive and negative weights are allowed. While arbitrary weights seem natural in the quantum setting, previous classical work simply assumed all weights were non-negative. The book [13] offers an excellent survey of this area. The more recent results of [24, 36] extend classical classification results to problems with non-binary variables (analogous to qu*d*its).

1.4 Open Questions

1. We have shown a relationship between when product state problems and general Hamiltonian problems are hard. This points toward an important question: can some "interesting" class of local Hamiltonians or a Hamiltonian problem for which we do not know an explicit efficient algorithm be proven *not* hard, e.g. neither NP- nor QMA-hard, by showing the corresponding product state problem is in P?
2. Is a more refined classification of the complexity of product state problems, taking into account allowable weights or spatial geometry in the vein of [33], or imposing other promises, possible?
3. While little progress has been made classifying the general $\mathcal{S}$-LH problem for higher-locality families, can we classify $\mathcal{S}$-prodLH for families of k-local Hamiltonians with $k > 2$?
4. Can we relate approximability instead of just complexity? For example, does the ability to α-approximate the product state problem imply the ability to β-approximate the general ground state problem on families defined by some sets of allowed interactions but not others?
5. As mentioned above, we make the first progress towards a conjecture of Lovász [29] that MC_k is NP-hard for any $k = O(1)$. We only focus on $k = 3$ here because of our interest in prodLH. Can our proof be generalized to other values of k?

2 Preliminaries

We assume familiarity with the conventions of quantum computation [38] and complexity theory [1, 26]. See also [16, 14] for surveys of Hamiltonian complexity.

2.1 Notation

I denotes the identity operator. $\lambda_{\min}(H)$ and $\lambda_{\max}(H)$ denote the minimum and maximum eigenvalues of an operator H. In the same manner as with asymptotic $\mathrm{O}(\cdot)$ notation, we use $\mathrm{poly}(n)$ to denote a term that can be bounded by some fixed polynomial in n.

For an operator A, we use superscripts such as A^{abc} to indicate A acts on individual qubits a, b, and c. Unless A is symmetric, the order matters and A^{ab} is different than A^{ba}. If no superscripts are used, then the action is implicit in the ordering of the terms (left versus right).

When clear from context, we will denote the tensor product of two operators $A \otimes B$ simply by AB. All terms implicitly are tensor the identity on any systems not specified.

SWAP will denote the 2-qubit operator exchanging $|01\rangle$ and $|10\rangle$ while $|00\rangle$ and $|11\rangle$ unchanged. We call a 2-qubit term H *symmetric* if $H = \mathrm{SWAP}(H)\,\mathrm{SWAP}$, meaning the ordering of the qubits does not matter. Alternatively, H is *antisymmetric* if $H = -\mathrm{SWAP}(H)\,\mathrm{SWAP}$.

The single-qubit Pauli matrices are denoted X, Y, Z or $\sigma_1, \sigma_2, \sigma_3$. Recall that $\{X, Y, Z, I\}$ is a basis for 2×2 Hermitian matrices. The *Pauli decomposition* of a 2-qubit Hermitian matrix H is H written in the Pauli basis,

$$H = \sum_{i,j=1}^{3} M_{ij}\sigma_i\sigma_j + \sum_{k=1}^{3}(v_k\sigma_k I + w_k I\sigma_k), \tag{5}$$

with all coefficients real and the 3×3 matrix M referred to as the *correlation matrix*. Generally, Equation (5) should include a term wII, but we will work with traceless terms such that $w = 0$.

Unless otherwise stated, all graphs are undirected and simple, meaning there are no self-loops and no multi-edges. We assume all graphs are connected, as it is straightforward to extend any of our constructions to disconnected graphs. When summing over edges, $\sum_{ij\in E}$, we do not double-count ij and ji. Finally, $S^i = \{x \in \mathbb{R}^{i+1} : \|x\| = 1\}$ denotes the unit sphere in $(i+1)$-dimensional space.

2.2 Definitions and Assumptions

A k-local Hamiltonian on n qubits is a Hermitian matrix $H \in \mathbb{R}^{2^n \times 2^n}$ that can be written as $H = \sum_{i=1}^{m} H_i$ such that each H_i is Hermitian and acts non-trivially on at most k qubits. More precisely, each H_i acts on some subset S_i of at most k qubits and each term in the sum is $H_i \otimes I^{[n]\setminus S_i}$, but we generally leave this implicit. We usually consider constant values of k, so each term is of constant size independent of n. The k-qubit terms H_i are often referred to as *interactions* between qubits. We may refer to eigenvalues and expectation values $\langle\psi| H |\psi\rangle$ as the *energy* of the state $|\psi\rangle$ in the system described by H. In particular, the *ground state energy* and *ground state* refer to the minimum eigenvalue and an associated eigenvector.

Estimating the minimum eigenvalue of a Hamiltonian is a natural quantum generalization of estimating the maximum number of satisfiable clauses in a Boolean formula.

▶ **Definition 7** (k-LH). *Given a k-local Hamiltonian $H = \sum_{i=1}^{m} H_i$ acting on n qubits with $m = \mathrm{poly}(n)$, the entries of each H_i specified by at most $\mathrm{poly}(n)$ bits, and the norms $\|H_i\|$ polynomially bounded in n, and two real parameters b, a such that $b - a \geq 1/\mathrm{poly}(n)$, decide whether $\lambda_{\min}(H)$ is at most a (YES) or at least b (NO), promised that one is the case.*

In this work, we are interested in k-LH restricted to families of local Hamiltonians, where the families are determined by sets of allowed interactions. In particular, we will be interested in sets of 2-qubit interactions.

▶ **Definition 8** ($\mathcal{S}$-LH). *For $\mathcal{S}$ any fixed set of Hamiltonian terms, define $\mathcal{S}$-LH as the problem k-LH with the additional promise that any input is of the form $\sum_{i=1}^{m} w_i H_i$ where each H_i is an element of $\mathcal{S}$ assigned to act on some subset of qubits and the* weights $w_i \in \mathbb{R}$ *have magnitude polynomially-bounded in n.*

▶ Remark 9. There are several standard assumptions implicit in our definition of $\mathcal{S}$-LH. Some are not physically realistic in the context of the condensed-matter literature but allow us to precisely characterize the complexity of these problems. First, although classical CSPs generally allow a constraint to take as input multiple copies of the same variable, this makes

less sense in the quantum setting and we do not allow it. Second, the definition of k-local only restricts the dimension of each term, it does not imply any spatial locality or geometry. Therefore, any term in H may be applied to any subset of qubits with the qubits arranged in any order. In particular this means that, if $\mathcal{S}$ contains a directed term H^{ab}, then the family of Hamiltonians allowed as input to $\mathcal{S}$-LH is equivalent to the family allowed to $\mathcal{S}'$-LH for $\mathcal{S}' = \mathcal{S} \cup \{H^{ba}\}$. Third, for the purpose of classifying the complexity of $\mathcal{S}$-LH, we may assume $I \in \mathcal{S}$, since adding or removing a term wI is equivalent to simply shifting the input parameters a, b by w. For $\mathcal{S}$ containing 2-qubit terms, this fact also implies we may assume all terms in $\mathcal{S}$ are traceless. Fourth, except when noted, we allow both positive and negative weights.

Classifying the complexity of systems under additional, more physically natural restrictions appears to be a significantly more difficult problem [14, 33].

Given this setup, our interest will be in the problems k-LH and $\mathcal{S}$-LH restricted to product states.

▶ **Definition 10** (Product state). *A state $\rho = \bigotimes_{i=1}^{n} \rho^i$ where each ρ^i is a single-qubit state.*

▶ **Definition 11** (PRODLH). *Given a k-local Hamiltonian $H = \sum_{i=1}^{m} H_i$ on n qubits with $m = \text{poly}(n)$, the entries of each H_i specified by at most $\text{poly}(n)$ bits, and the norms $\|H_i\|$ polynomially bounded in n, and two real parameters $b \geq a$, decide whether there exists a product state ρ with $\text{tr}(\rho H) \leq a$ (YES) or all product states satisfy $\text{tr}(\rho H) \geq b$ (NO), promised that one is the case.*

The problem $\mathcal{S}$-PRODLH is defined analogously. In both definitions, the fact that product states have concise classical descriptions allows us to naturally consider any choice of parameters, even an exact decision problem with $b = a$, in contrast to k-LH.

By convexity, a product state ρ achieves an extreme value of $\text{tr}(\rho H)$ if and only if there exists a pure product state $|\psi\rangle$ which achieves that value. Similarly, mixtures of product states, known as separable states, reduce to considering pure product states.

▶ Remark 12. In the context of $\mathcal{S}$-LH or $\mathcal{S}$-PRODLH given some fixed set S, we will describe $\mathcal{S}$ as "containing" a Hamiltonian term H, or that we "have access to" H, even when formally $H \notin \mathcal{S}$. As previously referenced in Remark 9, given a set $\mathcal{S}$, the family of Hamiltonians allowed as input to $\mathcal{S}$-LH may be equivalent to the family allowed given some other set $\mathcal{S}'$. For example, $\mathcal{S}'$ may include $\{PHP^\dagger\}$ for $H \in \mathcal{S}$ and any permutation of the qubits P. Similarly, adding constant multiples of the terms in $\mathcal{S}$ or any linear combinations of terms from $\mathcal{S}$ does not change the family of allowed Hamiltonians. So, when discussing $\mathcal{S}$-LH, we may implicitly refer to elements of the largest set $\mathcal{S}'$ such that $\mathcal{S}$-LH and $\mathcal{S}'$-LH each have the same family of allowed inputs.

Additionally, we note that $\mathcal{S}$-LH is reducible to $\mathcal{S}'$-LH for any $\mathcal{S}'$ which can be used to implement all elements of $\mathcal{S}$ – whether because formally $\mathcal{S} \subseteq \mathcal{S}'$ or through other means. In the opposite direction, if the terms in a set $\mathcal{S}$ can be used to construct some term H and we wish to show hardness of $\mathcal{S}$-LH, then it is sufficient to show hardness of $\{H\}$-LH.

Finally, for a 2-local Hamiltonian, we may refer to the *interaction graph*, with vertices associated with each qubit such that vertex i is adjacent to vertex j whenever a nonzero interaction exists on qubits i and j. When all interactions are symmetric, then the graph is undirected. Notably, when $\mathcal{S}$-LH is defined with $\mathcal{S}$ a singleton, then an input is fully specified by its weighted interaction graph.

3 Classification of $\mathcal{S}$-prodLH

In this section, we prove a dichotomy theorem classifying the complexity of estimating the minimum expectation of product states for given families of 2-local Hamiltonians. In particular, we show that for any set $\mathcal{S}$ of 2-qubit terms such that at least one term is not 1-local,[4] the problem $\mathcal{S}$-prodLH is NP-complete. These are precisely those 2-local families such that (as shown in [14]) $\mathcal{S}$-LH is QMA-, StoqMA-, or NP-complete. Conversely, if all terms are 1-local, then both $\mathcal{S}$-LH and $\mathcal{S}$-prodLH are in P.

▶ **Theorem 1.** *For any fixed set of 2-qubit Hamiltonian terms $\mathcal{S}$, if every matrix in $\mathcal{S}$ is 1-local then $\mathcal{S}$-prodLH is in P, and otherwise $\mathcal{S}$-prodLH is NP-complete.*

Our NP-hardness results hold with coupling strengths of at most constant magnitude.

For comparison with our classification, we recall the tetrachotomy theorem of Cubitt and Montanaro [14] classifying $\mathcal{S}$-LH for families of 2-local Hamiltonians. They proved that for every set of 2-qubit Hamiltonian terms $\mathcal{S}$, the problem $\mathcal{S}$-LH is either in P or NP-, StoqMA-, or QMA-complete, and described properties of the set $\mathcal{S}$ which determine the problem's complexity. We note that both Theorem 13 and our Theorem 1 classify the complexity of *all* sets of 2-qubit terms.

▶ **Theorem 13** (Theorem 7 of [14])**.** *For $\mathcal{S}$ any fixed set of 2-qubit Hamiltonian terms:*

- *If every matrix in $\mathcal{S}$ is 1-local, then $\mathcal{S}$-LH is in P.*
- *Otherwise, if there exists a single-qubit unitary U such that U locally diagonalizes all elements of $\mathcal{S}$ (i.e. $U^{\otimes 2} H U^{\dagger \otimes 2}$ is diagonal for each 2-qubit $H \in \mathcal{S}$), then $\mathcal{S}$-LH is NP-complete.*
- *Otherwise, if there exists a single-qubit unitary U such that for each 2-qubit $H \in \mathcal{S}$,*

$$U^{\otimes 2} H U^{\dagger \otimes 2} = \alpha Z^{\otimes 2} + A \otimes I + I \otimes B$$

for some $\alpha \in \mathbb{R}$ and 1-local Hermitian matrices A, B, then $\mathcal{S}$-LH is StoqMA-complete.

- *Otherwise, $\mathcal{S}$-LH is QMA-complete.*

Combining our classification of $\mathcal{S}$-prodLH with the classification of $\mathcal{S}$-LH gives us Corollaries 2 and 3.

To prove Theorem 1, showing containment in P and NP are straightforward, and our effort is to prove NP-hardness. In the proof, we will use a simple symmetrization trick that allows us to consider only antisymmetric or symmetric Hamiltonian terms. We then prove two lemmas, one for each case.

▶ **Lemma 14.** *If $\mathcal{S}$ contains a 2-qubit antisymmetric term that is not 1-local, then $\mathrm{MC}^{\mathrm{L}}_W$ with $W = \mathrm{diag}(1, 1, 0)$ is polynomial-time reducible to $\mathcal{S}$-prodLH.*

▶ **Lemma 15.** *If $\mathcal{S}$ contains a 2-qubit symmetric term that is not 1-local, then there exists a fixed non-negative $W = \mathrm{diag}(\alpha, \beta, \gamma)$ with at least one of α, β, γ nonzero such that $\mathrm{MC}^{\mathrm{L}}_W$ is polynomial-time reducible to $\mathcal{S}$-prodLH.*

In Section 4 we prove Theorem 4, that $\mathrm{MC}^{\mathrm{L}}_W$ is NP-complete for any $W = \mathrm{diag}(\alpha, \beta, \gamma)$ that is nonzero and non-negative.

We state some helpful facts in Section 3.1 below and then prove the two lemmas in Section 3.2. We will now use these lemmas to prove our main theorem.

[4] We would prefer a more concise name for 2-qubit terms that are not 1-local, but are unaware of any. One option is 2-qubit terms with *Pauli degree* 2. Alternatively, these are 2-qubit terms which have nonzero *Pauli rank*, referring to the rank of the correlation matrix M in the Pauli decomposition (Equation (5)).

Proof of Theorem 1. First consider the case where $\mathcal{S}$ only contains 1-local terms. Then we can write $H = \sum_i H_i$, where H_i acts only on the i^{th} qubit. If $|\psi_i\rangle$ is the single-qubit state minimizing $\langle\psi_i|H_i|\psi_i\rangle$, then $|\psi\rangle = \bigotimes_i |\psi_i\rangle$ minimizes $\langle\psi|H|\psi\rangle$, and so $\mathcal{S}$-prodLH can be solved by finding the ground state of n single-qubit Hamiltonians, which takes $O(n)$ time.

Now suppose $\mathcal{S}$ is a set of of 2-qubit Hamiltonian terms such that at least one element of $\mathcal{S}$ is not 1-local. Let H be any such element. As previously mentioned, $\mathcal{S}$-prodLH $\in$ NP for any fixed $\mathcal{S}$, as product states have concise classical descriptions which can be used to efficiently calculate expectation values for a given local Hamiltonian. If H is antisymmetric, then MC^{L}_W with fixed $W = \text{diag}(1, 1, 0)$ is reducible to $\mathcal{S}$-prodLH by Lemma 14, and MC^{L}_W with such a W is NP-hard by Theorem 4, so $\mathcal{S}$-prodLH is NP-complete. If H is symmetric, then by Lemma 15 there exists a non-negative nonzero matrix $W = \text{diag}(\alpha, \beta, \gamma)$ such that MC^{L}_W is reducible to $\mathcal{S}$-prodLH, it is NP-hard by Theorem 4, and so $\mathcal{S}$-prodLH is NP-complete. If H is neither of these, then we use our freedom to permute the direction H is applied to any pair of qubits a, b to apply both H^{ab} and H^{ba}, which is equivalent to implementing the symmetric term $H' = H + \text{SWAP}\,(H)\,\text{SWAP}$. So, we can say $\mathcal{S}$ effectively contains H', or formally that hardness of $\{H'\}$-prodLH implies hardness of $\mathcal{S}$-prodLH, and again referring to Lemma 15 concludes the proof. ◀

3.1 Closure Properties of $\mathcal{S}$-prodLH

Before proving the two lemmas required in the proof of Theorem 1, we review several more facts regarding 2-qubit Hamiltonian terms and operations under which the complexities of $\mathcal{S}$-LH and $\mathcal{S}$-prodLH are unaffected. This section mostly reviews observations made in [14].

First, for a single-qubit unitary U and an operator H, define *simultaneous conjugation* by U to mean $U^{\otimes n} H U^{\dagger \otimes n}$. When discussing sets $\mathcal{S}$ of k-qubit terms, we define simultaneous conjugation to mean $\{U^{\otimes k} H U^{\dagger \otimes k} | H \in \mathcal{S}\}$.

▶ **Fact 16.** *For any single-qubit unitary U, the complexities of $\mathcal{S}$-LH and $\mathcal{S}$-prodLH are equal to the complexities of $\mathcal{S}'$-LH and $\mathcal{S}'$-prodLH, respectively, where $\mathcal{S}'$ is $\mathcal{S}$ simultaneously conjugated by U.*

Observe that $U^{\otimes n} \left(\sum_{i=1}^m H_i\right) U^{\dagger \otimes n} = \sum_{i=1}^m U^{\otimes k} H_i U^{\dagger \otimes k}$. Simultaneous conjugation by U gives a bijection between Hamiltonians allowed in $\mathcal{S}$-LH and $\mathcal{S}'$-LH as well as $\mathcal{S}$-prodLH and $\mathcal{S}'$-prodLH. The above fact follows from observing that this bijection preserves expectation values, and that $U^{\otimes n} |\phi\rangle$ is a product state iff $|\phi\rangle$ is.

As an application of Fact 16, the following is based on an observation in [14].

▶ **Fact 17.** *For any choice of permutation π on $\{1, 2, 3\}$ and any choice of* two *of $c_1, c_2, c_3 = \pm 1$, there exists a single-qubit unitary U and corresponding third coefficient s.t. simultaneous conjugation by U maps the Pauli matrices $\{\sigma_1, \sigma_2, \sigma_3\}$ to $\{c_{\pi(1)}\sigma_{\pi(1)}, c_{\pi(2)}\sigma_{\pi(2)}, c_{\pi(3)}\sigma_{\pi(3)}\}$. So, writing every element of $\mathcal{S}$ in the Pauli basis, relabeling all σ_i with $c_{\pi(i)}\sigma_{\pi(i)}$ in the decompositions of each element of $\mathcal{S}$ does not change the complexity of $\mathcal{S}$-LH or $\mathcal{S}$-prodLH, where π and two of the coefficients can be chosen arbitrarily.*

To justify the above fact, consider simultaneously rotating the three axes of the Bloch sphere. Next, we quote the following, more involved, fact without proof.

▶ **Fact 18** ([14, 21]). *Let H be a 2-qubit Hamiltonian term with Pauli decomposition $H = \sum_{i,j=1}^{3} M_{ij}\sigma_i\sigma_j + \sum_{k=1}^{3}(v_k\sigma_k I + w_k I\sigma_k)$. For any single-qubit unitary U,*

$$U^{\otimes 2}H(U^\dagger)^{\otimes 2} = \sum_{i,j=1}^{3}(RMR^T)_{ij}\sigma_i\sigma_j + \sum_{k=1}^{3}((Rv)_k\sigma_k I + (Rw)_k I\sigma_k) \tag{6}$$

for some $R \in SO(3)$. Likewise, for any $R \in SO(3)$, there exists a single-qubit U such that the Pauli decomposition of $U^{\otimes 2}H(U^\dagger)^{\otimes 2}$ matches Equation (6).

A further straightforward observation from [14] is that in the Pauli decomposition (Equation (5)), if H is symmetric then the correlation matrix M is symmetric, and if H is antisymmetric then M is skewsymmetric, meaning $M = -M^\top$.

Finally, the below observation combines some of the above facts to establish a "normal form" for symmetric and antisymmetric terms.

▶ **Fact 19.** *If a 2-qubit Hamiltonian term H is symmetric, and so the associated correlation matrix M is symmetric, there exists $R \in SO(3)$ which diagonalizes M. Combining Facts 16 and 18, for any 2-qubit symmetric term H, there exists a symmetric term of the form $H' = \sum_{i=1}^{3} u_i\sigma_i\sigma_i + \sum_{j=1}^{3} v_j(\sigma_j I + I\sigma_j)$ such that the complexities of $\{H\}$-LH and $\{H\}$-prodLH are respectively the same as $\{H'\}$-LH and $\{H'\}$-prodLH.*

If H is a 2-qubit antisymmetric term that is not 1-local, then M is skewsymmetric and nonzero. Such an M may be block diagonalized via some $R \in SO(3)$ such that H is mapped to $a(\sigma_i\sigma_j - \sigma_j\sigma_i) + \sum_{k=1}^{3} v_k(\sigma_k I - I\sigma_k)$ [37, 14]. In particular, by Fact 17, H can be mapped to $a(XZ - ZX) + \sum_{k=1}^{3} v'_k(\sigma_k I - I\sigma_k)$. Therefore, the complexities of $\{H\}$-LH and $\{H\}$-prodLH are unaffected by assuming H has this form.

3.2 Proofs of Antisymmetric and Symmetric Lemmas

We now prove the two lemmas required in the proof of the main theorem, respectively handling the cases that $\mathcal{S}$ contains an antisymmetric term and that $\mathcal{S}$ contains a symmetric term. In both cases, it is sufficient for $\mathcal{S}$ to contain just a single term. Interestingly, our construction in Lemma 14 for antisymmetric terms is unweighted, meaning all weights are $+1$. In this case, the final Hamiltonian is fully determined by the specification of a single 2-qubit term and the interaction graph. Our construction in Lemma 15 uses positive and negative unit weights, ± 1. Intuitively, antisymmetric terms inherently allow negativity by simply permuting the qubits they act on, while for symmetric terms we must use negative weights.

Proof of Lemma 14. Consider an arbitrary instance of the problem $\mathrm{MC}^{\mathrm{L}}_W$ with $W = \mathrm{diag}(1, 1, 0)$. For a given graph $G = (V, E)$, the objective function is

$$\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}(G) = \frac{1}{2}\max_{i\in S^2}\sum_{ij\in E}\|W\hat{\imath} - W\hat{\jmath}\|.$$

Given a parameter C, we must decide whether $\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}(G)$ is at least C or at most $C - \varepsilon$. To reduce $\mathrm{MC}^{\mathrm{L}}_W$ to $\mathcal{S}$-prodLH, we first construct a gadget using the promised antisymmetric term. Then, we apply this gadget according to the graph G such that the minimum energy of a product state in our final Hamiltonian will equal $-\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}(G)$.

Denote the assumed 2-qubit antisymmetric term that is not 1-local in $\mathcal{S}$ by H. By Fact 18, our antisymmetric term H may be mapped to a term of the form

$$w\left(X^aZ^b - Z^aX^b\right) + \sum_{k=1}^{3} v_k\left(\sigma_k^a I^b - I^a\sigma_k^b\right)$$

where all coefficients are real, and we have $w \neq 0$ since the term is not 1-local. As explained in Remark 12 and Section 3.1, the complexity of $\{H\}$-PRODLH is equivalent to that of $\{H'\}$-PRODLH for H' derived using a variety of operations, including permutations and linear combinations. If w is negative, then we redefine the direction the term acts in, H^{ab} versus H^{ba}, so that w is positive. Finally, we scale[5] the term so that $w = 1$ and define a single-qubit Hermitian matrix $A = \sum_{k=1}^{3} v_k \sigma_k$. Given the complexity is unchanged using H or H', we simply redefine the original term, so that

$$H^{ab} = X^a Z^b - Z^a X^b + A^a I^b - I^a A^b.$$

Next, we use a symmetrization gadget to remove the 1-local terms $AI - IA$. For four qubits a, b, c, d, define

$$B = H^{ab} + H^{bc} + H^{cd} + H^{da}.$$

Note that here the direction of the interaction matters, since the terms are asymmetric. Then

$$B = (X^a - X^c)(Z^b - Z^d) - (Z^a - Z^c)(X^b - X^d).$$

Now we consider how B interacts with product states on four qubits. For $e = a, b, c, d$, define

$$\rho^e = \frac{1}{2}(I + r^e \cdot v^e)$$

with $v^e = (X^e, Y^e, Z^e)$ the Pauli operators and $r^e = (x^e, y^e, z^e)$ the Bloch vector. Then, writing any product state on qubits a, b, c, d as $\rho^a \rho^b \rho^c \rho^d$, the expectation value on B is $\operatorname{tr}\left(B \rho^a \rho^b \rho^c \rho^d\right)$. After eliminating terms, we find this equals

$$(r^a - r^c)^\top W' (r^b - r^d) \quad \text{for} \quad W' = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}.$$

It is helpful to note that $W'^\top = -W'$.

Now, consider the graph G given as input to $\mathrm{MC}^{\mathrm{L}}_W$. Associate a qubit with each vertex and call these the "vertex qubits". For each edge ij, construct a copy of B such that it acts on qubits ij and two ancilla qubits. The vertex qubits may be shared among several gadgets, while the ancilla qubits are part of only one gadget. In particular, we choose to associate the vertex qubits with qubits a and c in each copy of B, letting b and d be the ancilla. We will refer to the copy of B which acts on vertex qubits i and j as B^{ij}. Our Hamiltonian is then

$$H_{\text{final}} = \sum_{ij \in E} B^{ij}.$$

Before analyzing the full Hamiltonian H_{final}, consider the minimum expectation of a single gadget B^{ij} if the two vertex qubits are fixed, i.e. $\min_{r^b, r^d} (r^i - r^j)^\top W' (r^b - r^d)$. The minimum is achieved when $r^b = -W'^\top (r^i - r^j) / \left\| W'^\top (r^i - r^j) \right\|$ and $r^d = -r^b$, which yields an expectation of

$$-2 \left\| W'' (r^i - r^j) \right\| \quad \text{for} \quad W'' = \operatorname{diag}(1, 0, 1).$$

[5] If we want the terms to have unit weights, we could forgo scaling the term and reduce to $w\,\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}$ instead. As $w > 0$ this problem has the same complexity as $\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}$.

Therefore, given an arbitrary graph $G = (V, E)$, applying our gadget to every edge constructs a Hamiltonian such that the minimum expectation of any product state is equal to

$$2 \min_{r^i \in S^2} \sum_{ij \in E} -\left\|W''(r^i - r^j)\right\| = -2 \max_{r^i \in S^2} \sum_{ij \in E} \left\|W''(r^i - r^j)\right\|.$$

For any graph G, the objective $\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}(G)$ is equal for $W = \mathrm{diag}(1, 1, 0)$ and $W'' = \mathrm{diag}(1, 0, 1)$. Alternatively, we may use our freedom to relabel Paulis to redefine the 2-qubit term H such that the final weight matrix is W.

Finally, multiplying the full Hamiltonian H_{final} by $\frac{1}{2}$ gives us that the minimum expectation of any product state equals

$$-\max_{r^i \in S^2} \sum_{ij \in E} \left\|W(r^i - r^j)\right\| = -\,\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}(G).$$

We conclude deciding $\mathrm{MC}^{\mathrm{L}}_{W}$ reduces to deciding PRODLH on H_{final}. Since H_{final} is entirely constructed from the antisymmetric term $H \in \mathcal{S}$, this completes the desired reduction of $\mathrm{MC}^{\mathrm{L}}_{W}$ with $W = \mathrm{diag}(1, 1, 0)$ to $\mathcal{S}$-PRODLH. ◀

Next we prove the lemma dealing with $\mathcal{S}$ containing a symmetric term. The construction is nearly the same as the in the previous proof, but requires negative weights to implement the symmetrization gadget removing 1-local terms.

Proof of Lemma 15. Given fixed $\mathcal{S}$, we will show there exists some fixed W such that $\mathrm{MC}^{\mathrm{L}}_{W}$ reduces to $\mathcal{S}$-LH. But, before describing W, we must analyze $\mathcal{S}$.

Denote the assumed 2-qubit symmetric term that is not 1-local in $\mathcal{S}$ by H. As in the previous proof, without changing the complexity of $\{H\}$-PRODLH we may conjugate and scale as necessary so that

$$H^{ab} = \alpha^- X^a X^b + \beta^- Y^a Y^b + \gamma^- Z^a Z^b + \sum_{j=1}^{3} v_j(\sigma_j^a I^b + I^a \sigma_j^b),$$

where all coefficients are real and at least one of $\alpha^-, \beta^-, \gamma^-$ is nonzero since H is nonzero. The superscripts in the above equations are to differentiate the coefficients of H from the entries of W, which must be non-negative.

We again use a symmetrization gadget to remove the 1-local terms, but now require negative weights. Given four qubits a, b, c, d, define $B = H^{ab} + H^{cd} - H^{bd} - H^{ac}$. This is a rectangle with two positive edges and two negative edges. Then

$$B = \alpha(X^a - X^d)(X^b - X^c) + \beta(Y^a - Y^d)(Y^b - Y^c) + \gamma(Z^a - Z^d)(Z^b - Z^c).$$

Now we see how B interacts with product states on four qubits. For $e = a, b, c, d$, we again define $\rho^e = \frac{1}{2}(I + r^e \cdot v^e)$ with $v^e = (X^e, Y^e, Z^e)$ and $r^e = (x^e, y^e, z^e)$. Then, writing any product state on a, b, c, d as $\rho^a \rho^b \rho^c \rho^d$, the expectation value on B is $\mathrm{tr}\left[B \rho^a \rho^b \rho^c \rho^d\right]$, which equals

$$\alpha^-(x^a - x^d)(x^b - x^c) + \beta^-(y^a - y^d)(y^b - y^c) + \gamma^-(z^a - z^d)(z^b - z^c)$$

which is in turn equal to

$$\left(r^b - r^c\right)^\top W' \left(r^a - r^d\right) \quad \text{for} \quad W' = \mathrm{diag}\left(\alpha^-, \beta^-, \gamma^-\right).$$

If we fix the state of qubits a and d and minimize the expectation on B, the minimum is achieved when $r^b = -W'(r^a - r^d)/\|W'(r^a - r^d)\|$ and $r^c = -r^b$. This minimum expectation is $-2\|W'(r^a - r^d)\|$. Observe this expectation value is equivalent to $-2\|W(r^a - r^d)\|$ for $W = \mathrm{diag}(\alpha, \beta, \gamma)$ where $\alpha = |\alpha^-|, \beta = |\beta^-|, \gamma = |\gamma^-|$.

Now we are prepared to set up our reduction. For $W = \mathrm{diag}(\alpha, \beta, \gamma)$, which is non-negative and nonzero, consider an arbitrary instance of $\mathrm{MC}^{\mathrm{L}}_W$. For a given graph $G = (V, E)$, the objective function is again $\mathrm{MC}^{\mathrm{L}}_W(G)$, and given a parameter C, we must decide whether $\mathrm{MC}^{\mathrm{L}}_W(G)$ is at least C or at most $C - \varepsilon$.

Associate a "vertex qubit" with each vertex and construct a copy of the gadget B for each edge ij, such that it acts on i, j and two ancilla qubits, and denote it B^{ij}. The vertex qubits may be shared among several gadgets, while the ancilla qubits are part of only one gadget. In particular, we choose a and d in each gadget to be the vertex qubits.

Substituting our gadget for every edge constructs a Hamiltonian H_{final} such that the minimum expectation of any product state is equal to

$$2 \min_{r^i \in S^2} \sum_{ij \in E} -\|Wr^i - Wr^j\| = -2 \max_{r^i \in S^2} \sum_{ij \in E} \|Wr^i - Wr^j\|.$$

Simply multiplying H_{final} by $\frac{1}{2}$ makes this equal to $-\mathrm{MC}^{\mathrm{L}}_W(G)$.

We conclude that given $\mathcal{S}$ contains a 2-qubit symmetric term H that is not 1-local, there exists some non-negative $W = \mathrm{diag}(\alpha, \beta, \gamma)$ with at least one nonzero entry such that $\mathrm{MC}^{\mathrm{L}}_W(G)$ reduces to $\textsc{prodlh}(H_{\text{final}})$. Since H_{final} was constructed using only the symmetric term $H \in \mathcal{S}$, this is also a reduction to an instance of $\mathcal{S}$-$\textsc{prodlh}$, as desired. ◀

4 The Stretched Linear Max-Cut Problem

We study a generalization of the classical Max-Cut problem which arises naturally from our study of product states and which is likely of independent interest. Both Max-Cut and its generalization

$$\mathrm{MC}_k(G) = \frac{1}{2} \max_{\hat{\imath} \in S^{k-1}} \sum_{ij \in E} 1 - \hat{\imath} \cdot \hat{\jmath} = \frac{1}{4} \max_{\hat{\imath} \in S^{k-1}} \sum_{ij \in E} \|\hat{\imath} - \hat{\jmath}\|^2$$

were introduced in Section 1. As the above equation emphasizes, maximizing the distance between vectors is equivalent to maximizing the angle, optimally being anti-parallel.

Our new problem is defined with two significant changes. First, the sum is over distances rather than squared distances. Second, the distances are allowed to incorporate a linear stretch.

▶ **Definition 20** (W-linear-Max-Cut ($\mathrm{MC}^{\mathrm{L}}_W$)). *For a fixed $d \times d$ diagonal matrix W, given an n-vertex graph $G = (V, E)$ and thresholds $b > a \geq 0$ with $b - a \geq 1/\mathrm{poly}(n)$, decide whether*

$$\mathrm{MC}^{\mathrm{L}}_W(G) = \frac{1}{2} \max_{\hat{\imath} \in S^{d+1}} \sum_{ij \in E} \|W\hat{\imath} - W\hat{\jmath}\| = \frac{1}{2} \max_{\hat{\imath} \in S^{d+1}} \sum_{ij \in E} \sqrt{\|W\hat{\imath}\| + \|W\hat{\jmath}\| - 2(W\hat{\imath})^\top (W\hat{\jmath})}$$

is at least b or at most a.

A comparison of the geometric interpretations of MC_k and $\mathrm{MC}^{\mathrm{L}}_W$ was given in Section 1. A further interpretation comes from treating the edges of the graph as springs or rubber bands. As explored in [29], the potential energy of a spring is quadratic in its length, so the MC_k value represents the total potential energy of the system given a particular embedding. On

the other hand, the force or tension of each spring is linear in its length. So, $\mathrm{MC}^{\mathrm{L}}_W$ gives the total force, tension, or pressure such an arrangement of springs would apply to the surface of the sphere (or ellipsoid, more generally).

In both problems, the objective is a linear sum, of either the distances or the inner products. Both problems generalize the traditional Max-Cut problem, since when restricted to ± 1 labels, distances are directly proportional to squared distances. Previous work was likely motivated to focus on squared distances because approximation algorithms like SDPs naturally apply to inner products but not to square roots of inner products.

Our main theorem concerning W-linear-Max-Cut is the below.

▶ **Theorem 4.** *For any fixed non-negative $W = \mathrm{diag}(\alpha, \beta, \gamma)$ with at least one of α, β, γ nonzero, W-linear-Max-Cut is NP-complete.*

Our hardness proofs do not require any edge weights (unlike our Hamiltonian constructions in the previous section).

Containment in NP is immediate, and we break the proof of NP-hardness into three cases based on the entries of W. The three cases depend on how many entries of W are equal, requiring different approaches for dealing with degenerate solutions. We assume throughout that $1 = \alpha \geq \beta \geq \gamma$; as we show in the final proof of Theorem 4, this suffices by scaling and symmetry. Lemma 21 considers the case when all three entries are equal. Lemma 22 considers the case when the largest entry is unique. Lemma 23 finally considers the case when the two largest entries are equal and distinct from the third, combining techniques from the previous two proofs. Note that these cases are not entirely disjoint.

When $W = \mathrm{diag}(1, 1, 1)$, we prove hardness by reducing from the NP-complete 3-Coloring problem. We replace every edge in the graph with a 4-clique, or tetrahedron. To deal with the symmetry created by equally weighted axes, all of the gadgets are connected to a new *sink* vertex, removing a degree of freedom. We then argue that there is an assignment to the new graph that simultaneously (nearly) maximizes all of these cliques iff the original graph is 3-colorable.

▶ **Lemma 21.** *For $W = \mathrm{diag}(1, 1, 1)$, W-linear-Max-Cut is NP-hard.*

Proof. We will reduce from 3-Coloring. Consider an arbitrary graph $G = (V, E)$ on n vertices and m edges. We construct $H = (V', E')$ as follows: Start with G. For each edge $ij \in E$, add vertices k_{ij}, t_{ij} and connect i, j, k_{ij}, t_{ij} to form a 4-clique. Then add a *sink* vertex t and an edge tt_{ij} for each t_{ij}. H therefore consists of m edge-disjoint 4-cliques, each containing one edge from G, and m additional edges from vertices t_{ij} to t.

We claim that if G is 3-colorable, then $\mathrm{MC}^{\mathrm{L}}_W(H) \geq m\mathrm{MC}^{\mathrm{L}}_W(K_4) + m$. Conversely, we claim that if $\mathrm{MC}^{\mathrm{L}}_W(H) \geq m\mathrm{MC}^{\mathrm{L}}_W(K_4) + m - \varepsilon$, for an $\varepsilon = \Omega\big(1/m^2\big)$ we will choose later, then G is 3-colorable.

First, suppose G is 3-colorable. We will show how to derive a vector assignment to H attaining $m\mathrm{MC}^{\mathrm{L}}_W(K_4) + m$ from a 3-coloring of G. Define

$$\Delta_W(u, v, w, r) = \frac{1}{2} \sum_{ij \in \{uv, uw, ur, vw, vr, wr\}} \|Wi - Wj\|$$

Let u, v, w, r be vectors corresponding to a regular tetrahedron inscribed in the unit sphere, known to achieve the maximum perimeter of any inscribed tetrahedron at $4\sqrt{6}$ [30]. We 3-color G (and therefore the vertices of H other than $(k_{ij})_{ij \in E}$, $(t_{ij})_{ij \in E}$, and t) with u, v, w, assigning each vertex the vector matching its color. Then for each $ij \in E$, we assign k_{ij}

the vector in $\{u, v, w\}$ not assigned to i or j, and r to t_{ij}. Finally, we assign $-r$ to t. By construction, this assignment will yield a value of $\mathrm{MC}^{\mathrm{L}}_W(K_4)$ on each 4-clique gadget, and the edges tt_{ij} will each contribute exactly 1. Thus $\mathrm{MC}^{\mathrm{L}}_W(H) \geq m\mathrm{MC}^{\mathrm{L}}_W(K_4) + m$, as desired.

Now we will show the converse. Suppose there exists an assignment $(\hat{k})_{k\in V'}$ of vectors achieving greater than $m\mathrm{MC}^{\mathrm{L}}_W(K_4)+m-\varepsilon$ on H. We construct a 4-coloring of each connected component of $H \setminus \{t\}$ as follows: Choose an edge ij in the corresponding component of G (note that there is a 1-to-1 correspondence between components of G and $H \setminus \{t\}$). Let $\hat{\imath}, \hat{\jmath}, \hat{k}_{ij}, \hat{t}_{ij}$ be the vectors assigned to the vertices of the corresponding clique in H. Our coloring will use these four vectors as colors, which we denote as the set $\mathcal{C}$. We assign each vertex v the color corresponding to the element of $\mathcal{C}$ that is closest to the vector assigned to v in the $\mathrm{MC}^{\mathrm{L}}_W$ assignment. We will show that this is a proper coloring, and that it assigns the same color to every t_{ij}, and therefore gives a 3-coloring of G.

To show that this is a proper coloring, we need to show that every pair of adjacent vertices in $H \setminus \{t\}$ are assigned different colors, i.e. that the closest elements of $\mathcal{C}$ to the vectors assigned to them in the $\mathrm{MC}^{\mathrm{L}}_W$ assignment are distinct. As every edge in $H \setminus \{t\}$ is contained in some 4-clique corresponding to some edge xy of G, it will suffice to show the following: for every edge xy in the component of G containing ij, if d is the smallest number of edges in a path in G starting with ij and ending with xy, each of $\hat{x}$, $\hat{y}$, $\hat{k}_{xy}$ and $\hat{t}_{xy}$ is within $\mathrm{O}(d\sqrt{\varepsilon})$ of a *different* element of $\mathcal{C}$.

First we note that, as H consists of m edge-disjoint 4-cliques and m other edges, and the maximum any assignment can earn on an edge is 1, the lower bound on the total amount the assignment earns implies that every 4-clique earns at least $\mathrm{MC}^{\mathrm{L}}_W(K_4) - \varepsilon$ and the other edges $(t_{xy}t)_{xy\in E}$ earn at least $1-\varepsilon$ each. So by trigonometry we immediately have that every $\hat{t}_{xy}$ is within $\mathrm{O}(\sqrt{\varepsilon})$ of $-\hat{t}$, and therefore within $\mathrm{O}(\sqrt{\varepsilon})$ of each other, and $\hat{t}_{ij}$ in particular.

For the other vertices, we proceed by induction on d. We have the desired result trivially for $d = 1$, as in this case $ij = xy$. Now suppose it holds for d. Let xy be the end of a $(d+1)$-edge path starting with ij. Without loss of generality let y be the vertex of xy earlier in the path, and let z be the immediately prior vertex in the path, so by the inductive hypothesis $\hat{y}$, $\hat{z}$, $\hat{k}_{yz}$, and $\hat{t}_{yz}$ are within $\mathrm{O}(d\sqrt{\varepsilon})$ of different elements of $\mathcal{C}$. Furthermore, as both $\{\hat{x}, \hat{y}, \hat{k}_{xy}, \hat{t}_{xy}\}$ and $\{\hat{y}, \hat{z}, \hat{k}_{yz}, \hat{t}_{yz}\}$ are vertices of tetrahedra with perimeters at least $4\sqrt{6} - \varepsilon$, by Lemma 29 the edge lengths of these tetrahedra are all in the interval $[\frac{4\sqrt{6}}{6} - \mathrm{O}(\sqrt{\varepsilon}), \frac{4\sqrt{6}}{6} + \mathrm{O}(\sqrt{\varepsilon})]$. So as we have already shown that $\hat{t}_{yz}$, $\hat{t}_{xy}$, and $\hat{t}_{ij}$ are all within $\mathrm{O}(\sqrt{\varepsilon})$ of each other, then the criteria of Lemma 30 are satisfied with $ABCD = \hat{y}\hat{t}_{xy}\hat{x}\hat{k}_{xy}$ and $AEFG = \hat{y}\hat{t}_{yz}\hat{z}\hat{k}_{yz}$ and so $\hat{x}$ and $\hat{k}_{xy}$ are each within $\mathrm{O}(\sqrt{\varepsilon})$ of (different) elements of $\{\hat{z}, \hat{k}_{yz}\}$. So then the result follows by the triangle inequality.

We now have that every vertex in H is assigned a vector within $\mathrm{O}(m\sqrt{\varepsilon})$ of an element of $\mathcal{C}$, and so taking ε to be a sufficiently small constant times $1/m^2$, we can take these distances to be at most 0.1. Moreover, for any two adjacent vertices, the choice of color will be different, and so as by Lemma 29 the distances between these four vectors are at least $4/\sqrt{6} - \mathrm{O}(\varepsilon)$, this implies any two adjacent vertices are assigned different colors, and so we have a proper 4-coloring of $H \setminus \{t\}$. Finally, note that, as every $\hat{t}_{xy}$ is within $\mathrm{O}(\varepsilon)$ of every other one, this implies all the t_{xy} were assigned the same color, and so no vertex in G uses this color. So this gives us a proper 3-coloring of G. ◀

Next, in Lemma 22 we consider the case $W = \mathrm{diag}(1, \beta, \gamma)$ and $1 > \beta, \gamma \geq 0$, in which the maximum weight is unique. The approach in the proof of Lemma 21, reducing from 3-Coloring, can in fact be modified to cover this case, but analyzing triangles inscribed in ellipses instead of circles is more technical. Instead, we take a different approach and in the case of a unique maximum (whether β equals γ or not) give a reduction from Max-Cut instead of 3-Coloring.

We insert ancilla vertices so that every vertex in the original graph is the center of a large star. These star gadgets amplify any deviation from the highest-weight axis such that any near-optimal solution must approximate a standard 1-dimensional labeling, as in Max-Cut.

▶ **Lemma 22.** *For any* $W = \mathrm{diag}(1, \beta, \gamma)$ *with* $1 > \beta \geq \gamma \geq 0$, W-linear-Max-Cut *is* NP*-hard.*

Proof. We reduce from the standard NP-complete Max-Cut problem. For any graph $G = (V, E)$ with n vertices and m edges, we construct a graph $H = (V', E')$ with $V \subseteq V'$ and $E \subseteq E'$ by, for each $v \in V$, adding $K = m^3 n$ ancilla vertices v_i, and then adding an edge from each of these vertices to v so that v is the center of a K-star. Now $|V'| = n(1+K)$ and $|E'| = m + Kn$.

We claim that, for any $C > 1$ and for large enough n, $\mathrm{MC}(G) \geq C$ implies $\mathrm{MC}^{\mathrm{L}}_W(H) \geq C + Kn$ and $\mathrm{MC}^{\mathrm{L}}_W(H) > C + Kn - 1/2$ implies $\mathrm{MC}(G) \geq C$.

First, suppose there is a cut of G with value at least C. We construct a corresponding assignment of vectors to vertices in V'. First assign the vector $(1, 0, 0)$ to all vertices in V which are labeled $+1$ in G and $(-1, 0, 0)$ to those with labels -1. Then, for every vertex in $v \in V$, which by construction is at the center of a star of ancilla qubits in H, assign the vector opposite the one assigned to v to each of the ancilla vertices. This assignment of vectors gives an objective value of at least C on the edges from the original graph and Kn on the edges of the star gadgets, and so the $\mathrm{MC}^{\mathrm{L}}_W$ value of this assignment is $C + Kn$.

Now suppose there exists an assignment of vectors achieving $\mathrm{MC}^{\mathrm{L}}_W$ value greater than $C + Kn - 1$ on H. We will show that the cut given by $\mathrm{sgn}(\hat{v}_1)$ for each $v \in V$ (i.e. projecting $\hat{v}$ to the x-axis and checking whether it is ≥ 0 or < 0) has value at least C.

First, for each $v \in V$, let $\widehat{\mathrm{sgn}}(\hat{v}) = (\mathrm{sgn}(\hat{v}_1), 0, 0)$. We will show that this is close to $\widehat{v}$. Because the original graph can contribute at most m to the $\mathrm{MC}^{\mathrm{L}}_W$ objective, and each star gadget can contribute at most K, each star gadget must contribute at least $K + (C - 1/2 - m) \geq K - 1/2 - m > K\big(1 - \frac{2m}{K}\big)$. By Lemma 26, for the K-star to achieve at least $K\big(1 - \frac{2m}{K}\big)$, the vector $\hat{v}$ assigned to v must satisfy

$$\|W\hat{v} - \widehat{\mathrm{sgn}}(v)\| \leq \delta \quad \text{for} \quad \delta = 2\sqrt{\frac{2m}{K}}\sqrt{\frac{1+\beta^2}{1-\beta^2}}.$$

We will use this fact to show that the $\mathrm{MC}^{\mathrm{L}}_W$ value earned by the vector assignment on the original graph G is close to the value of the cut we defined. We have

$$\left|\sum_{ij\in E} \|W\hat{\imath} - W\hat{\jmath}\| - \sum_{ij\in E} \|\widehat{\mathrm{sgn}}(\hat{\imath}) - \widehat{\mathrm{sgn}}(\hat{\jmath})\|\right| = \left|\sum_{ij\in E} (\|W\hat{\imath} - W\hat{\jmath}\| - \|\widehat{\mathrm{sgn}}(\hat{\imath}) - \widehat{\mathrm{sgn}}(\hat{\jmath})\|)\right|$$
$$\leq \left|\sum_{ij\in E} (\|W\hat{\imath} - \widehat{\mathrm{sgn}}(\hat{\imath})\| + \|W\hat{\jmath} - \widehat{\mathrm{sgn}}(\hat{\jmath})\|)\right| \leq 2m\delta = 4\sqrt{\frac{2}{n}}\sqrt{\frac{1+\beta^2}{1-\beta^2}} = \mathrm{O}\big(1/\sqrt{n}\big)$$

which is $< 1/2$ for large enough n.

Using for a second time the fact that the edges of the star gadgets can contribute at most Kn to the $\mathrm{MC}^{\mathrm{L}}_W(H)$, the vector assignment must achieve at least $C - 1/2$ on G, and so

$$C - 1/2 \leq \frac{1}{2}\sum_{ij\in E} \|W\hat{\imath} - W\hat{\jmath}\| \leq \frac{1}{2}\sum_{ij\in E} |\widehat{\mathrm{sgn}}(\hat{\imath}) - \widehat{\mathrm{sgn}}(\hat{\jmath})| + \mathrm{O}\big(1/\sqrt{n}\big)$$

implying that the value of our cut is strictly greater than $C - 1$ for sufficiently large n. Therefore, as it is integer-valued it is at least C, concluding the proof. ◀

Finally, we give Lemma 23, in which $W = \mathrm{diag}(1, 1, \gamma)$ and $1 > \gamma$. Our proof combines the techniques in the previous two proofs. Similar to the proof of Lemma 21, we show hardness by reducing from 3-Coloring. Our construction begins by replacing every edge in the graph with a 3-clique. Then, as in the proof of Lemma 22, we insert large star gadgets on every vertex, forcing solutions away from the low-weight z-axis. With solutions restricted to two dimensions, we are able to argue there is an assignment to the new graph that simultaneously (nearly) maximizes all of the cliques iff the original graph was 3-colorable.

▶ **Lemma 23.** *For any $W = \mathrm{diag}(1, 1, \gamma)$ with $1 > \gamma \geq 0$, W-linear-Max-Cut is* NP*-hard.*

Proof. Consider an arbitrary instance of 3-Coloring on a graph $G = (V, E)$ with n vertices and m edges. Construct a new graph $H' = (V', E')$ by taking G and, for each edge $ij \in E$, adding a vertex k_{ij} and edges ik_{ij} and jk_{ij}, so that each edge in G corresponds to a 3-clique (K_3) in H'. Note that the m cliques constructed this way are edge-disjoint. Next, for each vertex v in H', add $K = m^6$ ancilla vertices v_i, each connected to v so that v is the center of a K-star. Call the final graph $H'' = (V'', E'')$, for which we have $(K + 1)(n + m)$ vertices and $3m + K(n + m)$ edges.

We claim that if G is 3-colorable, then $\mathrm{MC}^{\mathrm{L}}_W(H'') \geq K(n + m) + 3\sqrt{3}m$. Conversely, we claim that if $\mathrm{MC}^{\mathrm{L}}_W(H'') > K(n + m) + 3\sqrt{3}m - \varepsilon$, for an $\varepsilon = \Omega(1/m^2)$ we will choose later, then G is 3-colorable. As testing 3-colorability is NP-hard, this will prove the theorem.

First, suppose G is 3-colorable. Let C be any set of three vectors in the xy-plane achieving the maximum value of $\mathrm{MC}^{\mathrm{L}}_W(K_3) = 3\sqrt{3}$. Given any 3-coloring of G, we assign one of these vectors to each color, and thus assign vectors from C to G with no two adjacent vertices having the same vector. We extend this assignment to H' by, for each of our constructed 3-cliques i, j, k_{ij}, assigning the vector in C that was not assigned to either i or j to k_{ij}. Now each of these cliques contributes $3\sqrt{3}$ to the objective, and as there are m of them and they are edge-disjoint, this contributes $3\sqrt{3}m$ to the objective. To extend to H'' and its star gadgets, every edge of a star centered on vertex v can trivially contribute $\|v\|$ to the objective value. Because all vectors are in the xy-plane, and so the unit circle, this means they all contribute 1. So in total, there exists an assignment with value $3\sqrt{3}m + K(n + m)$.

Now suppose there exists a vector assignment achieving greater than $K(n+m)+3\sqrt{3}m-\varepsilon$ on H''. For a vector $\hat{v} = (\hat{v}_x, \hat{v}_y, \hat{v}_z)$, let $\widehat{\mathrm{sgn}}(\hat{v})$ denote the vector $(\hat{v}_x, \hat{v}_y, 0)/\|(\hat{v}_x, \hat{v}_y, 0)\|$. We assign colors as follows. Choose any of the K_3 gadgets i, j, k_{ij} in H'' corresponding to an edge ij in the original graph. Let $\hat{\imath}, \hat{\jmath}, \hat{k}_{ij}$ be the vectors assigned to the vertices, respectively. Assign three colors to the vertices. Then, choose any K_3 gadget $i, \ell, k_{i\ell}$ adjacent to the first. For each vertex in the second clique, consider its assigned vector and round it via $\widehat{\mathrm{sgn}}$, determine which of the rounded vectors $\widehat{\mathrm{sgn}}(\hat{\imath}), \widehat{\mathrm{sgn}}(\hat{\jmath}), \widehat{\mathrm{sgn}}(\hat{k}_{ij})$ it is closest to, and assign the same color. We continue coloring adjacent cliques in this way, comparing rounded vectors to the original set $\widehat{\mathrm{sgn}}(\hat{\imath}), \widehat{\mathrm{sgn}}(\hat{\jmath}), \widehat{\mathrm{sgn}}(\hat{k}_{ij})$, until the coloring is propagated to the entire graph. This colors all of the vertices in H', which are the centers of star gadgets in H''. We will show no adjacent vertices in H' were assigned the same color. Since G is a subgraph of H', this also implies a proper coloring of G, as desired.

We first show that just as the assigned vectors must achieve a large objective value on the clique gadgets, the star gadgets force the rounded vectors to achieve a similar score. Because the m clique gadgets can each contribute at most $3\sqrt{3}$ to the objective value, the star gadgets in H'' must contribute at least $K(n + m) - \varepsilon$. Similarly, because the edges of each star can contribute at most K, each star must achieve at least $K - \varepsilon$. For any vertex in s in H', consider the star gadget centered on it in H''. As shown in Lemma 26, for the star to achieve $K(1 - \varepsilon/K)$, the vector $\hat{s}$ assigned to s must satisfy

$$\|W\hat{s} - \widehat{\operatorname{sgn}}(\hat{s})\| \leq \delta \quad \text{for} \quad \delta = 2\sqrt{\frac{\varepsilon}{K}}\sqrt{\frac{1+\gamma^2}{1-\gamma^2}}.$$

On the other hand, because the star gadgets can contribute at most $K(n+m)$ to the objective value, the vector assignment must achieve at least $3\sqrt{3}m - \varepsilon$ on the remaining edges, the ones in H'. Given the vectors are close, we have that

$$\left|\sum_{ij\in E'} \|W\hat{\imath} - W\hat{\jmath}\| - \sum_{ij\in E'} \|\widehat{\operatorname{sgn}}(\hat{\imath}) - \widehat{\operatorname{sgn}}(\hat{\jmath})\|\right| = \left|\sum_{ij\in E'} \|W\hat{\imath} - W\hat{\jmath}\| - \|\widehat{\operatorname{sgn}}(\hat{\imath}) - \widehat{\operatorname{sgn}}(\hat{\jmath})\|\right|$$

$$\leq \left|\sum_{ij\in E'} \|W\hat{\imath} - \widehat{\operatorname{sgn}}(\hat{\imath})\| + \|W\hat{\jmath} - \widehat{\operatorname{sgn}}(\hat{\jmath})\|\right| \leq 2|E'|\delta = 12\sqrt{\varepsilon}m^{-2}.$$

Therefore, for $\mu = \varepsilon + 12\sqrt{\varepsilon}m^{-2}$, the set of rounded vectors must achieve at least $3\sqrt{3}m - \mu$ on the rest of H'. Because each of the m clique gadgets can contribute at most $3\sqrt{3}$, the rounded vectors must achieve at least $3\sqrt{3} - \mu$ on each individual clique.

The rounded vectors exist in the xy-plane, which means they are inscribed in the unit circle. Maximizing the sum of the edge lengths on a K_3 gadget is equivalent to maximizing the perimeter of a triangle. For a triangle in the unit circle to have nearly maximal perimeter, it must be nearly regular; as shown in Lemma 27, if the perimeter is at least $3\sqrt{3} - \mu$, then each edge length must be in the interval $\sqrt{3} \pm 4\sqrt{\mu}$.

Now, consider any two adjacent vertices u, v in H', which are at the center of star gadgets in H''. We must show they were assigned different colors. The two vertices exist in some K_3 gadget in H', and the coloring procedure, starting from i, j, k_{ij}, must have reached them in at most m rounds. At each step in the procedure, there is a colored clique and a successive adjacent clique which share a clique. Two nearly regular triangles which share a vertex must nearly share their other vertices; as shown in Lemma 28, the distance between the first clique's rounded vectors and the second clique's is at most $\mathrm{O}\big(\sqrt{\mu}\big)$.

After m rounds of the coloring procedure, we conclude u and v must be assigned colors such that $\widehat{\operatorname{sgn}}(\hat{u}), \widehat{\operatorname{sgn}}(\hat{v})$ are each at distance at most $m \times \mathrm{O}\big(\sqrt{\mu}\big) = \mathrm{O}\big(\varepsilon^{1/4}\big)$ away from the same-colored rounded vectors in the initial clique. Both the initial clique and the clique containing u and v must have vectors separated by at least $\sqrt{3} - \varepsilon$. For sufficiently small ε ($\varepsilon = \Theta\big(1/m^2\big)$ with a sufficiently small constant suffices), $\widehat{\operatorname{sgn}}(\hat{u}), \widehat{\operatorname{sgn}}(\hat{v})$ cannot also be within $\mathrm{O}(\sqrt{\varepsilon})$ of the same initial vector, so we conclude they must be associated with different colors, and this is a proper coloring. ◀

We now conclude with a proof of the main theorem of this section, that W-LINEAR-MAX-CUT is NP-complete for any fixed diagonal 3×3 non-negative nonzero W.

Proof of Theorem 4. The containment of W-LINEAR-MAX-CUT for any diagonal W is straightforward. With W a constant, given a claimed vector assignment, the value $\mathrm{MC}^{\mathrm{L}}_W(G)$ can be verified in time linear in the number of edges.

To show hardness, we make two simplifications. First, because $\mathrm{MC}^{\mathrm{L}}_{cW} = c\mathrm{MC}^{\mathrm{L}}_W$ for a constant c, we can easily reduce to an instance in which we assume the largest entry of W equals 1. Second, although rearranging the entries of diagonal W requires changing any vector assignment, it does not change the objective value. So $\mathrm{MC}^{\mathrm{L}}_{\mathrm{diag}(\alpha,\beta,\gamma)} = \mathrm{MC}^{\mathrm{L}}_{\mathrm{diag}(\beta,\alpha,\gamma)}$ and any other rearrangement of W, and we can assume the entries are ordered $1 \geq \alpha \geq \beta \geq \gamma \geq 0$. With this, the theorem follows by Lemmas 21–23. ◀

5 NP-Hardness of Unweighted Quantum Max-Cut

Our lower bounds elsewhere in this paper are for local Hamiltonian problems in which terms can be given positive *or* negative weight. They apply in the model where all weights are restricted to being of constant weight but require some terms with negative weight to work. In this section we show that this restriction can be removed for one of the best-studied local Hamiltonian problems: the Quantum Max-Cut (or QMC) problem.

The Quantum Max-Cut problem can be defined as $\mathcal{S}$-LH with $\mathcal{F} = \{XX + YY + ZZ\}$. We will write h for $XX + YY + ZZ$.[6]

▶ **Definition 24** (Quantum Max-Cut (QMC)). *Given a Hamiltonian $H = \sum_{ij}^{n} w_{ij} H_{ij}$ acting on n qubits where each $H_{ij} = I - h_{ij}$, all w_{ij} are real, polynomially-bounded, and specified by at most* $\mathrm{poly}(n)$ *bits, and two real parameters b, a such that $b - a \geq 1/\mathrm{poly}(n)$, decide whether $\lambda_{\max}(H)$ is at least b (YES) or at most a (NO).*

Note that we have written QMC as a *maximum* eigenvalue problem (with a flip and shift of the local terms) rather than in terms of the minimum eigenvalue as for k-LH in Definition 7. This is to follow the norm in previous work; note that as both the terms *and* the objective function are flipped, an instance of the problem defined this way will be equivalent to an instance of the corresponding k-LH problem with the same weights.

When minimizing (maximizing) the eigenvalue and the weights are restricted to be non-negative (non-positive), it is referred to as the anti-ferromagnetic Heisenberg model. Flipping the restrictions, e.g. minimizing with non-positive weights, is referred to as the ferromagnetic Heisenberg model. The latter case is trivial when viewed as an optimization problem (as $|00\rangle$ earns 0 on the local term, and so the problem is optimized by assigning $|0\rangle$ to every qubit), so we will be interested in hardness results for the former.

It is straightforward to verify that for two qubits a, b with *pure* states ρ^a, ρ^b,

$$\mathrm{tr}\left(\rho^a \otimes \rho^b \; h\right) = r^a \cdot r^b = 1 - \frac{1}{2}\left\|r^a - r^b\right\|^2,$$

where r^a, r^b are the corresponding Bloch vectors. This shows that deciding QMC restricted to product states, which we denote prodQMC, is equivalent to the standard Vector Max-Cut problem in three dimensions:

▶ **Definition 25** ($\text{Max-Cut}_k(\text{MC}_k)$). *Given an n-vertex graph $G = (V, E)$ and thresholds $b > a \geq 0$ such that $b - a \geq 1/\mathrm{poly}(n)$, decide whether*

$$\text{MC}_k = \frac{1}{2} \max_{\hat{\imath} \in S^{k-1}} \sum_{ij \in E} 1 - \hat{\imath} \cdot \hat{\jmath} = \frac{1}{4} \max_{\hat{\imath} \in S^{k-1}} \sum_{ij \in E} \left\|\hat{\imath} - \hat{\jmath}\right\|^2$$

is at least b or less than a.

Note that this is different from the W-linear-Max-Cut we studied in Section 4 as it considers *squared* distances. Furthermore, while Max-Cut is a classic NP-complete problem, MC_k is not expected to be hard for all values of k, and in particular is tractable when $k = n = |V|$.

Our main result classifying $\mathcal{S}$-prodLH immediately implies $\{h\}$-prodLH, and therefore also prodQMC and MC_3, is NP-complete. However, the proof of Lemma 15 utilizes Hamiltonian gadgets involving negative weights. This leaves open whether prodQMC and

[6] Other work frequently includes a multiplicative factor, e.g. 1/2, in the definition of h and/or of QMC.

MC_3 remain NP-hard on unweighted graphs. We now prove that MC_3 is NP-hard even when restricted to positive unit weights. This is the first published proof of this fact, although we note that a sketch of a different proof was previously known for this specific problem [40].

Our approach is similar to the proof of Lemma 21, which demonstrated hardness of $\mathrm{MC}_W^{\mathrm{L}}$ for $W = I$ by replacing every edge with a 4-clique (with one vertex connected to a source vertex), and showing that the resulting graph could simultaneously optimize all of these 4-cliques (by assigning vectors corresponding to a regular tetrahedron) if and only if the original graph was 3-colorable.

However, the change in objective function from distances to *squared* distances causes a problem: while the regular tetrahedron is the unique optimal solution for $\mathrm{MC}_W^{\mathrm{L}}(K_4)$, in the case of $\mathrm{MC}_3(K_4)$, setting $v_1 = v_2 = (1, 0, 0), v_3 = v_4 = (-1, 0, 0)$ would also be optimal. So we replace each edge of the 4-cliques with triangles, which penalize the degenerate solution and forces the vectors assigned to the vertices of the 4-cliques toward regular tetrahedra.

▶ **Theorem 5.** MC_3 *is* NP*-complete.*

Proof of Theorem 5. Clearly MC_3 is in NP. To show hardness, we reduce 3-Coloring to MC_3. Given an instance $G = (V, E)$ on n vertices and m edges, we first replace every edge with a copy of K_4. That is, for each edge $ij \in E$, we add vertices q_{ij}, t_{ij} and add edges to form a 4-clique. Call the new graph $H' = (V', E')$. Then in H', we replace every edge with a copy of K_3. That is, for all $ij \in E'$ we add a vertex r_{ij} and edges ir_{ij} and jr_{ij}. Finally, we add a *sink* vertex t and for every edge $ij \in E$, in the original graph, add edge tt_{ij}. Call the resulting graph $H'' = (V'', E'')$. For future reference, let R denote a copy of K_4 with each edge replaced by a copy of K_3, which we may call a "tetrahedron with adjoined triangles".

We claim that if G has a proper 3-coloring then $\mathrm{MC}_3(H'') \geq m\,\mathrm{MC}_3(R) + m$. Conversely, we claim that if $\mathrm{MC}_3(H'') \geq m\,\mathrm{MC}_3(R) + m - \varepsilon$, for an $\varepsilon = \Omega\big(1/m^2\big)$ we will choose later, then there is a 3-coloring of G. Later, we will show $\mathrm{MC}_3(R) = 10 + 2\sqrt{3}$.

First, suppose G is 3-colorable. Let S consist of the following three unit vectors in $\mathbb{R}^3$:

$$\left(\sqrt{8/9}, 0, -1/3\right), \left(-\sqrt{2/9}, \sqrt{2/3}, -1/3\right), \left(-\sqrt{2/9}, -\sqrt{2/3}, -1/3\right).$$

Along with $(0, 0, 1)$, these are the four vertices of a regular tetrahedron inscribed in the unit sphere. Assign each one of the vectors of S to each color, such that every vertex in G is assigned a vector and no adjacent vertices have the same vector. We copy those vectors to the vertices of H', and for each vertex q_{ij} we assign the vector in S not assigned to i or j. We assign $(0, 0, 1)$ to each vertex t_{ij}. We copy these vectors to the vertices of H''. In H'' we assign $(0, 0, -1)$ to t. Finally, for each edge $ij \in E'$ and 3-clique i, j, r_{ij}, the only uncolored vertex is r_{ij}, and we assign the unique unit vector that is antiparallel to the sum of the vectors assigned to i and j.

Now we calculate the objective value that this assignment achieves. For vertex i, let $\hat{\imath}$ denote the assigned vector. For any edge ij in G, the vectors assigned to the associated K_4 gadget correspond to vertices of a regular tetrahedron, and it can be directly calculated that for any of the six edges ab, we have $\hat{a} \cdot \hat{b} = -1/3$. For edges tt_{ij}, we have $t \cdot t_{ij} = -1$. For edges ir_{ij}, given $\hat{\imath} \cdot \hat{\jmath} = -1/3$ and $\hat{r}_{ij} \propto -(\hat{\imath} + \hat{\jmath})$, the inner product is $\hat{\imath} \cdot \hat{r}_{ij} = \hat{\jmath} \cdot \hat{r}_{ij} = -1/\sqrt{3}$. In total, for each edge ij in G, the graph H'' has a gadget with six edges in a K_4 gadget, twelves edges in K_3 gadgets, and one edge incident on t. Plugging the inner products we calculated into the definition of MC_k gives a total objective value of $m(10 + 2\sqrt{3}) + m$.

Second, suppose there is a set of unit vectors $\hat{\imath}$ for $i \in V''$ such that $\mathrm{MC}_3(H'') \geq m\,\mathrm{MC}_3(R) + m - \varepsilon$. We color the vertices of H', comprised of the K_4 gadgets, as follows. Pick some color and assign it to all of the vertices t_{ab}. Then, choose any K_4 gadget i, j, k_{ij}, t_{ij}

in H' corresponding to an edge ij in the original graph. Let $\hat{\imath}, \hat{\jmath}, \hat{k}_{ij}, \hat{t}_{ij}$ be the vectors assigned to the vertices, respectively. Assign three colors to the uncolored vertices arbitrarily. Then, choose any K_4 gadget $i, \ell, k_{i\ell}, t_{i\ell}$ adjacent to the first. For each uncolored vertex in the second clique, consider its assigned vector, determine which of $\hat{\imath}, \hat{\jmath}, \hat{k}_{ij}$ it is closet to, and assign the same color. We continue coloring adjacent cliques in this way, comparing vectors to the original set $\hat{\imath}, \hat{\jmath}, \hat{k}_{ij}$, until the coloring is propagated to the entire graph.

Although we assigned four colors, one of them is used exclusively by the vertices t_{ab}. Since these vertices are not in G, if this is a proper 4-coloring of H', then it gives a proper 3-coloring of the subgraph G. So, in the remainder of the proof we will show that no two adjacent vertices in H' were assigned the same color.

The graph H'' is comprised of m edge-disjoint gadgets R as well as m edges tt_{ab}. In order for the total objective value to be greater than $m\,\mathrm{MC}_3(K_3) + m - \varepsilon$, each gadget R must contribute at least $\mathrm{MC}_3(R) - \varepsilon$ and each edge tt_{ab} must contribute at least $1 - \varepsilon$. As shown in Lemma 31, the sum of squared edge lengths of a "tetrahedron with adjoined triangles" is nearly maximized iff the lengths of the edges forming the tetrahedron are nearly regular. Specifically, for a gadget R to achieve at least $10 + 2\sqrt{3} - \varepsilon$, each side of the tetrahedron must be in the interval $4/\sqrt{6} \pm \mathrm{O}(\sqrt{\varepsilon})$.

Now consider any two adjacent vertices v, u in H'. The two vertices exist in some K_4 gadget in H', and the coloring procedure must have reached it after at most m rounds. At each step in the procedure, there is a colored clique and a successive adjacent clique. There are two copies of R in H'' which contain those two 4-cliques from H'. The gadgets share one vertex. Additionally, each gadget has a vertex adjacent to the sink t, whose assigned vectors must be at distance at least $2 - 2\varepsilon$ from $\hat{t}$. Since all of these are in the unit sphere, standard trigonometry shows those two assigned vectors must themselves be within $\mathrm{O}(\varepsilon)$ of each other. As verified in Lemma 30, considering two nearly regular tetrahedra which share one vertex and have a pair of similar vertices, the other vertices of the tetrahedra must each be within $\mathrm{O}(\sqrt{\varepsilon})$ of each other.

After m rounds of the coloring procedure, we conclude that $\hat{v}, \hat{u}$ are each at distance at most $m \times \mathrm{O}(\sqrt{\varepsilon})$ away from the same-colored vectors in the original set $\hat{\imath}, \hat{\jmath}, \hat{k}_{ij}$. Since all of the edges of the 4-clique gadgets must be at least $4/\sqrt{6} - \mathrm{O}(\sqrt{\varepsilon})$, for $\varepsilon = \Theta(1/m^2)$ sufficiently small the vectors $\hat{v}, \hat{u}$ cannot be near the same original vector. Therefore, they must been colored differently, as desired. ◀

By the correspondence described earlier in the section, we immediately have the desired bound on $\mathcal{S}$-prodLH.

▶ **Corollary 6.** Quantum Max-Cut *restricted to product states,* prodQMC*, is* NP*-complete, even when all terms are restricted to have positive unit weight.*

References

1. Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
2. Afonso S. Bandeira, Christopher Kennedy, and Amit Singer. Approximating the little Grothendieck problem over the orthogonal and unitary groups. *Mathematical programming*, 160:433–475, 2016. doi:10.1007/s10107-016-0993-7.
3. Alexander I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. *Discrete & Computational Geometry*, 13:189–202, 1995. doi:10.1007/BF02574037.
4. Fernando G.S.L. Brandão and Aram W. Harrow. Product-state approximations to quantum states. *Commun. Math. Phys.*, 342:47–80, 2016. doi:10.1007/s00220-016-2575-1.

5 Sergey Bravyi, David Gosset, Robert König, and Kristan Temme. Approximation algorithms for quantum many-body problems. *Journal of Mathematical Physics*, 60(3):032203, 2019. doi:10.1063/1.5085428.

6 Sergey Bravyi and Matthew Hastings. On complexity of the quantum Ising model. *Communications in Mathematical Physics*, 349(1):1–45, 2017. doi:10.1007/s00220-016-2787-4.

7 Jop Briët, Harry Buhrman, and Ben Toner. A generalized Grothendieck inequality and nonlocal correlations that require high entanglement. *Communications in mathematical physics*, 305(3):827–843, 2011. doi:10.1007/s00220-011-1280-3.

8 Jop Briët, Fernando Mário de Oliveira Filho, and Frank Vallentin. The positive semidefinite Grothendieck problem with rank constraint. In *International Colloquium on Automata, Languages, and Programming*, pages 31–42. Springer, 2010. doi:10.1007/978-3-642-14165-2_4.

9 Jop Briët, Oded Regev, and Rishi Saket. Tight hardness of the non-commutative Grothendieck problem. *Theory of Computing*, 13(15):1–24, 2017. doi:10.4086/toc.2017.v013a015.

10 Anne Broadbent and Alex Bredariol Grilo. QMA-hardness of consistency of local density matrices with applications to quantum zero-knowledge. *SIAM Journal on Computing*, 51(4):1400–1450, 2022. doi:10.1137/21M140729X.

11 Samuel Burer and Renato D.C. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical programming*, 95(2):329–357, 2003. doi:10.1007/s10107-002-0352-8.

12 Nadia Creignou. A dichotomy theorem for maximum generalized satisfiability problems. *Journal of Computer and System Sciences*, 51(3):511–522, 1995. doi:10.1006/jcss.1995.1087.

13 Nadia Creignou, Sanjeev Khanna, and Madhu Sudan. *Complexity classifications of Boolean constraint satisfaction problems*. SIAM, 2001.

14 Toby Cubitt and Ashley Montanaro. Complexity classification of local Hamiltonian problems. *SIAM Journal on Computing*, 45(2):268–316, 2016. doi:10.1137/140998287.

15 Toby S. Cubitt, Ashley Montanaro, and Stephen Piddock. Universal quantum Hamiltonians. *Proceedings of the National Academy of Sciences*, 115(38):9497–9502, August 2018. doi:10.1073/pnas.1804949115.

16 Sevag Gharibian, Yichen Huang, Zeph Landau, Seung Woo Shin, et al. Quantum Hamiltonian complexity. *Foundations and Trends in Theoretical Computer Science*, 10(3):159–282, 2015. doi:10.1561/0400000066.

17 Sevag Gharibian and Julia Kempe. Approximation algorithms for QMA-complete problems. *SIAM Journal on Computing*, 41(4):1028–1050, 2012. doi:10.1137/110842272.

18 Sevag Gharibian and Ojas Parekh. Almost Optimal Classical Approximation Algorithms for a Quantum Generalization of Max-Cut. In *Approximation, Randomization, and Combinatorial Optimization (APPROX/RANDOM 2019)*, volume 145 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 31:1–31:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.APPROX-RANDOM.2019.31.

19 Michel X. Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995. doi:10.1145/227683.227684.

20 Sean Hallgren, Eunou Lee, and Ojas Parekh. An approximation algorithm for the max-2-local Hamiltonian problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.APPROX/RANDOM.2020.59.

21 Ryszard Horodecki et al. Information-theoretic aspects of inseparability of mixed states. *Phys. Rev. A*, 54(3):1838–1843, 1996. doi:10.1103/PhysRevA.54.1838.

22 Yeongwoo Hwang, Joe Neeman, Ojas Parekh, Kevin Thompson, and John Wright. Unique games hardness of Quantum Max-Cut, and a conjectured vector-valued Borell's inequality. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1319–1384. SIAM, 2023. doi:10.1137/1.9781611977554.ch48.

23 Peter Jonsson. Boolean constraint satisfaction: complexity results for optimization problems with arbitrary weights. *Theoretical Computer Science*, 244(1-2):189–203, 2000. doi:10.1016/S0304-3975(98)00343-0.

24 Peter Jonsson, Mikael Klasson, and Andrei Krokhin. The approximability of three-valued Max CSP. *SIAM Journal on Computing*, 35(6):1329–1349, 2006. doi:10.1137/S009753970444644X.

25 Sanjeev Khanna, Madhu Sudan, and David P Williamson. A complete classification of the approximability of maximization problems derived from boolean constraint satisfaction. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 11–20, 1997. doi:10.1145/258533.258538.

26 Alexei Yu Kitaev, Alexander Shen, and Mikhail N Vyalyi. *Classical and quantum computation*. Number 47 in Graduate Studies in Mathematics. American Mathematical Soc., 2002.

27 Yi-Kai Liu. Consistency of local density matrices is QMA-complete. In Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 438–449, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. doi:10.1007/11830924_40.

28 L. Lovász. *Semidefinite Programs and Combinatorial Optimization*, pages 137–194. Springer New York, New York, NY, 2003. doi:10.1007/0-387-22444-0_6.

29 László Lovász. *Graphs and geometry*, volume 65. American Mathematical Soc., 2019.

30 Hiroshi Maehara. On the total edge-length of a tetrahedron. *The American Mathematical Monthly*, 108(10):967–969, 2001. URL: http://www.jstor.org/stable/2695418.

31 Ojas Parekh and Kevin Thompson. Beating random assignment for approximating quantum 2-local Hamiltonian problems. In *29th Annual European Symposium on Algorithms (ESA 2021)*, volume 204 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 74:1–74:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ESA.2021.74.

32 Ojas Parekh and Kevin Thompson. An optimal product-state approximation for 2-local quantum Hamiltonians with positive terms, 2022. arXiv:2206.08342v1.

33 Stephen Piddock and Ashley Montanaro. The complexity of antiferromagnetic interactions and 2D lattices. *Quantum Info. Comput.*, 17(7-8):636–672, 2017. doi:10.5555/3179553.3179559.

34 Stephen Piddock and Ashley Montanaro. Universal qudit Hamiltonians, 2018. arXiv:1802.07130.

35 Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, pages 216–226, 1978. doi:10.1145/800133.804350.

36 Johan Thapper and Stanislav Živnỳ. The complexity of finite-valued CSPs. *Journal of the ACM*, 63(4):1–33, 2016. doi:10.1145/2974019.

37 G. Thompson. Normal forms for skew-symmetric matrices and Hamiltonian systems with first integrals linear in momenta. *Proc. of the Amer. Math. Soc.*, 104(3):910–916, 1988. doi:10.2307/2046815.

38 John Watrous. *The theory of quantum information*. Cambridge university press, 2018.

39 Pawel Wocjan and Thomas Beth. The 2-local Hamiltonian problem encompasses NP. *International Journal of Quantum Information*, 1(03):349–357, 2003. doi:10.1142/S021974990300022X.

40 John Wright. Personal communication, 2022.

A Geometry Lemmas

First, we give a simple lemma regarding star graphs. Our goal is to show that inserting star gadgets into a graph forces maximal solutions to $\mathrm{MC}^{\mathrm{L}}_W$ to be close to the highest weighted axes.

▶ **Lemma 26.** *Consider a star graph S_K with center vertex v and K neighbors. Consider any $0 \leq \varepsilon \leq 1$ and any $W = \mathrm{diag}(1, w_2, w_3)$ with $1 \geq w_2 \geq w_3 \geq 0$ and $w_3 < 1$.*

Let Π *denote the projector onto the axes for which* $w_i = 1$ *(so the projector onto the* x*-axis or the* xy*-plane). Similarly, let* λ *be the largest* w_i *which is not equal to 1 (either* w_2 *or* w_3*). Let* $\widehat{\operatorname{sgn}}(\hat{v}) = \Pi v / \|\Pi v\|$.

If vectors are assigned to S_K *which achieve at least* $\mathrm{MC}^{\mathrm{L}}_{\mathrm{W}}(S_K) \geq K(1-\varepsilon)$*, and* $\hat{v}$ *is the vector assigned to* v*, then*

$$\|W\hat{v} - \widehat{\operatorname{sgn}}(\hat{v})\| \leq \delta \quad \textit{for} \quad \delta = 2\sqrt{\varepsilon}\sqrt{\frac{1+\lambda^2}{1-\lambda^2}}.$$

Proof. Suppose to the contrary that $\|W\hat{v} - \widehat{\operatorname{sgn}}(\hat{v})\| > \delta$. We know that

$$\begin{aligned}
&\|W\hat{v} - \widehat{\operatorname{sgn}}(\hat{v})\| = \sqrt{(\widehat{\operatorname{sgn}}(\hat{v})_x - \hat{v}_x)^2 + (\widehat{\operatorname{sgn}}(\hat{v})_y - w_2\hat{v}_y)^2 + w_3^2\hat{v}_z^2} \\
&= \sqrt{(\widehat{\operatorname{sgn}}(\hat{v})_x - \hat{v}_x)^2 + (\widehat{\operatorname{sgn}}(\hat{v})_y - w_2\hat{v}_y)^2 + w_3^2\left(1 - \hat{v}_x^2 - \hat{v}_y^2\right)} \\
&\leq \sqrt{\left(1 - \hat{v}_x^2 - \hat{v}_y^2\right) + \lambda^2(1 - \hat{v}_x^2 - \hat{v}_y^2)} = \sqrt{(1+\lambda^2)(1 - \hat{v}_x^2 - \hat{v}_y^2)},
\end{aligned}$$

where the inequality holds for $\hat{v}_x, \hat{v}_y \in [0,1]$. Combining, we have that $(1+\lambda^2)(1-\hat{v}_x^2-\hat{v}_y^2) > \delta^2$, which implies $\hat{v}_x^2 + \hat{v}_y^2 < 1 - \delta^2/(1+\lambda^2)$. Now we are ready to find that the maximum objective value earned on the star is less than half of

$$\begin{aligned}
&K(1 + \|W\hat{v}\|) = K + K\sqrt{\hat{v}_x^2 + \hat{v}_y^2 + \lambda^2\hat{v}_z^2} = K + K\sqrt{\hat{v}_x^2 + \hat{v}_y^2 + \lambda^2(1 - \hat{v}_x^2 - \hat{v}_y^2)} \\
&= K + K\sqrt{(\hat{v}_x^2 + \hat{v}_y^2)(1-\lambda^2) + \lambda^2} \leq K + K\sqrt{\left(1 - \frac{\delta^2}{1+\lambda^2}\right)(1-\lambda^2) + \lambda^2} \\
&= K + K\sqrt{1 - 4\varepsilon} \leq K + K(1 - 4\varepsilon) = K(2 - 4\varepsilon).
\end{aligned}$$

So, the objective value is less than $K(1-2\varepsilon)$, which is less than $K(1-\varepsilon)$, a contradiction. ◀

Next, we study the geometry of triangles. For a triangle ABC, let $\mathrm{L}(ABC)$ denote the sum of the edge lengths. It is straightforward to show that for a triangle inscribed in a circle of radius r, $\mathrm{L}(ABC) \leq 3\sqrt{3}r$, which is uniquely achieved by an equilateral triangle [30].

▶ **Lemma 27.** *Consider a triangle* ABC *inscribed in the unit circle and any* $0 \leq \varepsilon < 1$*. If* $\mathrm{L}(ABC) \geq 3\sqrt{3} - \varepsilon$*, then each edge length is in the interval* $\left[\sqrt{3} - 4\sqrt{\varepsilon}, \sqrt{3} + 4\sqrt{\varepsilon}\right]$.

Proof. For the sake of contradiction, suppose $\mathrm{L}(ABC) \geq 3\sqrt{3} - \varepsilon$ and that there exists an edge length outside the interval $\sqrt{3} \pm 4\sqrt{\varepsilon}$.

First suppose that some edge length is greater than $\sqrt{3} + 4\sqrt{\varepsilon}$. Because the maximum value of L is $3\sqrt{3}$, this implies at least one edge length is less than $\sqrt{3} - 2\sqrt{\varepsilon}$. So whether an edge length is above or below the interval, some edge length is less than $\sqrt{3} - 2\sqrt{\varepsilon}$.

Given $\mathrm{L}(ABC) \geq 3\sqrt{3} - \varepsilon$ and some edge length is less than $\sqrt{3} - 2\sqrt{\varepsilon}$, some edge length must be greater than $(3\sqrt{3} - \varepsilon - \sqrt{3} + 2\sqrt{\varepsilon})/2 \geq \sqrt{3} + \sqrt{\varepsilon}$, so there exists a pair of edge lengths whose difference is greater than $3\sqrt{\varepsilon}$.

We relabel the triangle so that $\|AC\| \geq \|AB\| \geq \|BC\|$ and $\|AC\| - \|BC\| > 3\sqrt{\varepsilon}$. We can rotate the unit circle as necessary so that ABC is of the form

$$A = (a, b) \quad B = (-a, b) \quad C = (c_x, c_y).$$

Define $C' = (0, -1)$. We will show $\mathrm{L}(ABC') - \mathrm{L}(ABC) > \varepsilon$, implying $\mathrm{L}(ABC)$ is less than $3\sqrt{3} - \varepsilon$, a contradiction.

Both $\mathrm{L}(ABC)$ and $\mathrm{L}(ABC')$ are sums of 3 edge lengths. The edge AB is shared and we have that $\|AC'\| = \|BC'\|$, so we are interested in $2\|AC'\| - \|BC\| - \|AC\|$. We may calculate

$$\|AC'\| = \sqrt{2+2b}, \quad \|BC\| = \sqrt{2+2bc_y+2ac_x}, \quad \|AC\| = \sqrt{2+2bc_y-2ac_x}.$$

From here, we may verify

$$4\|AC'\|^2 = (\|BC\| + \|AC\|)^2 + (\|BC\| - \|AC\|)^2,$$

and so

$$(2\|AC'\| + \|BC\| + \|AC\|)(2\|AC'\| - \|BC\| - \|AC\|) \geq (\|BC\| - \|AC\|)^2,$$

and

$$2\|AC'\| - \|BC\| - \|AC\| \geq (\|BC\| - \|AC\|)^2 (2\|AC'\| + \|BC\| + \|AC\|)^{-1}.$$

To lower bound the left-hand side, we use our lower bound on the difference between the edge lengths and the upper bound that all lengths in the unit circle are less than 2. We find

$$2\|AC'\| - \|BC\| - \|AC\| > \left(3\sqrt{\varepsilon}\right)^2 8^{-1} = 9\varepsilon/8.$$

Overall, we have found that $\mathrm{L}(ABC') - \mathrm{L}(ABC) > \varepsilon$. ◀

▶ **Lemma 28.** *Consider triangles ABC and ADE inscribed in the unit circle. If all edges of the triangle have lengths in the interval $\sqrt{3} \pm \delta$, then the points $\{B, C\}$ are each within $\mathrm{O}(\delta)$ distance of (different) points in $\{D, E\}$.*

Proof. Because these are vectors restricted to a unit circle, we can assume $A = (1, 0, 0)$. Given A is fixed, we can characterize the constraints on the other points as follows. B, C, D, E must lie in the intersection of the unit circle with a circular shell bounded by radii $\sqrt{3} - \delta$ and $\sqrt{3} + \delta$ centered at A. Given $\sqrt{3} + \delta < 2$, it is clear this intersection is two disjoint segments of the unit circle.

We wish to upper bound the distance between points in either of these regions. This is the distance between the ends of the regions, along a chord from the point at distance $\sqrt{3} - \delta$ to the point at distance $\sqrt{3} + \delta$.

The length of the chord can be bound as follows. Given a chord length d, the internal angle is $2\arcsin(d)$. Given an angle θ, the chord length is $2\sin(\theta/2)$. So, we can convert the known distances to angles, take the difference, and convert back to a distance: $2\sin(\arcsin((\sqrt{3} + \delta)/2) - \arcsin((\sqrt{3} - \delta)/2))$. This is equivalent to $2\sin\left(\arcsin(d/2) - \arcsin\left((d - 2\delta)/2\right)\right)$ for $d = \sqrt{3} + \delta$. Observing this function is increasing in d, we can upper bound the value by taking $d < \sqrt{3} + 0.1$. The Taylor series of $2\sin(\arcsin(\frac{\sqrt{3}+0.1}{2}) - \arcsin(\frac{\sqrt{3}+0.1-\delta}{2}))$ gives that this is at most 2.5δ. ◀

Next, we transition to considering tetrahedra. For a tetrahedron $ABCD$, let $\mathrm{L}(ABCD)$ denote the sum of the edge lengths. It is known that for a tetrahedron inscribed in a sphere of radius r, $\mathrm{L}(ABCD) \leq 4\sqrt{6}r$, and the maximum is uniquely achieved by a regular tetrahedron [30].

▶ **Lemma 29.** *Consider a tetrahedron $ABCD$ inscribed in the unit sphere and any $\varepsilon \geq 0$. If $\mathrm{L}(ABCD) \geq 4\sqrt{6} - \varepsilon$, then each edge length is in the interval $\left[\frac{4\sqrt{6}}{6} - \frac{50\sqrt{2\varepsilon}}{3}, \frac{4\sqrt{6}}{6} + \frac{50\sqrt{2\varepsilon}}{3}\right]$.*

Proof. For the sake of contradiction, suppose $\mathrm{L}(ABCD) \geq 4\sqrt{6} - \varepsilon$ and that there exists an edge length outside the given interval.

First consider the case that some edge length is greater than $\frac{4\sqrt{6}}{6} + \frac{50\sqrt{2\varepsilon}}{3}$. Because the maximum of L in the unit sphere is $4\sqrt{6}$, at least one of the other edge lengths must be less than $\frac{4\sqrt{6}}{6} - \frac{10\sqrt{2\varepsilon}}{3}$. So, whether an edge length is above or below the interval, in both cases some edge length is less than $\frac{4\sqrt{6}}{6} - \frac{10\sqrt{2\varepsilon}}{3}$.

If $\mathrm{L}(ABCD) \geq 4\sqrt{6} - \varepsilon$ and some edge length is less than $\frac{4\sqrt{6}}{6} - \frac{10\sqrt{2\varepsilon}}{3}$, then some edge length must be greater than $\left(4\sqrt{6} - \varepsilon - \frac{4\sqrt{6}}{6} + \frac{10\sqrt{2\varepsilon}}{3}\right)/5 \geq \frac{4\sqrt{6}}{6} + \frac{2\sqrt{2\varepsilon}}{3}$. So, there exists a pair of edges such that the difference in their lengths is greater than $\frac{12\sqrt{2\varepsilon}}{3} = 4\sqrt{2\varepsilon}$.

If two edge lengths e_1, e_2 in a tetrahedron differ by more than $4\sqrt{2\varepsilon}$, then there must exist a pair of adjacent edges which differ by more than $2\sqrt{2\varepsilon}$. This is because either the two edges are adjacent, or they are both adjacent to a third edge e_3 and that length must be closer to e_1 or to e_2, i.e. $\min\{\|e_3 - e_1\|, \|e_2 - e_1\|\} \leq \|e_1 - e_2\|/2$ We relabel the tetrahedron so that $\|AC\| - \|BC\| > 2\sqrt{2\varepsilon}$.

Next, without changing the value of L, we can rotate $ABCD$ such that A, B are of the form $A = (a, -b, 0) \quad B = (-a, -b, 0)$. Let $C = (c_x, c_y, c_z) \quad D = (d_x, d_y, d_z)$ and $C' = (0, c_y, \sqrt{1 - c_y^2}) \quad D' = (0, d_y, \sqrt{1 - d_y^2})$, We will show $\mathrm{L}(ABC'D') - \mathrm{L}(ABCD) > \varepsilon$, implying $\mathrm{L}(ABCD)$ must in fact be less than $4\sqrt{6} - \varepsilon$.

Both $\mathrm{L}(ABCD)$ and $\mathrm{L}(ABC'D')$ are sums of 6 edge lengths. We can ignore AB. We can directly verify that $\|C'D'\| - \|CD\|$ is

$$\sqrt{(c_y - d_y)^2 + (c_x - d_x)^2 + (c_z - d_z)^2 + 2c_x d_x + 2c_z d_z + 2\sqrt{c_x^2 + c_z^2}\sqrt{d_x + d_z^2}} - \sqrt{(c_x - d_x)^2 + (c_y - d_y)^2 + (c_z - d_z)^2},$$

and so is clearly non-negative. We may calculate

$$\|AC'\| = \|BC'\| = \sqrt{2 + 2bc_y}$$
$$\|BC\| = \sqrt{2 + 2bc_y + 2ac_x}$$
$$\|AC\| = \sqrt{2 + 2bc_y - 2ac_x}$$

and similar expressions for $\|AD'\|, \|BD'\|, \|BD\|, \|AD\|$.

Just as in the proof of Lemma 27, we may verify

$$\begin{aligned} 2\|AC'\| - \|BC\| - \|AC\| &= (\|BC\| - \|AC\|)^2 \left(2\|AC'\| + \|BC\| + \|AC\|\right)^{-1} \\ &> \left(2\sqrt{2\varepsilon}\right)^2 8^{-1} \\ &= \varepsilon. \end{aligned}$$

Similarly, $2\|AD'\| - \|AD\| - \|BD\|$ is twice the difference of the quadratic mean and arithmetic mean of $\|AD\|, \|BD\|$, which is non-negative by the generalized mean inequality.

Overall, we have found that $\mathrm{L}(ABC'D') - \mathrm{L}(ABCD) > \varepsilon$. ◀

▶ **Lemma 30.** *Consider tetrahedra $ABCD$ and $AEFG$ inscribed in the unit sphere with $\|B - E\| \leq \varepsilon$. If all edges of the tetrahedra have lengths in the interval $\frac{4\sqrt{6}}{6} \pm \delta$, then the points $\{C, D\}$ are each within $\mathrm{O}(\delta + \varepsilon)$ distance of (different) points in $\{F, G\}$.*

Proof. Because this is the unit sphere, we can rotate so that A, B are in the xy-plane and equidistant from the y-axis.

Given A, B, E are fixed, we can characterize the constraints on the other points as follows. Let R_1 denote the region bounded by a spherical shell centered at A with radii $\frac{4\sqrt{6}}{6} - \delta$ and $\frac{4\sqrt{6}}{6} + \delta$. Similarly, let R_2 denote a region of the same size and shape but centered at B. The points C, D must exist in the intersection of R_1, R_2, and the unit sphere. The constraints imply the intersection is two disjoint sectors of the unit sphere opposite A and B, call them S_1, S_2. Similarly, the points F, G must exist within ε of S_1 and S_2.

We wish to show that any two points in S_1 must be close, and similarly for S_2. Consider any two points P_1, P_2 in S_1. Fixing P_1 and letting $d(P_2) = \|P_1 - P_2\|$, the function d is convex in P_2 over S_1. Therefore, the extremes of the distance $\|P_1 - P_2\|$ will occur at the four extremal points of S_1, identified by their distances $\frac{4\sqrt{6}}{6} \pm \delta$ from A and B.

If P_1 is at $\frac{4\sqrt{6}}{6} + \delta$ from A and from B, then $d(P_2)$ is maximized with P_2 at $\frac{4\sqrt{6}}{6} - \delta$ from A and B. From point P_1, moving 2δ towards A and then 2δ towards B is an upper bound on the distance to P_2, so $\|P_1 - P_2\| \leq 4\delta$.

The other extreme may occur with $\frac{4\sqrt{6}}{6} + \delta$ and $\frac{4\sqrt{6}}{6} - \delta$ as the distances from P_1 to A and B, respectively, and as the distances from P_2 to B and A, respectively. Recalling that we were able to assume A, B are in the xy-plane and symmetric about the z-axis, it is clear that P_1, P_2 must be symmetric such that $P_1 = (0, P_y, P_z), P_2 = (0, -P_y, P_z)$. This implies A, B, P_1, P_2 are coplanar, and so ABP_1P_2 is a cyclic quadrilateral inscribed in a circular cross-section of the unit sphere. We know the lengths of three sides and the diagonals of this quadrilateral and want to know the fourth side, $\|P_1P_2\|$. Ptolmey's Theorem for cyclic quadrilaterals gives us that

$$\|AP_1\| \times \|BP_2\| = \|AB\| \times \|P_1P_2\| + \|AP_2\| \times \|BP_1\|.$$

Simplifying, we find

$$\|P_1P_2\| \leq \frac{\left(4\sqrt{6} + \delta\right)^2 - \left(4\sqrt{6} - \delta\right)^2}{4\sqrt{6} - \delta} = \frac{8\sqrt{6}\delta}{2\sqrt{6} - 3\delta} \leq \frac{8\sqrt{6}\delta}{2\sqrt{6}} = 4\delta.$$

Overall, we conclude that given any two points in S_1, the distance between the points is at most $\mathrm{O}(\delta)$, and similarly for S_2.

Finally, we return to the shapes $ABCD$ and $AEFG$. Since the distance $\|CD\|$ must be in $\frac{4\sqrt{6}}{6} \pm \delta$, one point must be in S_1 and one point in S_2. Similarly, one of F, G must be near S_1 and one near S_2. So arbitrarily, C, F are within ε of S_1 and D, G are within ε of S_2. We can conclude $\|C - F\|$ and $\|D - G\|$ are at most $\mathrm{O}(\delta + \varepsilon)$, as desired. ◀

We will use the name "tetrahedron with adjoined triangles" to refer to a three-dimensional tetrahedron such that for each edge there is an additional point with which the ends of the edge are joined to form a triangle. In total, there are 10 vertices and 12 edges. The shape will be assumed to be inscribed in the unit sphere, so a vector describing any vertex of the tetrahedron or a triangle is a unit vector. For some instance R of this shape with assigned vertex locations, let $\mathrm{L}(R)$ denote the sum of the *squared* edge lengths. The following lemma says that for $\mathrm{L}(R)$ to be maximized, the edges of the tetrahedron must approximate a regular tetrahedron.

▶ **Lemma 31.** *Given a tetrahedron with adjoined triangles R, $\mathrm{L}(R) \leq 10 + 2\sqrt{3}$. Furthermore, If any edge of the tetrahedron has a length outside of the interval $[\frac{4}{\sqrt{6}} - \varepsilon, \frac{4}{\sqrt{6}} + \varepsilon]$, then*

$$\mathrm{L}(R) < 10 + 2\sqrt{3} - \Omega\left(\varepsilon^2\right).$$

Proof. Suppose the vertices of the tetrahedron are V_1, V_2, V_3, V_4, and the additional vertex forming a triangle with edge V_iV_j is v_{ij}. Let s_{ij} denote the length of edge ij.

Note that the optimal location (which maximizes L) of a vertex v_{ij} is entirely determined by the locations of V_i and V_j, since it has no other neighbors. Given V_i, V_j, the point which maximizes the squared distances is proportional to $-(V_i + V_j)$. Moreover, because we are in the unit sphere, the maximum objective achieved by a triangle $V_iV_jv_{ij}$ when V_i, V_j are fixed is entirely determined by the length of V_iV_j alone; their exact coordinates do not matter, due to rotational symmetry. Combining these two facts and applying trigonometry, we may calculate that given an edge length of the tetrahedron s_{ij}, the maximum objective value of the triangle $V_iV_jv_{ij}$ is

$$\mathrm{t}(s_{ij}) = 1 + \frac{s_{ij}^2}{4} + \sqrt{1 - \frac{s_{ij}^2}{4}}.$$

Now, because R is the disjoint union of the six triangles which each have one of the tetrahedron's edges as a side, we can express

$$\mathrm{L}(R) = \sum_{i<j}^{4} t(s_{ij}).$$

A regular tetrahedron inscribed in the unit sphere has edge lengths equal to $4/\sqrt{6}$. So, if R is formed by a regular tetrahedron and the ideal points for the triangles, then $\mathrm{L}(R) = 6\,\mathrm{t}(4/\sqrt{6}) = 10 + 2\sqrt{3}$.

Now, we suppose the tetrahedron is not regular and some side length is not in the interval $\frac{4}{\sqrt{6}} \pm \varepsilon$. As noted in prior lemmas, it is known that the maximum sum of the (unsquared) edge-lengths of a tetrahedron is $4\sqrt{6}$ [30], so we have $\sum_{i<j} s_{ij} \leq 4\sqrt{6}$. If any edge length is greater than $\frac{4}{\sqrt{6}} + \varepsilon$, then the sum of the other five sides must be greater $5 \times \frac{4}{\sqrt{6}} - \varepsilon$ (otherwise, $\mathrm{L}(R)$ could be trivially increased, meaning our assumption gives a lower bound on the difference), so some side is less than $\frac{4}{\sqrt{6}} - \varepsilon/5$. So whether some side is above or below the interval, there exists a side $s_{ij} < \frac{4}{\sqrt{6}} - \varepsilon/5$.

Analyzing $\mathrm{t}(s)$, we see it is increasing and concave down on the interval $[1.22, 2)$. The maximum occurs at $\mathrm{t}(\sqrt{3})$. We can assume all edge lengths are at least 1.22 and in the concave-down region, as otherwise the total edge length is less than $\mathrm{t}(1.22)+5\,\mathrm{t}(\sqrt{3}) = 13.4145$, which is bounded away from the optimum $10 + 2\sqrt{3} \approx 13.464$. Therefore, the sum of t over the edges of a regular tetrahedron are greater than over a non-regular tetrahedron:

$$6\,\mathrm{t}\left(\frac{4}{\sqrt{6}}\right) > \sum_{i<j}^{4} \mathrm{t}(s_{ij}).$$

In particular,

$$6\,\mathrm{t}\left(\frac{4}{\sqrt{6}}\right) - \sum_{i<j}^{4} \mathrm{t}(s_{ij}) \geq 2\,\mathrm{t}\left(\frac{4}{\sqrt{6}}\right) - \mathrm{t}\left(\frac{4}{\sqrt{6}} - \frac{\varepsilon}{5}\right) - \mathrm{t}\left(\frac{4}{\sqrt{6}} + \frac{\varepsilon}{5}\right).$$

After extensive calculation, the difference of the sum over a regular tetrahedron (the LHS) minus the sum over the non-regular tetrahedron (the RHS) is $\Omega(\varepsilon^2)$. ◀