

Two-local qubit Hamiltonians: when are they stoquastic?

Joel Klassen¹ and Barbara M. Terhal^{1,2}

¹QuTech, Delft University of Technology, P.O. Box 5046, 2600 GA Delft, The Netherlands

²Institute for Theoretical Nanoelectronics, Forschungszentrum Juelich, D-52425 Juelich, Germany

May 1, 2019

We examine the problem of determining if a 2-local Hamiltonian is stoquastic by local basis changes. We analyze this problem for two-qubit Hamiltonians, presenting some basic tools and giving a concrete example where using unitaries beyond Clifford rotations is required in order to decide stoquasticity. We report on simple results for n -qubit Hamiltonians with identical 2-local terms on bipartite graphs. Our most significant result is that we give an efficient algorithm to determine whether an arbitrary n -qubit XYZ Heisenberg Hamiltonian is stoquastic by local basis changes.

1 Introduction and Motivation

The notion of stoquastic Hamiltonians was introduced in [1] in the context of quantum complexity theory. The definition of stoquastic Hamiltonians aims to capture Hamiltonians which do not suffer from the sign problem: the definition ensures that the ground state of the Hamiltonian has nonnegative amplitudes in some local basis while the matrix $\exp(-H/kT)$ is an entry-wise nonnegative matrix for any temperature T in this basis. Stoquastic Hamiltonians are quite ubiquitous. Any (mechanical) Hamiltonian defined with conjugate variables $[\hat{x}_i, \hat{p}_j] = i\hbar\delta_{ij}$ of the form

$$H = \frac{1}{2}\vec{p}^T C^{-1}\vec{p} + U(\vec{x}). \quad (1)$$

with invertible (mass) matrix $C > 0$, is stoquastic, since the kinetic term is off-diagonal in the position \vec{x} -basis. The paradigm of circuit-QED produces such Hamiltonians by the canonical quantization of an electric circuit represented by its Lagrangian density.

For the subclass of stoquastic Hamiltonians the problem of determining the lowest eigenvalue might be easier than for general Hamiltonians. In complexity terms the stoquastic lowest-eigenvalue problem is StoqMA-complete while it is QMA-complete for general Hamiltonians.

This difference in complexity goes hand-in-hand with the existence of heuristic computational methods for de-

termining the lowest eigenvalue of stoquastic Hamiltonians using Monte Carlo techniques [2, 3]. These methods can be viewed as stochastic realizations of the power method, namely a stochastic efficient implementation of the repeated application of the nonnegative matrix $\exp(-H)$. However, there is no proof that these heuristic methods are as efficient as quantum phase estimation for estimating ground-state energies (some results are obtained in [4]).

Beyond the lowest eigenvalue problem, one can also consider the complexity of evaluating the partition function $Z = \text{Tr}[\exp(-H/kT)]$ of a stoquastic Hamiltonian. The nonnegativity of $\exp(-H/kT)$ directly ensures that one can rewrite Z as the partition function of a classical model in one dimension higher, which can then be sampled using classical stochastic techniques. Again, rigorous results on this well-known path-integral quantum Monte Carlo method (see e.g. [5]) are sparse: [6] shows how to estimate the partition function with $kT = \Omega(1/\log n)$ for 1D stoquastic Hamiltonians using a rigorous analysis of the path-integral Monte Carlo method. In [7] the authors gave a fully-polynomial randomized approximating scheme for estimating the partition function of stoquastic XY models (in a transverse magnetic field).

Stoquastic Hamiltonians play an important role in adiabatic computation [8]. It has been proven that there is no intrinsic quantum speed-up in stoquastic adiabatic computation which uses frustration-free stoquastic Hamiltonians: a classical polynomial-time algorithm for simulating such adiabatic computation was described in [9]. The quantum power of the well-known quantum annealing method which uses the transverse field Ising model is still an intriguing open question, given the road-block cases that have been thrown up for the path-integral Quantum Monte Carlo method in [10] as well as the diffusion Monte Carlo method in [11]. However, given that stoquastic adiabatic computation such as quantum annealing is amenable to heuristic Monte Carlo methods, research has also moved into the direction of finding ways of engineering non-stoquastic Hamiltonians (see e.g. [12], [13], [14]).

If it is true that questions pertaining to stoquastic

Hamiltonians are easier to answer than those pertaining to general Hamiltonians, then it is clearly important to be able to recognize *which Hamiltonians are stoquastic*.

Work aimed at recognizing where and when a sign problem does not exist, is not new, but few systematic approaches to this problem exist. Most of the work has been focused on understanding the sign problem for specific classes of Hamiltonians of physical interest, for example Fermi-Hubbard models [15]. An example of a recent exploration relating stoquasticity to time-reversal for fermionic systems is [16] (see [17] for a recent review). [18] has attempted to generalize the so-called Peierls-Marshall sign rules for isotropic spin-1/2 Heisenberg models to more general XYZ Hamiltonians, falling short of formulating any algorithmic approach. More recently, [19] found the first negative algorithmic results. They showed that for 3-local qubit Hamiltonians determining whether single-qubit Clifford rotations can indeed "cure the sign problem" is NP-complete (and similarly for orthogonal rotations).

In this paper we report on our first results in this direction (see e.g. [20]): they are summarized in Section 1.2. Our main result is an efficient algorithm to determine whether an XYZ Hamiltonian is stoquastic.

The definition of stoquasticity, see Definition 2, is primarily motivated by computational efficiency. A local change of basis can be efficiently executed in that the description of the k -local Hamiltonian in this new basis can be easily obtained. It is however clear that the subclass of single-qubit unitary transformations are only the simplest example of a mapping or encoding of one Hamiltonian H to another Hamiltonian H_{sim} as has been formalized in [21]. Such a more general mapping could allow: constant-depth circuits; ancilla qubits; perturbative gadgets and approximations such that only $\lambda(H) \approx \lambda(H_{\text{sim}})$ or $Z_H \approx cZ_{H_{\text{sim}}}$ for some known constant c . The goal of such an encoding is to map the original Hamiltonian H onto a simulator Hamiltonian H_{sim} where H_{sim} is stoquastic and its low-energy dynamics effectively capture those of H . The upshot is that the encoded Hamiltonian H_{sim} is directly computationally useful. If for a class of Hamiltonians H , simulator Hamiltonians H_{sim} can be found, then the ground-state energy problem or the adiabatic computation of H is amenable to quantum Monte Carlo techniques and its complexity is in StoqMA. Additionally, the definition of encoding employed in [21] ensures that the partition function of the simulator Hamiltonian can be used to estimate that of the original Hamiltonian. An example is the mapping of any k -local Hamiltonian onto a real $k + 1$ -local Hamiltonian [21].

Surprisingly, there are no known no-go's for the use of perturbative gadgets to map seemingly non-stoquastic

Hamiltonians onto stoquastic simulator Hamiltonians. Even though mapping all Hamiltonians onto stoquastic Hamiltonians seems to be excluded from a computational complexity perspective, it is possible that there are Hamiltonians which are not stoquastic by a local basis change but for which a stoquastic simulator Hamiltonian exists.

1.1 Definitions

To make contact with some of the mathematics literature, we will use some standard terminology:

Definition 1 (Symmetric Z -matrix). *A matrix is a symmetric Z -matrix if all of its matrix elements are real, it is symmetric, and its off-diagonal elements are non-positive.*

Note that for a symmetric Z -matrix H the matrix exponential $\exp(-H/kT)$ is entrywise non-negative for all $kT \geq 0$ ¹.

We reproduce the definition of a Hamiltonian being termwise stoquastic from [1]². We will refer to it as stoquastic for simplicity.

Definition 2 (Stoquastic [1]). *A k -local Hamiltonian H is stoquastic if there exists a local basis \mathcal{B} and a decomposition $H = \sum_i D_i$ such that each D_i is a symmetric Z -matrix in the basis \mathcal{B} and acts non-trivially on at most k qudits.*

Throughout the text the terms local basis, local change of basis, local unitary, or local Clifford, are all 1-local in the sense that they preserve the 1-local tensor product structure of the Hilbert space.

Given the discussion in the introduction and following [19], we also introduce the following definition (loosely stated):

Definition 3. *A family of Hamiltonians $\{H\}$ on n qudits is said to be **computationally stoquastic** if there is a polynomial-time algorithm for (approximately) mapping $\{H\}$ onto a family of stoquastic simulator Hamiltonians $\{H_{\text{sim}}\}$. The description of H_{sim} in terms of its entries should be efficiently given.*

¹It is worthwhile noting that the entrywise non-negativity of the matrix $\exp(-H/kT)$ for a given temperature T does not imply that H is a symmetric Z -matrix. It is simple to generate numerical 4×4 cases where $G = A^T A$ is a non-negative matrix by taking A as a non-negative matrix, but $H = -\log G$ is not a symmetric Z -matrix. Thus when H is a symmetric Z -matrix, the non-negativity of $\exp(-H/kT)$ is guaranteed for all values of kT , while in other cases $\exp(-H/kT)$ might be non-negative only for sufficiently small values of kT .

²[1] points out that for 2-local Hamiltonians one can prove that termwise stoquastic is the same as stoquastic. This is however not true for $k > 2$ -local Hamiltonians generally.

Another aspect of stoquasticity is the complexity of determining whether H is (computationally) stoquastic. The question is only clearly formulated when we restrict ourselves to specific mappings. For example, one can ask: is there an efficient algorithm for determining whether a 2-local Hamiltonian on n qubits is stoquastic in a basis obtained by single-qubit Clifford rotations. In this light, it was shown in [19] that for 3-local qubit Hamiltonians determining whether single-qubit Clifford rotations can indeed "cure the sign problem" is NP-complete (similarly for orthogonal rotations). For two-local Hamiltonians the question how hard it is to decide whether a Hamiltonian is stoquastic by local basis changes is entirely open. The next section reviews our main results on this matter.

1.2 Summary of Results

In Section 2 we consider, as a warm-up, the two-qubit Hamiltonian. We introduce some concepts which we will make use of in the multi-qubit setting, and we analyze under what conditions a two-qubit Hamiltonian can be transformed into a Z -matrix by a local basis change. This turns out to be surprisingly complex. In Theorem 7 we show that one can determine whether a 2-qubit Hamiltonian is real under local basis changes by inspecting a subset of non-local invariants. This thus gives a simple necessary condition to check if one wants to decide whether a two-qubit Hamiltonian is stoquastic. In Section 2.2 we establish that basis changes beyond single-qubit Clifford rotations are, perhaps unsurprisingly, strictly stronger than Clifford rotations. For the general two-qubit case we show that a two-qubit Hamiltonian H is stoquastic iff one can find a solution to a set of degree-4 polynomials in two variables (the polynomials are quadratic in each variable), with details in Appendix C.

In Section 3 we move to the problem of determining if an n -qubit 2-local Hamiltonian is stoquastic. The challenge in this case is twofold. Recalling Definition 2 one must find an appropriate *decomposition* into 2-qubit terms D_i . Additionally one must find a consistent assignment of rotations to bring the two-qubit terms into Z -matrix form. The interplay of these two tasks seems to make the problem quite difficult. As such, we do not try to solve the most general case. Instead we consider tractable cases of this problem. The first case we consider is deciding if a multi-qubit 2-local Hamiltonian admits a decomposition into 2-qubit terms each of which are Z -matrices. This case is tractable because we are ignoring any local unitary transformations. A more general case of this problem is also discussed in [19], but here we specifically discuss the two-local case. Next we consider a family of natural Hamiltonians, which

have identical two-local terms on a bipartite interaction graph. We prove that in this case the problem of finding a set of one-local unitary rotations in concert with finding an appropriate decomposition is tractable, due to the special structure of the Hamiltonian.

The final case we consider is a subclass of Hamiltonians, called XYZ Heisenberg Hamiltonians, with an arbitrary interaction graph. Here we present an efficient algorithm to decide if a given Hamiltonian of this type is stoquastic, and to construct the unitary rotation which transforms the Hamiltonian into a Z -matrix. This constitutes the most significant result of this paper. For such Hamiltonians, finding a decomposition is trivial, because all 1-local terms are zero. This leaves us only needing to consider how to apply one-local unitary rotations. The Hamiltonian of the XYZ Heisenberg model for n qubits is $H = \sum_{u,v} H_{uv}$ for $u, v = 1, \dots, n$ where

$$H_{uv} = a_{XX}^{uv} X_u X_v + a_{YY}^{uv} Y_u Y_v + a_{ZZ}^{uv} Z_u Z_v. \quad (2)$$

Here the coefficients a_{PP}^{uv} , which can be different for each uv , are specified with some k bits. The theorem that we prove in Section 3.3 is:

Theorem 4. *There is an efficient algorithm that runs in time $O(n^3)$ to decide whether an n -qubit XYZ Heisenberg model Hamiltonian is stoquastic. The algorithm finds the local basis change such that the resulting H is a symmetric Z -matrix or decides that it does not exist.*

This result relies critically on Lemmas 22 and 21 (proven in Appendix D) which show that when an XYZ Heisenberg model is stoquastic, it can always be transformed into a Z -matrix by single-qubit Clifford rotations. Clifford rotations are significant here because single-qubit Clifford rotations are a finite group, and thus dramatically simplify the analysis.

Going beyond the XYZ Heisenberg Hamiltonian, one can ask about the complexity of deciding stoquasticity by local Clifford rotations. Even though this is not the most general class of local basis changes, see for example Equation (6) in Section 2, our algorithm for the XYZ Hamiltonian suggests progress could most readily be made in this direction. We are able to prove that determining whether a Hamiltonian is real by local Clifford rotations is NP-complete:

Theorem 5. *Deciding if a 2-local n -qubit Hamiltonian is real under single-qubit Clifford operations is NP-complete: there exists a subclass of 2-local n -qubit Hamiltonians for which deciding if they are real under single-qubit Clifford operations is equivalent to solving the restricted exact covering by 3-sets problem (RXC3), which is known to be NP-complete.*

This is significant because making a Hamiltonian real is a necessary condition for transforming it into a Z -matrix. Note that this result does not imply that deciding if a Hamiltonian can be transformed into a Z -matrix by single-qubit Clifford rotations or even general local basis changes for 2-local Hamiltonians is NP-complete. We give the proof of Theorem 5 in Appendix E since it is not central to the paper.

2 Warm-Up: Two Qubit Hamiltonians

In this section we establish some basic facts about two-qubit Hamiltonians which we will make use of throughout the paper. We begin by establishing under what conditions a two-qubit Hamiltonian is a Z -matrix. We then present a convenient representation of two-qubit Hamiltonians which we will make heavy use of throughout the paper. We briefly discuss the nature of Clifford transformations in this representation, and give an example where Clifford transformations are not sufficient to decide if a Hamiltonian is stoquastic. Next we present a set of invariants which completely indicate if a Hamiltonian can be made real under local unitary rotations, a necessary condition for a Hamiltonian to be stoquastic. Finally we explain how the problem of deciding if a two-qubit Hamiltonian is stoquastic can be reduced to deciding if a set of degree-4 polynomial inequalities admit a solution, which suggests that even in the two-qubit case deciding if a Hamiltonian is stoquastic is a non-trivial task.

Consider a general two qubit Hamiltonian

$$H = \sum_{ij=I,X,Y,Z} a_{ij} \sigma_i \otimes \sigma_j, \quad a_{ij} \in \mathbb{R}, \quad (3)$$

where we may take $a_{II} = 0$ without loss of generality. It is simple to show that

Proposition 6. *H is a symmetric Z -matrix if and only if $a_{IY} = a_{YI} = a_{XY} = a_{YX} = a_{ZY} = a_{YZ} = 0$ (the matrix is real) and $a_{XX} \leq -|a_{YY}|$ and $a_{IX} \leq -|a_{ZX}|$, $a_{XI} \leq -|a_{XZ}|$ (the matrix has non-positive off-diagonal elements).*

Proof Sketch. Visually, the positive off-diagonal contributions of XX and YY terms need to be removed leading to $a_{XX} \leq -|a_{YY}|$. Similarly, the positive off-diagonal contributions of IX and ZX terms need to be removed, leading to $a_{IX} \leq -|a_{ZX}|$ (and similarly $a_{XI} \leq -|a_{XZ}|$). \square

The set of traceless symmetric Z -matrices forms a polyhedral cone \mathcal{C}_2 , that is, a cone supported by a finite number of hyperplanes. In the case of 4×4 matrices, each matrix is determined by 9 parameters (three of

them real and 6 of them nonnegative). Viewed as a polyhedral cone, \mathcal{C}_2 has 12 extremal vectors [22]

$$\begin{aligned} \mathbf{o}_1 &= -X \otimes X - Y \otimes Y & \mathbf{o}_2 &= -X \otimes X + Y \otimes Y \\ \mathbf{o}_3 &= -|0\rangle\langle 0| \otimes X & \mathbf{o}_4 &= -|1\rangle\langle 1| \otimes X \\ \mathbf{o}_5 &= -X \otimes |0\rangle\langle 0| & \mathbf{o}_6 &= -X \otimes |1\rangle\langle 1| \\ \mathbf{d}_1^\pm &= \pm Z \otimes I & \mathbf{d}_2^\pm &= \pm I \otimes Z \\ \mathbf{d}_3^\pm &= \pm Z \otimes Z. \end{aligned}$$

Appendix A discusses some of this structure and its generalization to two qudits.

2.1 A Convenient Representation of 2-qubit Hamiltonians

Two-qubit Hamiltonians are conveniently parametrized by a 3×3 real matrix and two three-dimensional real vectors:

$$\beta = \begin{pmatrix} a_{XX} & a_{XY} & a_{XZ} \\ a_{YX} & a_{YY} & a_{YZ} \\ a_{ZX} & a_{ZY} & a_{ZZ} \end{pmatrix} \quad S = \begin{pmatrix} a_{XI} \\ a_{YI} \\ a_{ZI} \end{pmatrix} \quad P = \begin{pmatrix} a_{IX} \\ a_{IY} \\ a_{IZ} \end{pmatrix}. \quad (4)$$

A basis transformation $H \rightarrow (U_1 \otimes U_2)H(U_1 \otimes U_2)^\dagger$ corresponds to a pair of $SO(3)$ rotations O_1 and O_2 acting as

$$\beta \rightarrow O_1 \beta O_2^T \quad S \rightarrow O_1 S \quad P \rightarrow O_2 P. \quad (5)$$

This correspondence between $SU(2)$ and $SO(3)$ is well known in the context of one qubit, and holds also in the two qubit case,

$$\begin{aligned} H' &= (U_1 \otimes U_2)H(U_1 \otimes U_2)^\dagger \\ H' &= \sum_{i \in XYZ} \alpha'_{iI} \sigma_i \otimes I + \sum_{j \in XYZ} \alpha'_{Ij} I \otimes \sigma_j \\ &+ \sum_{ij \in X,Y,Z} \alpha'_{ij} \sigma_i \otimes \sigma_j \\ \alpha'_{iI} &= \frac{1}{4} \text{Tr}[H' \sigma_i \otimes I] = \sum_m [O_1]_{im} \alpha_m I \\ \alpha'_{Ij} &= \frac{1}{4} \text{Tr}[H' I \otimes \sigma_j] = \sum_n \alpha_{In} [O_2]_{jn} \\ \alpha'_{ij} &= \frac{1}{4} \text{Tr}[H' \sigma_i \otimes \sigma_j] = \sum_{mn} [O_1]_{im} \alpha_{mn} [O_2]_{jn} \\ [O_x]_{ij} &= \frac{1}{2} \text{Tr}[\sigma_i U_x \sigma_j U_x^\dagger]. \end{aligned}$$

This implies that for 2-local qubit Hamiltonians it suffices to consider pairs of local $SO(3)$ rotations and their effect on β, P, S . It is particularly useful to consider β in diagonal form. Every real square matrix admits a singular value decomposition $\beta = O_L \Sigma O_R^T$ with

$O_L, O_R \in O(3)$, $\Sigma \geq 0$. If we do not enforce that Σ be positive-semidefinite, then $O_L, O_R \in SO(3)$. Hence there always exists local unitary rotations acting on H which put the β matrix into diagonal form.

2.2 Clifford Rotations and Beyond

The single-qubit Clifford rotations play a special role as local basis changes. The action of single-qubit Clifford rotations on the Pauli-matrices form a discrete subgroup of $SO(3)$ representing the symmetries of the cube [23]. The Clifford rotations realize any permutation of the Paulis (the group S_3) as well as sign-flips $\sigma_i \rightarrow \pm\sigma_i$ (with determinant 1). As we shall see later in the paper, there are some cases where it is sufficient to consider Clifford rotations alone to decide if a Hamiltonian is stoquastic. It is important to establish that this is not always the case to get some intuition about the problem. Here we will present a two-qubit Hamiltonian for which it does not suffice to consider only the single-qubit Clifford rotations. Consider the two-qubit Hamiltonian

$$H = a_Z(Z_1 + Z_2) - a_X(X_1 + X_2) + a_{XX}X_1X_2 + Z_1Z_2. \quad (6)$$

with $1 \geq a_{XX} \geq 0$ and $1 > a_Z > 0$, $a_X > 0$. It is easy to argue that no single-qubit Clifford basis change can make this Hamiltonian into a symmetric Z -matrix. Since H has to be real, no permutations to Y -components are allowed. A sign-flip which makes the XX term negative must also flip the sign of one of the single-qubit X terms. Interchanging X and Z for both qubits leads to the same problem for the ZZ and Z terms. Applying a sign flip on the Z of one of the qubits, followed by interchanging X and Z on that same qubit leads to the requirement that $-a_Z \leq -1$ and $-a_X \leq -|a_{XX}|$. This is not satisfiable as long as $0 < a_Z < 1$.

However, we could create off-diagonal XZ and ZX terms by single-qubit non-Clifford rotations. For example, one could apply $U_1 = e^{i\pi Y/8} e^{i\pi X/2}$ on the first qubit and $U_2 = e^{-i\pi Y/8} e^{i\pi X/2}$ on the second qubit, so that $U_1 X U_1^\dagger = -U_2 Z U_2^\dagger = (Z + X)/\sqrt{2}$, $U_1 Z U_1^\dagger = U_2 X U_2^\dagger = (Z - X)/\sqrt{2}$ and $U_1 Y U_1^\dagger = U_2 Y U_2^\dagger = -Y$. In this new basis H is a symmetric Z -matrix when

$$a_X \geq a_Z, \quad 2(a_Z - a_X)^2 \geq (a_{XX} + 1)^2. \quad (7)$$

For $a_{XX} \leq \sqrt{2} - 1$, and sufficiently large $|a_Z - a_X|$, these inequalities can certainly be satisfied.

In Figure 1 we show in what region of the parameter space (a_X, a_Z, a_{XX}) H is stoquastic using the general set of inequalities given in Appendix C.




Figure 1: A coarse-grained visualization of the approximate region where the Hamiltonian in Equation (6) is stoquastic. The purple cubes correspond to points in parameter space where a solution is found to the system of equations (28), (29) and (30) when sampling a uniform lattice of points. The region obscured by the blocks is filled.

2.3 Deciding Realness With Makhlin Invariants

To characterize the set of stoquastic 2-qubit Hamiltonians one needs to consider the orbit of \mathcal{C}_2 under the local rotations. For a traceless Hermitian matrix acting on two qubits, there is a complete set of 18 polynomial invariants $\{I_l\}_{l=1}^{18}$ specifying the Hermitian matrix up to local unitary rotations, which we will refer to as the Makhlin invariants [24]. Any two traceless Hamiltonians which have the same Makhlin invariants must be equivalent up to local unitary rotations, and vice versa. Since the invariants are straightforward to compute, if we are interested in some property only up to local unitary rotations, such as being stoquastic, then if we can find a way of characterizing that property in terms Makhlin invariants, it may be easier to test for that property. Note that the fact that there are 18 polynomial invariants can be consistent with the existence of only $d_{\text{inv}} = 15 - 6 = 9$ independent non-local parameters in a two-qubit traceless Hermitian matrix [25, 26]. One can see these 18 Makhlin invariants simply as a convenient set of invariants.

Stoquastic two-qubit Hamiltonians thus form a volume in the $d_{\text{inv}} = 9$ -dimensional non-local parameter space: the question is how the invariants are constrained in this stoquastic subspace.

We have not been able to express the stoquasticity of a two-qubit Hamiltonian in terms of inequalities involving the Makhlin or other invariants. However, we

are able to capture the condition that H can be made real by local rotations in terms of some of the Makhlin invariants needing to be zero, see Theorem 7. This is useful because realness is a necessary condition for a Hamiltonian to be stoquastic³.

The convenience of the Makhlin invariants is that they can be expressed in terms of inner products and triple products constructed from β , S and P which are invariant under $SO(3)$ rotations. Of particular interest are the triple product Makhlin invariants I_{10} , I_{11} and $I_{15} - I_{18}$. These take the following form (using the notation that $(\vec{a}, \vec{b}, \vec{c})$ stands for the scalar triple product $\vec{a} \cdot (\vec{b} \times \vec{c})$):

$$\begin{aligned} I_{10} &= (S, \beta\beta^T S, [\beta\beta^T]^2 S) \\ I_{11} &= (P, \beta^T \beta P, [\beta^T \beta]^2 P) \\ I_{15} &= (S, \beta\beta^T S, \beta P) \\ I_{16} &= (\beta^T S, P, \beta^T \beta P) \\ I_{17} &= (\beta^T S, \beta^T \beta \beta^T S, P) \\ I_{18} &= (S, \beta P, \beta\beta^T \beta P) \end{aligned} \quad (8)$$

One can prove that

Theorem 7. *A two-qubit Hamiltonian H is real under local unitary rotations if and only if all of the triple product invariants given in Table 8 are equal to zero.*

Proof Sketch. The proof of this Theorem is simple in one direction, namely when H is real, then the triple product invariants of the corresponding β , S , P are zero. This is because the Y coefficients in S and P are zero, and so they lie in a common two-dimensional subspace and the repeated multiplication by β or β^T keeps them in this subspace. Any triple of vectors, all lying in a two-dimensional space have zero triple product.

In order to prove the other direction, one can observe that the triple product invariants are all equal to zero if and only if the vectors transforming under O_L are co-planar, and the vectors transforming under O_R are co-planar. The vectors transforming under O_L are

$$\Gamma_L = \{S, \beta\beta^T S, [\beta\beta^T]^2 S, \beta P, \beta\beta^T \beta P\},$$

and the vectors transforming under O_R are

$$\Gamma_R = \{P, \beta^T \beta P, [\beta^T \beta]^2 P, \beta^T S, \beta^T \beta \beta^T S\}.$$

To prove that zero triple product invariants implies realness under local rotations is more work since β can be degenerate and have rank less than 3, so the proof has to incorporate the various cases. We have included a full proof of Theorem 7 in the Appendix B. \square

³It is well known that any k -local Hamiltonian can be mapped onto a real $k + 1$ -local Hamiltonian, see e.g. [21]. This is not directly useful since deciding whether local basis changes exist which make a 3-local Hamiltonian have non-positive off-diagonal elements will be harder than the same problem for 2-local Hamiltonians.

2.4 Reducing the Two-Qubit Stoquastic Decision Problem to Polynomial Inequalities

Here we present a general analytical strategy to determine if a two qubit Hamiltonian is stoquastic, and in what basis.

First apply the Makhlin invariants from the previous section to determine if the Hamiltonian can be made real. If not, then it is not stoquastic.

We would now like to put our Hamiltonian into a *standard form*. Recalling the discussion in section 2.1 we know that it is always possible to bring β into a diagonal form:

$$\beta = \begin{pmatrix} a_{XX} & 0 & 0 \\ 0 & a_{YY} & 0 \\ 0 & 0 & a_{ZZ} \end{pmatrix}.$$

A trivial consequence of Theorem 7 is that H is real under local unitary rotations if and only if there exists a transformation which diagonalizes β and puts S and P into the form:

$$S = \begin{pmatrix} a_{XI} \\ 0 \\ a_{ZI} \end{pmatrix}, \quad P = \begin{pmatrix} a_{IX} \\ 0 \\ a_{IZ} \end{pmatrix}. \quad (9)$$

So, given that we know our Hamiltonian can be made real under local unitary rotations, we say our standard form must have this structure. In this form H is real, because $a_{IY} = a_{YI} = a_{YX} = a_{ZY} = a_{YZ} = 0$, and since this is a necessary condition for H to be a Z -matrix we demand that any further transformation we perform preserves this realness.

We would now like to add some additional structure to our standard form. Before we can do this, we must handle some simple special cases.

The first special case is when $S = P = 0$. In this case we know H is stoquastic, because one can always apply an appropriate choice of permutations and sign-flips on β so that $a'_{XX} \leq -|a'_{YY}|$.

As a second special case, if $\beta = 0$ then H is stoquastic, since S and P may be freely rotated independently.

Finally, as a third special case, if $a_{ZZ} = a_{XX} = 0$, $a_{YY} \neq 0$, at least one of a_{XI} and a_{IX} is non-zero, and at least one of a_{ZI} and a_{IZ} is non-zero, then H is not stoquastic. This is true by the following reasoning. Any transformation which makes H a Z -matrix must satisfy $a'_{XX} \leq -|a'_{YY}|$ so that $a'_{XX} \neq 0$ or $a'_{YY} = 0$. If $a'_{YY} \neq 0$ then $a'_{XX} \neq 0$ and since β is rank-1, $a'_{XY} \neq 0$ and $a'_{YX} \neq 0$, which would disqualify H from being a Z -matrix. Therefore any transformation which makes H a Z -matrix must set $a'_{YY} = 0$. However any such transformation will set at least one other $a'_{\bullet Y}$ or $a'_{Y \bullet}$ to be non-zero. Therefore H is not stoquastic.

We may proceed under the assumption that these special cases have been handled, and argue that we can always put S , P and β into a standard form with the following properties, namely at least one S or P is non-zero and $\beta \neq 0$. Furthermore we may assume that $a_{ZZ} \geq a_{XX} \geq 0$. This follows from the fact that a_{ZZ} and a_{XX} can always be made non-negative by perform sign flips on the eigenvalues of β and dumping any extra sign into a_{YY} so as to preserve the determinant, furthermore any permutation between the a_{XX} and a_{ZZ} eigenvalues will preserve the standard form. Finally we may assume H can be normalized so that $a_{ZZ} = 1$. This follows from the fact that the exclusion of the second and third special cases allows us to always put S , P , and β into a standard form such that $a_{ZZ} > 0$, since if $a_{ZZ} = a_{XX} = 0$, but $a_{YY} \neq 0$ then either $a_{IX} = a_{XI} = 0$ or $a_{IZ} = a_{ZI} = 0$, in which case a permutation can always be performed that preserves the standard form but sets $a'_{ZZ} \neq 0$.

Having put S , P , and β into standard form with the above structure, we wish to know if there exists a pair of $SO(3)$ rotations O_1 and O_2 such that $\beta' = O_1\beta O_2^T$, $S' = O_1S$ and $P' = O_2P$ where β' , S' , P' are associated with a symmetric Z -matrix Hamiltonian. As discussed above, any such transformation must preserve the realness of the Hamiltonian. If it is not the case that either $a_{XI} = a_{IX} = 0$ or $a_{ZI} = a_{IZ} = 0$ then the only transformations which preserve this realness are $SO(3)$ rotations in the X - Z subspace combined with reflections in the X - or Z -axis, given by

$$O = \begin{pmatrix} \cos(\theta) & 0 & \sin(\theta) \\ 0 & \gamma & 0 \\ -\gamma \sin(\theta) & 0 & \gamma \cos(\theta) \end{pmatrix}, \quad \gamma = \pm 1 \quad (10)$$

(If it is the case that either $a_{XI} = a_{IX} = 0$ or $a_{ZI} = a_{IZ} = 0$ then there is an additional transformation which can be performed, see the extra step in Appendix C.)

Recalling Proposition 6, three inequalities must be satisfied in order for the rotated H with β' , S' , P' to be a symmetric Z -matrix. As we show in Appendix C these inequalities can be re-expressed as systems of two-variable polynomials which are at most quadratic in either variable, so analytic solutions to their roots can be constructed and solutions can be found using graphical methods. However, the complexity of these inequalities makes it unclear whether their interest goes beyond that of numerically finding a set of local basis changes.

3 The Stoquastic Problem for 2-local Hamiltonians

Consider a general n -qubit 2-local Hamiltonian

$$H = \sum_{uv \in E} H_{uv} + \sum_{v \in V} H_v, \quad (11)$$

where u and v are the vertices of the interaction graph $G = (V, E)$ of the Hamiltonian H . Each vertex $v \in V$ corresponds to a qubit so that H_v are the 1-local terms (corresponding to S and P in the two-qubit case). Each edge $e = uv \in E$ is represented as H_{uv} and can be associated with the 3×3 matrix β^{uv} for that edge.

One wants to determine when such a Hamiltonian is stoquastic. Recalling Definition 2, we want to find both a basis as well as a local decomposition $H = \sum_i D_i$ so that each D_i is a Z -matrix acting on at most 2 qubits. We begin with two simple cases. In the first case we consider how to find a local decomposition so that each D_i is a Z -matrix, ignoring any basis transformations. Next we consider a family of natural Hamiltonians which have identical two-local terms on a bipartite graph. In this case the problem of finding an appropriate decomposition and basis transformation becomes much simpler. The final case we consider is the XYZ Heisenberg Hamiltonian, where we show that there exists a polynomial time algorithm for deciding if such a Hamiltonian is stoquastic. This constitutes our most significant result.

3.1 Finding a Decomposition in a Fixed Basis

Even given a fixed basis it remains non-trivial to identify if a decomposition of the form $H = \sum_i D_i$ exists.

We can define the set of 2-local symmetric Z -matrices as the polyhedral cone \mathcal{C}_n which is generated by the extremal vectors of each of the 2-qubit cones \mathcal{C}_2 , which are now embedded in a higher dimensional n -qubit Hilbert-Schmidt space.

Given a fixed basis, the cone \mathcal{C}_n corresponds to those H which admit such a decomposition $H = \sum_i D_i$ so that each D_i is a Z -matrix acting on at most 2 qubits, and we are interested in deciding our Hamiltonian lies in that cone. A necessary condition is that the Hamiltonian is real, however this is straightforward to check. So going forward we assume that H is real and want to decide whether $H \in \mathcal{C}_n$.

It is important to note that for any v , \mathcal{C}_2^{uv} and $\mathcal{C}_2^{vu'}$ intersect: for any $u, u' \neq v$, the extremal vectors obey $\mathbf{o}_3^{uv} + \mathbf{o}_4^{uv} = \mathbf{o}_5^{vu'} + \mathbf{o}_6^{vu'}$. This element in the intersection of \mathcal{C}_2^{uv} with $\mathcal{C}_2^{vu'}$ is the 1-local $-X_v$ term. This non-empty intersection shows the freedom in the decomposition D_i for a fixed basis. The parsimonious strategy below shows how to find a decomposition D_i of a n -qubit Hamiltonian H in terms of 2-local terms each

of which is a 4×4 symmetric Z-matrix or decide that it does not exist.

Let H_u^X be the 1-local term proportional to X_u and similarly H_u^Z . First, note that a real H can only be in \mathcal{C}_n when for all u , $H_u^X \leq 0$. This follows directly from Proposition 6 demanding that $a_{IX}, a_{XI} \leq 0$. Hence we assume this to be the case (otherwise we conclude that no decomposition exists in the given basis).

Efficient Parsimonious Strategy:

Repeat the following for all edges uv in H .

1. Given the current Hamiltonian H , pick a pair of vertices u and v and consider $h(\alpha, \beta) = H_{uv} + H_u^Z + H_v^Z - \alpha X_u - \beta X_v$ which includes *all* current single-qubit terms which act on vertices u and v and wlog $\alpha \geq 0, \beta \geq 0$.

Find the minimal $\alpha_{\min} \leq \alpha, \beta_{\min} \leq \beta$ such that $h(\alpha_{\min}, \beta_{\min}) \in \mathcal{C}_2$, or decide that $h(\alpha, \beta) \notin \mathcal{C}_2$. In the latter case, decide that $H \notin \mathcal{C}_n$ and exit. Note that when $h(\alpha, \beta) \notin \mathcal{C}_2$ then $\forall \alpha' \leq \alpha, \beta' \leq \beta$, $h(\alpha', \beta') \notin \mathcal{C}_2$ since $a_{IX} \leq -|a_{ZX}|$ and $a_{XI} \leq -|a_{XZ}|$, aka it is easier to satisfy the stoquastic condition for large α, β , more negative 1-local X -terms. Hence the goal is to use as little α and β as possible, be parsimonious, so that a lot of $-X_u$ and $-X_v$ will be left over for another edge. This is then clearly an optimal strategy.

2. Define the left-over Hamiltonian $H_{\text{left-over}} = H - h(\alpha_{\min}, \beta_{\min})$ and repeat the previous step with the left-over Hamiltonian as the current Hamiltonian.
3. Stop the iteration when you have done all the edges uv . Either the left-over Hamiltonian is now 0 or it is the sum of 1-local X -terms, each of which has a negative sign, hence $H \in \mathcal{C}_n$.

This problem is also discussed in [19], where they argue that it can in general always be solved using linear programming.

3.2 Uniform Bipartite Graphs

Let us next consider a particularly simple example for which a resolution of the two-qubit problem analyzed in section 2 suffices for deciding if the Hamiltonian is stoquastic.

Proposition 8. *Let the interaction graph $G = (V, E)$ of a Hamiltonian H be bipartite, i.e. $V = V_A \cup V_B$, $V_A \cap V_B = \emptyset$, and there are only edges $uv \in E$ with $u \in V_A, v \in V_B$. Furthermore, $H = \sum_{uv \in E} h_{uv}$ where h_{uv} acts with both one and two local terms on sites u and v , and $h_{uv} = h$ for all $u \in V_A, v \in V_B$. If the two-qubit Hamiltonian h is stoquastic, then H is stoquastic. If h is not stoquastic then H will not be stoquastic under any*

local basis change which acts identically on all qubits in a partition.

Bipartite graphs include linear arrays, square lattices, cubic lattices and hexagonal lattices, all of which are very natural structures to consider.

Proof. If there exists a $U_A \otimes U_B$ such that $(U_A \otimes U_B)h(U_A \otimes U_B)^\dagger$ is a symmetric Z-matrix, then apply this rotation U_A to all $u \in V_A$ and apply U_B to all $v \in V_B$ and $H' = \sum_{uv \in E} (U_A \otimes U_B)h_{uv}(U_A \otimes U_B)^\dagger$ with $(U_A \otimes U_B)h_{uv}(U_A \otimes U_B)^\dagger$ a symmetric Z-matrix.

Suppose h were not stoquastic, but there existed a pair of unitaries U_A and U_B such that a decomposition $H = \sum_{uv \in E} D_{uv}$ exists where $D'_{uv} = (U_A \otimes U_B)D_{uv}(U_A \otimes U_B)^\dagger$ is a symmetric Z-matrix. $h'_{uv} = (U_A \otimes U_B)h_{uv}(U_A \otimes U_B)^\dagger$ cannot be a symmetric Z-matrix, and so $D_{uv} \neq h_{uv}$. However both D_{uv} and h_{uv} must share the same purely two local parts. As such, D'_{uv} must differ from h'_{uv} by its 1-local terms, and in particular one or both of the $-X_u$ and $-X_v$ terms. Since h'_{uv} is not stoquastic, D'_{uv} must have more support on either $-X_u$ or $-X_v$ than h'_{uv} . Suppose wlog that D'_{uv} has more support on $-X_u$, then, by the parsimonious reasoning outlined earlier, there must exist some D'_{ux} which has less support on $-X_u$ than h'_{ux} does, and consequently D'_{ux} is not a symmetric Z-matrix, leading to a contradiction. \square

3.3 Efficient Algorithm for XYZ Heisenberg Models

In this section we present and prove our main result, namely that there is an efficient algorithm to decide if an XYZ Heisenberg Hamiltonian is stoquastic.

We define an XYZ Heisenberg Hamiltonian as an n qubit Hamiltonian of the form

$$H = \sum_{uv} H_{uv}, \quad (12)$$

$$H_{uv} = a_{XX}^{uv} X_u X_v + a_{YY}^{uv} Y_u Y_v + a_{ZZ}^{uv} Z_u Z_v. \quad (13)$$

The interaction graph of the Hamiltonian may be of any form. Our main theorem is as follows.

Theorem 9. *There exists a constructive algorithm, which runs in $O(n^3)$ time, that decides if an XYZ Hamiltonian H is stoquastic.*

3.3.1 Preliminaries and Proof Outline

Our proof consists of first describing an algorithm, and secondly proving the correctness of the algorithm.

The general aim of the algorithm is to find a set of single qubit Clifford rotations which transforms the Hamiltonian into a Z-matrix. If such a set is found, then our

Hamiltonian is stoquastic. We claim that if our algorithm cannot find such a set, then no such set exists, and furthermore that this implies that no local unitary rotation exists which makes H a Z-matrix, and therefore the Hamiltonian is not stoquastic.

We find that the most natural way to represent the problem is as a matrix-weighted graph, with each vertex corresponding to a qubit, and each edge corresponding to an interaction term H_{uv} . The matrix weight of the edge corresponding to the term H_{uv} is the β^{uv} matrix discussed in Section 2.1, i.e. a 3×3 real diagonal matrix whose diagonal entries correspond to the coefficients of the XX, YY and ZZ terms. As discussed in Section 2.1, single qubit unitary rotations on qubits u and v correspond to $SO(3)$ rotations O_u and O_v acting by left and right matrix multiplication: $O_u \beta^{uv} O_v^T$. In this picture single-qubit Clifford rotations correspond to signed permutations with determinant 1, and so the goal of our algorithm is to assign signed permutations Π_u to each vertex so that they transform all of the weights $\Pi_u \beta^{uv} \Pi_v^T$ into a form which satisfies the stoquasticity conditions given by proposition 6. We refer to such an assignment $\{\Pi_u\}$ as a **solution**.

Our algorithm relies on three important points. The first is that it is sufficient to consider only Clifford transformations, instead of all possible unitaries, so in our algorithm we only need to consider signed permutations. The second point is that any solution which transforms the Hamiltonian into a Z-matrix must preserve the diagonal form of the β matrices. This follows from Proposition 6 and the fact that there are no single qubit terms. Thus an assignment of signed permutations $\{\Pi_u\}$, with $\tilde{\beta}^{uv} = \Pi_u \beta^{uv} \Pi_v^T$ is considered a solution if and only if

$$\tilde{\beta}_{ij}^{uv} = \delta_{ij} \tilde{\beta}_{ii}^{uv}, \text{ ie } \tilde{\beta} \text{ remains diagonal,} \quad (14)$$

and

$$\tilde{\beta}_{11}^{uv} \leq -|\tilde{\beta}_{22}^{uv}| \quad (15)$$

the latter of which is equivalent to

$$[\tilde{\beta}_{11}^{uv}]^2 \geq [\tilde{\beta}_{22}^{uv}]^2 \text{ and} \quad (16)$$

$$\tilde{\beta}_{11}^{uv} \leq 0 \quad (17)$$

The third point is that when an edge has a matrix weight β which has a rank greater than 1, then in order to preserve the diagonal form of the matrix, the signed permutation acting from the left must be the same as to the signed permutation acting on the right, up to a difference of signs.

Any solution of signed permutations $\{\Pi_u\}$ admits a decomposition into a specification of permutations $\{\Pi_u \in S_3\}$, and signed diagonal matrices $\{R_u = \text{diag}(\pm 1, \pm 1, \pm 1)\}$ so that $\{\Pi_u = R_u \Pi_u\}$. The algorithm thus breaks up into two parts. The goal of the

first part is to restrict the possible permutations $\{\Pi_u\}$ to those which satisfy condition 16. The goal of the second part is to determine for which of those possible permutations an assignment of signs $\{R_u\}$ can be made so that condition 17 is also satisfied.

The general strategy of the first part of the algorithm is to ‘quotient out’ clusters of vertices which are connected by rank>1 edges, since the permutations applied to those vertices need to be all identical. One is left with a graph composed only of rank-1 edges. The goal is then to find permutations at the vertices which transform away any rank-1 β matrix proportional to a YY term in the Hamiltonian, since such lone (not accompanied by XX) YY terms are forbidden by inequality 16. By temporarily ignoring particular vertices, this task can be reduced to deciding if there exists an exact satisfying assignment to a classical Ising problem, which is an XOR-SAT problem and can be decided straightforwardly. The solution to this Ising problem translates back to the original graph in the form of sets of compatible permutation assignments.

We then proceed to the second part of the algorithm, in which we determine for which (candidate) selections of permutations one can choose signs so that condition 17 is also satisfied.

For any given choice of permutations, deciding the existence of such appropriate signs again reduces to deciding if there exists an exact solution to a classical Ising problem, and the number of such Ising problems one needs to check is polynomially bounded.

For the sake of clarity we present the definitions, the algorithm, and the proof of correctness separately, but the reader may find it useful to read these sections in parallel.

3.3.2 Definitions

Definition 10 (weighted interaction graph). *Given a Hamiltonian, the **weighted interaction graph** is a matrix weighted graph, with each vertex corresponding to a qubit, and each edge (u, v) weighted by the β^{uv} matrix of the H_{uv} term in the Hamiltonian.*

In this algorithm we will be considering weighted interactions graphs of n qubit XYZ Heisenberg Hamiltonians.

Definition 11 (rank-1 and rank>1 edge). *An edge in a weighted interaction graph is **rank-1** if its matrix weight has a rank 1. Conversely, an edge in a weighted interaction graph is **rank>1** if its matrix weight has a rank greater than 1.*

Definition 12 (rank>1 connected component). *Consider a weighted interaction graph G . Remove all rank-1*

edges. One is left with a set of distinct connected components which are composed entirely of rank >1 edges. Each of these connected components is a subgraph of G which we call a rank >1 connected component Γ .

Note that a rank >1 connected component may contain a single vertex with no edges, in the case where some vertex is connected to only rank-1 edges. Note also that every vertex in the weighted interaction graph belongs to exactly one rank >1 connected component.

Definition 13 (rank-1 quotient graph). *Given a weighted interaction graph of an XYZ Heisenberg Hamiltonian, a **rank-1 quotient graph** is a multi-graph with labelled edges which is constructed as follows. For each rank >1 connected component in the weighted interaction graph populate the rank-1 quotient graph with a corresponding vertex. Since every vertex in the weighted interaction graph belongs to exactly one rank >1 connected component, every vertex in the weighted interaction graph has a corresponding representative in the rank-1 quotient graph. For every pair of vertices in the weighted interaction graph connected by a rank-1 edge, connect their corresponding representative vertices in the rank-1 quotient graph by an edge. Label this new edge either 1, 2, or 3, corresponding to the index i for which $[\beta]_{ii} \neq 0$.*

This construction effectively quotients out all the rank >1 connected components. It is however a multi-graph because there may be multiple rank-1 edges connecting the vertices in a pair of rank >1 connected components.

Definition 14 (single-label vertex). *Given a rank-1 quotient graph a **single-label vertex** is a vertex incident to edges of a single label.*

Definition 15 (single-label connected component). *Consider a rank-1 quotient graph G_Q with all vertices removed which are not single-label, call this graph G_Q^- . For each subset of vertices S_Q associated with a connected component of G_Q^- , the subgraph Γ_Q of G_Q induced⁴ by S_Q is called a **single-label connected component**.*

Note that because every vertex in a given single-label connected component is incident to edges of a single label i , every edge in the single-label connected component must have a common label i , and we may say that a single-label connected component is labelled i . Every single-label vertex in the rank-1 quotient graph belongs to exactly one single-label connected component.

⁴A subgraph of a graph $G = (V, E)$ induced by a subset of vertices $S \subseteq V$ is defined as a graph whose vertex set is S and whose edge set is all edges in E that have both endpoints in S .

Definition 16 (heterogeneous rank-1 quotient graph). *Given a rank-1 quotient graph, the **heterogeneous rank-1 quotient graph** is a copy of the rank-1 quotient graph modified in the following way. For each single-label connected component, with label i , remove all vertices in the single-label connected component from the rank-1 quotient graph, and connect every vertex in the boundary of the single-label connected component to every other vertex in the boundary, as well as to itself, by an edge labelled by i .*

By construction, every single-label vertex is removed from the heterogeneous rank-1 quotient graph. Therefore every vertex in this new heterogeneous rank-1 quotient graph will be incident to edges labelled by more than one label.

Definition 17 (Admissible Permutations). *The set of **admissible permutations** \tilde{S} of a rank >1 connected component Γ is a subset of permutations $\tilde{S} \subseteq S_3$ such that for each permutation $\Pi \in \tilde{S}$, when applied identically to every vertex in Γ , transforms every edge weight in Γ to a form which satisfies condition 16. Since every vertex u in both the rank-1 quotient graph and the heterogeneous rank-1 quotient graph is associated with a rank >1 connected component, each such vertex is associated with a set of **admissible permutations** \tilde{S}_u .*

Definition 18 (Compatible Set of Permutation Assignments). *Given a rank-1 quotient graph, or a heterogeneous rank-1 quotient graph, a **compatible set of permutation assignments** is an assignment of sets of possible permutations to each vertex in the graph: $A = \{\Sigma_u \subseteq \tilde{S}_u\}$ satisfying the following property. For every edge with label i and end vertices u and v , there exists a number $j \neq 2$ such that every permutation Π in Σ_u or Σ_v satisfies $\Pi(i) = j$.*

Definition 19 (partitioned weighted interaction graph). *Given a rank-1 quotient graph, a weighted interaction graph and a compatible set of permutation assignments: $A = \{\Sigma_u \subseteq S_3\}$, the **partitioned weighted interaction graph** is a copy of the weighted interaction graph modified in the following way. For each edge in the rank-1 quotient graph labelled i , with end vertices u and v , any permutation Π in Σ_u or Σ_v will map i to specific label $\Pi(i) = j \neq 2$. If $j \neq 1$, remove the corresponding edge from the weighted interaction graph.*

Note that the partitioned weighted interaction graph may or may not be fully connected.

Definition 20 (Ising Partition). *Given a connected component Γ_W of a partitioned weighted interaction graph and an assignment of permutations $\{\Pi_u\}$ to each vertex u in the connected component, an **Ising Partition** is a copy of Γ_W , with each edge labelled by the value $-\text{sign}([\Pi_u \beta^{uv} \Pi_v^T]_{11})$.*

The Ising partition can be thought of as a classical Ising model, since it is a graph whose edges are labelled by $+1$ or -1 . One can ask if there exists an assignment of $+1$ or -1 to each vertex in the graph so that for every edge the product of its neighbouring vertices is equal to the sign associated with that edge. If such an assignment exists, we call it an exact satisfying solution to the classical Ising problem. Note that this problem is distinct from the generic classical Ising problem, which asks for a minimizing assignment instead of an exact assignment. It is straightforward to decide if such an exact satisfying solution exists. Simply assign a sign to one vertex, and propagate that choice through the graph, according to the edge constraints, until you either encounter a term which is not satisfied under this procedure, or have successfully assigned signs to every vertex. Since the procedure runs over all edges it runs in $O(n^2)$ time. If a solution exists, then two solutions must exist, since one can simply flip every sign and still have solution.


3.3.3 The Algorithm

The algorithm presented here takes as input an n qubit XYZ Heisenberg Hamiltonian, and outputs either true, indicating that the Hamiltonian is stoquastic, or false, indicating that the Hamiltonian is not stoquastic. In the event that the algorithm outputs true, the algorithm also provides a set of local unitary rotations which transform the Hamiltonian into a Z-matrix. In this sense it is a constructive algorithm.

Part 1 of the Algorithm


Step 1: Construct the weighted interaction graph of the Hamiltonian. $O(n^2)$

Step 2: Identify all rank >1 connected components of the weighted interaction graph and determine their admissible permutations. $O(n)$




Step 3: Construct the rank-1 quotient graph of the weighted interaction graph. $O(n^2)$

Step 4: If there are any vertices in the rank-1 quotient graph which connect to at least three edges, each with a different label,



then return **false**. If not, then every vertex in our graph connects to edges which are labelled by at most two distinct labels. $O(n^2)$


Step 5 edge case: If the whole rank-1 quotient graph is itself a single-label connected cluster, then and only then will the heterogeneous rank-1 quotient graph be empty. If so proceed to the step 9 edge case. $O(n)$



Step 5: Construct the heterogeneous rank-1 quotient graph from the rank-1 quotient graph. Since every vertex in the rank-1 quotient graph connects to edges which are labelled by at most two distinct labels,

every vertex in the heterogeneous rank-1 quotient graph must be incident to edges with exactly two distinct labels. Let us denote a vertex incident to edges labelled i and j as an ij -vertex. $O(n)$


Step 6: A classical Ising problem may be specified by a graph with $+1$ and -1 labels on the edges, and the solution corresponds to an assignment of $+1$ and -1 to the vertices so that the product of an edge with its vertex assignments produces $+1$. Using the heterogeneous rank-1 quotient graph, construct an



Ising problem in the following way. Copy the graph structure of the heterogeneous rank-1 quotient graph. Label the edges of this new graph using the following prescription. $O(n^2)$

- For each edge in the heterogeneous rank-1 quotient graph labelled 2, with one incident 12-vertex and one incident 23-vertex, label the corresponding edge in the Ising problem with a -1 .
- Label all other edges in the Ising problem with a $+1$.

Step 7: If there does not exist an exactly satisfying solution to this classical Ising problem then return **false**. If there does, then there must exist two possible solutions I_1 and I_2 . $O(n)$



Step 8: A solution to the Ising problem described above uniquely specifies an assignment of permutations to the vertices of the heterogeneous rank-1 quotient graph, when it is combined with the information about the edges to which each of those vertices connect. This is most intuitively illustrated by figure 2, and a concise description of the unique assignment is given by table 1.

Ising solution:	+1	-1
13-vertex	(1)(2)(3)	(13)(2)
12-vertex	(1)(23)	(132)
23-vertex	(12)(3)	(123)

Table 1: Translation between permutation and Ising solution





Figure 2: A complete illustration of which permutations to apply to a vertex, given a solution to an Ising problem on the heterogeneous rank-1 quotient graph

Using this rubric, and the two solutions I_1 and I_2 , construct two possible assignments of permutations to the vertices of the heterogeneous rank-1 quotient graph $A_1^h = \{\Sigma_u^h = \{\Pi_u^1\}\}$ and $A_2^h = \{\Sigma_u^h = \{\Pi_u^2\}\}$. If for either assignment, the permutation assigned to some vertex is not in the admissible permutations of that vertex, then discard that assignment. If both assignments are discarded in this way, return **false**. Any remaining A^h constitute compatible sets of permutations assignments to the heterogeneous rank-1 quotient graph, with each Σ_u^h containing a single element. $O(n)$




Step 9 preliminaries We now wish to construct all compatible sets of permutation assignments for the rank-1 quotient graph. Note that unlike in the case of the heterogeneous rank-1 quotient graph, the permutations at the single label vertices may not be completely specified in a given compatible set of permutations assignment, so that some Σ_u may have more than one element.

Step 9 edge case: In the case where the whole rank-1 quotient graph is a single-label connected component with label i , as is the case in the step 5 edge case, construct two compatible sets of permutation assignments A_1 and A_2 . $A_1 = \{\Sigma_u^1\}$, and Σ_u^1 is all permutations Π in the admissible permutations of vertex u which satisfy $\Pi(i) = 1$. $A_2 = \{\Sigma_u^2\}$, and Σ_u^2 is all permutations Π in the admissible permutations of vertex u which satisfy $\Pi(i) = 3$. If either A_1 or A_2 contains an element Σ_u

which is empty, then discard it. If both A_1 and A_2 have been discarded, return **false**. $O(n)$

Step 9: Each compatible set of permutation assignments to the heterogeneous rank-1 quotient graph translates into an assignment of permutations to all of the non-single-label vertices in the rank-1 quotient graph. This partially




specifies up to two compatible sets of permutation assignments $A_1 = \{\Sigma_u^1\}$ and $A_2 = \{\Sigma_u^2\}$, where $\Sigma_u^1 = \Sigma_u^h \in A_1^h$ and $\Sigma_u^2 = \Sigma_u^h \in A_2^h$ are single element sets defined only on non-single label vertices. For each compatible set of permutation assignments A_x , perform the following procedure. For each single-label connected component in the rank-1 quotient graph, with label i , choose a vertex u from its boundary, which will have a permutation Π_u^x assigned to it. There exists only one other permutation $\Pi' \neq \Pi_u^x$ such that $\Pi_u^x(i) = \Pi'(i)$. For each vertex v in the single-label connected component, let Σ_v^x be all permutations in $\{\Pi', \Pi_u^x\}$ which are also in the admissible permutations of vertex v . If either A_1 or A_2 contains an element Σ_u which is empty, then discard it. If both A_1 and A_2 have been discarded, return **false**. $O(n)$

Conclusion of Part 1 of the Algorithm: We have now constructed up to two compatible sets of permutation assignments to the vertices in the rank-1 quotient graph. Together these sets represent all possible assignments of permutations to the vertices which satisfy conditions 16 and 14. The aim of the second part of the algorithm is to determine whether or not there exists any assignments of permutations from these sets which also admit an assignment of signs at every vertex $\{R_u = \text{diag}(\pm 1, \pm 1, \pm 1)\}$ so that condition 17 is satisfied.

Part 2 of the Algorithm

Step 10: Let $\text{sol} = \{\Pi_u = R_u \Pi_u\}$ store the potential solution. For each set of compatible permutation assignments $A_x = \{\Sigma_u^x\}$, $x \in 1, 2$, from part 1, perform steps 11 through 13, discarding those A_x which fail.

Step 11: Construct the partitioned weighted interaction graph. $O(n^2)$




Step 12: By construction, every single-label vertex u in the rank-1 quotient graph, with label i , for which the permutations Π in Σ_u^x satisfy $\Pi(i) = 3$, will correspond to a unique connected component in

the partitioned weighted interaction graph. For each such vertex u , do the following. For each permutation $\Pi \in \Sigma_u^x$, the permutation assignment for each vertex v in the connected component is given by Π . Construct the corresponding Ising Partition and decide if it admits a solution. If yes, then choose a solution and let $\Pi_u = \text{diag}(\delta_u, 1, \delta_u \det(\Pi))\Pi$, and $\delta_u \in \{+1, -1\}$ is given by the solution to the Ising model. If neither permutation in Σ_u^x admits a solution, then discard A_x . In the worst case there are n single-label vertices, and in the worst case constructing and solving an Ising Partition is $O(n^2)$, so in the naive worst case the runtime is $O(n^3)$.

Step 13: For the remaining connected components Γ_P in the partitioned weighted interaction graph do the following. For each vertex v in Γ_P , choose a permutation $\Pi_v \in \Sigma_v^x$.

Equipped with this permutation assignment, construct the corresponding Ising partition and decide if it admits a solution. If not, then discard A_x . If yes, then choose a solution and let $\Pi_v = \text{diag}(\delta_v, 1, \delta_v \det(\Pi_v))\Pi_v$ for all v in Γ_P , and $\delta_v \in \{+1, -1\}$ is given by the solution to the Ising model. By similar reasoning to step 12, the worst case runtime is $O(n^3)$.

Step 14: If all A_x have been discarded, return **false**. Otherwise return **true**, along with sol , which will have been completely specified for every vertex in the weighted interaction graph.



3.3.4 Correctness of the algorithm

It is obvious that the algorithm is complete, as long as it is supplied an XYZ Heisenberg Hamiltonian, it will always either return true or false. We need only prove soundness. The proof of soundness of the algorithm is twofold. We must first show that when the algorithm returns false then H is not stoquastic. Secondly we must show that when the algorithm returns true H is stoquastic, and sol corresponds to the rotation which transforms it into a Z-matrix.

Lemma 21 (All Diagonal Matrix Weighted Graphs are Connected by Signed Permutations). *Given a set of diagonal 3×3 matrices $\{\beta^{uv}\}$ which are diagonal in some basis $b = \{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$. For any set of $SO(3)$ rotations $\{O_u\}$ such that $O_u \beta^{uv} O_v^T$ is also diagonal in basis b for all β^{uv} , there exists a set of signed permutations $\{\Pi_u\}$ such that $\Pi_u \beta^{uv} \Pi_v^T = O_u \beta^{uv} O_v^T$*

The proof of lemma 21 can be found in appendix D

Lemma 22 (Single-Qubit Clifford Rotations Suffice). *The XYZ Heisenberg Hamiltonian can be transformed into a Z-matrix by single-qubit unitary transformations if and only if it can be transformed into a Z-matrix by single-qubit Clifford transformations.*

Proof. The XYZ Heisenberg Hamiltonian corresponds to a set of diagonal 3×3 matrices $\{\beta^{uv}\}$, and any transformation which transforms one into a Z-matrix corresponds to a set of $SO(3)$ rotations $\{O_u\}$ such that $\tilde{\beta}_{uv} = O_u \beta^{uv} O_v^T$ satisfies conditions 14, 16 and 17 in the basis b for all u and v . Therefore $\tilde{\beta}_{uv}$ is diagonal, and by lemma 21 there exists a set of signed permutations $\{\Pi_u\}$ such that $\Pi_u \beta^{uv} \Pi_v^T = O_u \beta^{uv} O_v^T$ also satisfies those conditions. These signed permutations may not have determinant 1, but there exist determinant 1 signed permutations $\Pi'_u = \text{diag}(1, 1, \det \Pi_u)\Pi_u$ such that the matrices $\{\Pi'_u \beta^{uv} \Pi'_v\}$ also satisfy conditions 14, 16 and 17 in the basis b for all u and v . The other direction of the biconditional is trivial. \square

Lemma 23 (Permutations must be identical on rank >1 connected components). *If a set of signed permutations $\{\Pi_u = R_u \Pi_u\}$ is a solution, then for any rank >1 connected component Γ , there exists a permutation Π such that $\Pi_u = \Pi$ for every vertex u in Γ .*

Proof. If $\{\Pi_u\}$ is a solution, then $\Pi_u \beta^{uv} \Pi_v^T = \sum_m \beta_{mm}^{uv} R_{\Pi_u(m)}^u R_{\Pi_v(m)}^v |\Pi_u(m)\rangle \langle \Pi_v(m)|$ is diagonal. This implies that if $\beta_{mm}^{uv} \neq 0$ then $\Pi_u(m) = \Pi_v(m)$. Since $\beta_{mm}^{uv} \neq 0$ for two of the three values of m , $\Pi_u = \Pi_v$. This holds for every pair of connected vertices in the connected components, hence for all $u \in \Gamma$, $\Pi_u = \Pi$ for some Π . \square

Lemma 24 (rank-1 quotient graph permutation assignments are necessary). *If there does not exist a compatible set of permutation assignments for the rank-1 quotient graph, then H is not stoquastic.*

Proof. If H is stoquastic, then by lemma 22 there exists a set of signed permutations $\{\Pi_u = R_u \Pi_u\}$ which transform H into a Z-matrix, and so for every edge weight β^{uv} , $\tilde{\beta}^{uv} = \Pi_u \beta^{uv} \Pi_v^T$ satisfies conditions 14, 16 and 17. Every vertex x in the rank-1 quotient graph is associated with a rigid connected component Γ_x in the weighted interaction graph. For every edge in the rank-1 quotient graph, with label i , end vertices x and y and associated rigid connected components Γ_x and Γ_y , there exists an edge in the weighted interaction graph with end vertices $u \in \Gamma_x$ and $v \in \Gamma_y$ and edge weight β^{uv} with rank 1 and $\beta_i^{uv} \neq 0$. By lemma 23, there must exist permutations Π_x and Π_y such that for every vertex u in Γ_x and vertex v in Γ_y , $\Pi_u = R_u \Pi_x$ and $\Pi_v = R_v \Pi_y$. By conditions 14 and 16, $\Pi_x(i) = \Pi_y(i) \neq 2$. Thus such

a specification of permutations $A = \{\Gamma_x = \{\Pi_x\}\}$ corresponds to a compatible set of permutation assignments. The lemma follows contrapositively. \square

Lemma 25 (Single label paths are transformed uniformly). *Given a connected path of edges in a rank-1 quotient graph which are identically labelled i , for any compatible set of permutation assignments $A = \{\Sigma_u\}$ there exists a label $j \neq 2$ such that for every vertex u in the path, and for every $\Pi \in \Sigma_u$, $\Pi(i) = j$.*

Proof. For any vertex u in the path, there exists a label $j \neq 2$ such that for every $\Pi \in \Sigma_u$, $\Pi(i) = j$. By construction, every vertex in the path neighbouring u satisfies this condition as well. The proof follows by induction. \square

Lemma 26 (Heterogeneous rank-1 quotient graph permutation assignments are necessary). *If there does not exist a compatible set of permutation assignments for the heterogeneous rank-1 quotient graph, then there does not exist one for the rank-1 quotient graph.*

Proof. Suppose there exists a compatible set of permutation assignments $A = \{\Sigma_u\}$ for the rank-1 quotient graph. Let $A^h = \{\Sigma_u\} \subset A$ be the assignment of permutations to those vertices u which are also in the heterogeneous rank-1 quotient graph. A^h is a compatible set of permutation assignments for the heterogeneous rank-1 quotient graph. To see that this is true, one only needs to note that the only edges that are added in the construction of the heterogeneous rank-1 quotient graph are those which connect boundary vertices of connected components. For every pair of boundary vertices v and w of a connected component labelled by i , it follows from lemma 25 that there exists a j such that $\Pi(i) = j$ for all $\Pi \in \Sigma_v, \Sigma_w$, because there exists a path of edges labelled i connecting v and w in the rank-1 quotient graph. Thus A^h is indeed a compatible set of permutation assignments, and the proof follows contrapositively. \square

Lemma 27 (The heterogeneous rank-1 quotient graph translates faithfully to an Ising problem). *Given a heterogeneous rank-1 quotient graph with vertices which connect to edges with at most two labels. If there exists a compatible set of permutation assignments A^h , then there exists a solution to the Ising problem described in step 6, which when translated using table 1 will produce A^h .*

Proof. Suppose $A^h = \{\Sigma_u^h\}$ exists. For every edge in the heterogeneous rank-1 quotient graph, with end vertices u and v and label i , there exists a $j \neq 2$ s.t. for all $\Pi_u \in \Sigma_u^h$, $\Pi_v \in \Sigma_v^h$: $\Pi_u(i) = \Pi_v(i) = j$. Let u be an ix -vertex and v be an iy -vertex. Then $\Pi_u(x) \neq 2$ and

$\Pi_v(y) \neq 2$. Since $x \neq i$, there must exist a third index k so that $\Pi_u(k) = 2$, this uniquely specifies two possible permutations, given by the rows of table 1, similarly for Π_v . If $i \neq 2$ then both Π_u and Π_v must belong to the same column in table 1. If $i = 2$ then Π_u and Π_v must belong to distinct columns in table 1. We see that any choice of assignments of permutations to the vertices satisfying the above restrictions specify an exact satisfying solution to the Ising model described in step 6, given by the labelling of the columns in table 1. \square

Theorem 28. *If the algorithm returns false, the Hamiltonian is not stoquastic.*

Proof. There are six steps where the algorithm returns false. We shall address them individually.

Step 4: If a vertex in the rank-1 quotient graph connects to at least three edges, each with a different label, then any assignment of permutation to that vertex will leave one label equal to 2. Thus no compatible set of permutation assignments is possible, and by lemma 24 the Hamiltonian is not stoquastic.

Step 7: If there does not exist an exact satisfying solution to the Ising problem given in step 6, then by lemmas 27, 26, and 24 the Hamiltonian is not stoquastic.

Step 8: If both A_1^h and A_2^h contain permutations at some vertices which are not in the admissible permutations of those vertices, then they are not compatible sets of permutations assignments. Furthermore no others exist, by lemma 27, because an Ising problem has only two solutions. So by 26 and 24 the Hamiltonian is not stoquastic.

Step 9 edge case: Since, in the step 5 edge case, there is a label- i path between every vertex in the rank-1 quotient graph, by lemma 25 any set of compatible permutation assignments must be of the form A_1 or A_2 . Thus if both are incompatible with the admissible permutations of the vertices, then by lemma 24 the Hamiltonian is not stoquastic.

Step 9: By lemma 25, any assignment of possible permutations Σ_v for a vertex v in a single-label connected component with labelling i must agree with the action $\Pi(i)$ of any permutation $\Pi \in \Sigma_u$ of a boundary vertex u of that single-label connected component. If, for A_x^h the admissible permutations of v do not contain any permutations which agree in this way with the assignments of permutations to the boundary vertices, then no set of compatible permutation assignments exists on the rank-1 quotient graph which corresponds to A_x^h . If both A_1^h and A_2^h are ruled out in this way, then by lemmas 26 and 24 H is not stoquastic.

Step 14: A compatible set of permutation assignments A_x can not correspond to a possible solution

if, for all of the possible permutations producible by A_x , there does not exist an assignment of signs $\{R_u = \text{diag}(\pm 1, \pm 1, \pm 1)\}$ satisfying condition 17. Given a fixed choice of permutations $\{\Pi_u\}$, the question of whether there exists an assignment of signs is dependent only on the values $[\Pi_u \beta^{uv} \Pi_v^T]_{11}$ at each edge. Therefore the question of whether there exists an assignment of signs can be answered independently for each subset of vertices in the weighted interaction graph which do not connect to any vertices outside of the subset with an edge satisfying $[\Pi_u \beta^{uv} \Pi_v^T]_{11} \neq 0$. Since for a given A_x , any choice of permutation assignment $\{\Pi_u \in \Sigma_u^x\}$ will, by definition, have the same action on all low-rank edges, it suffices to consider each connected component of the partitioned weighted interaction graph independently. Steps 12 and 13 check each connected component of the partitioned weighted interaction graph in this fashion.

For each connected component of the partitioned weighted interaction graph one must rule out all possible permutation assignments one could draw from $A_x = \{\Sigma_v^x\}$ and assign to the vertices in that connected component. In the case of step 12, each single-label vertex corresponds to its own connected component, and there are only two possible permutations to consider. In the case of step 13, which covers the other connected components of the partitioned weighted interaction graph it may seem that there are an exponential number of possible assignments one could draw from. However each vertex u in one of these connected components necessarily corresponds to either a non-single label vertex v in the rank-1 quotient graph, for which Σ_v^x has only one element, or a single-label vertex v in the rank-1 quotient graph, labelled by i , for which any permutation $\Pi_v \in \Sigma_v^x$ satisfies $\Pi_v(i) = 1$. As such, the Ising model produced is independent of the choice of permutation assignment, so we choose a single representative.

By the reasoning outlined for the earlier steps, the only possible compatible sets of permutation assignments A_x are the ones produced in part 1. If none of the A_x correspond to a possible solution, then there is none and the Hamiltonian is not stoquastic. \square

Theorem 29. *When the algorithm returns true, along with a solution sol , the solution $sol = \{\Pi_u = R_u \Pi_u\}$ prescribes a set of single qubit unitary transformations $\{U_v\}$ such that $(\otimes_v U_v) H (\otimes_v U_v)^\dagger$ is a Z -matrix, and therefore H is stoquastic.*

Proof. As discussed in the preliminaries, single qubit unitary transformations correspond to $SO(3)$ rotations, and H is stoquastic by those rotations if conditions 14, 16, and 17 are satisfied. Therefore it suffices to show

that sol is a set of $SO(3)$ rotations satisfying those conditions. The construction of sol is primarily described in steps 12 and 13.

Conditions 14 and 16 are only determined by the assignment of permutations $\{\Pi_u\}$. The assignment of permutations is determined by a compatible set of permutation assignments $A = \{\Sigma_v\}$, so for every low-rank edge, conditions 14 and 16 will be satisfied by definition. For every rank >1 connected component Γ , associated with a vertex v in the rank-1 quotient graph, every vertex in Γ is assigned the same permutation $\Pi \in \Sigma_v$. Thus by the argument given in lemma 23, condition 14 is satisfied for all edges in Γ . Furthermore, since Σ_v only contains permutations compatible with the admissible permutations of the rank >1 connected component, 16 must also be satisfied. Thus all edges satisfy conditions 16 and 14.

For every connected component of the partitioned weighted interaction graph and every vertex u in the connected component, the sign assignments $R_u = \text{diag}(\delta_u, 1, \delta_u \det(\Pi_u))$ are determined by the solution $\{\delta_u\}$ to an Ising problem given by definition 20. The solution must satisfy $\delta_u (-\text{sign}([\Pi_u \beta^{uv} \Pi_v^T]_{11})) \delta_v = 1$. Therefore $[R_u \Pi_u \beta^{uv} \Pi_v^T R_v^T]_{11} \leq 0$ for every connected component in the partitioned weighted interaction graph. The only edges which do not belong to a connected component in the partitioned weighted interaction graph are those edges for which $[R_u \Pi_u \beta^{uv} \Pi_v^T R_v^T]_{11} = [\Pi_u \beta^{uv} \Pi_v^T]_{11} = 0$. Therefore for all edges, condition 17 is satisfied.

Finally, every Π_u is an $SO(3)$ transformation, since $\det(\Pi_u) = \det(R_u) \det(\Pi_u) = \delta_u^2 \det(\Pi_u)^2 = +1$. \square

4 Discussion

There remain many open questions with regards to stoquasticity. It would be interesting to capture the two-qubit (or two-qudit) stoquastic space in terms of inequalities involving invariants. It would be desirable if the problem of deciding whether a Hamiltonian is computationally stoquastic is easier when allowing more transformations than just local basis changes.

Many of the proofs in this paper are tedious and involved since the reasoning is different for different subcase Hamiltonians. We did not see a way of simplifying these arguments. It is an interesting question whether one could supply such proofs using a proof assistant, automatically ensuring their validity.

The presence of single-qubit terms causes our algorithm for deciding the stoquasticity of the XYZ Heisenberg model to fail. Furthermore, the existence of single-qubit terms makes the class of single-qubit Clifford rotations strictly weaker than that of single-qubit rotations.

Nevertheless, it seems likely that a generalization of the algorithm described here could be constructed to decide stoquasticity under single-qubit Clifford rotations of an XYZ model with single qubit terms included. It remains an open question whether it is computationally hard to decide whether such a Hamiltonian is stoquastic by local basis changes beyond single-qubit Clifford rotations.

We expect that the problem of deciding whether a Hamiltonian is stoquastic by local basis changes is easy for n -qubit Hamiltonians with an underlying line or tree interaction graph G by employing a combination of an efficient dynamic programming strategy with the parsimonious strategy (i.e. discretize the set of local basis changes, show how the optimization of a partial problem giving a basis change at the boundary can be used to solve a next-larger partial problem etc.).

An intriguing idea is that results showing hardness of deciding whether a Hamiltonian is stoquastic could be used for quantum computer verification. A general adiabatic computation with frustration-free Hamiltonians requires a quantum computer to run it, but let's assume that Alice, the verifier, picks a stoquastic frustration-free Hamiltonian, but then hides this stoquasticity by local basis changes or, say, the application of a constant depth circuit. Bob with a quantum computer can run the adiabatic computation and provide Alice with samples from its output which she can efficiently verify to be statistically correct using the algorithm in [9]. On the other hand, a QC-pretender Charlie may have to either find the constant-depth hiding circuit or classically simulate the quantum adiabatic algorithm, and either task could be too daunting. A similar idea has been suggested in [19].

5 Funding

We acknowledge support through the EU via the ERC GRANT EQEC No. 682726. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Economic Development & Innovation.

6 Acknowledgements

BMT would like to thank David DiVincenzo for interesting discussions. BMT and JDK would also like to thank Sergey Bravyi as well as Christophe Vuillot for helpful discussions.

References

- [1] S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal, "The Complexity of Stoquastic Local Hamiltonian Problems," *Quantum Information and Computation* **8** no. 5, (2008) 0361–0385, [arXiv:0606140 \[quant-ph\]](#).
- [2] N. J. Cerf and O. C. Martin, "Projection Monte Carlo methods : an algorithmic analysis," *International Journal of Modern Physics C* **6** no. 5, (1995) 693–723.
- [3] S. Sorella and L. Capriotti, "Green function Monte Carlo with stochastic reconfiguration: An effective remedy for the sign problem," *Physical Review B - Condensed Matter and Materials Physics* **61** no. 4, (2000) 2599–2612, [arXiv:9902211 \[cond-mat\]](#).
- [4] S. Bravyi, "Monte Carlo Simulation of Stoquastic Hamiltonians," *Quantum Information and Computation* **15** no. 13/14, (2015) 1122–1140, [arXiv:1402.2295](#).
- [5] S. Wessel, "Monte Carlo Simulations of Quantum Spin Models Institute for Theoretical Solid State Physics," in *Autumn School on Correlated Electrons*. 2013. <https://www.cond-mat.de/events/correl13/manuscripts/>.
- [6] E. Crosson, *Classical and Quantum Computation in Ground States and Beyond*. PhD thesis, University of Washington, 2015. <http://hdl.handle.net/1773/34128>.
- [7] S. Bravyi and D. Gosset, "Polynomial-Time Classical Simulation of Quantum Ferromagnets," *Physical Review Letters* **119** no. 10, (2017) , [arXiv:1612.05602](#).
- [8] T. Albash and D. A. Lidar, "Adiabatic quantum computation," *Reviews of Modern Physics* **90** no. 1, (2018) 015002, [arXiv:1611.04471 \[quant-ph\]](#).
- [9] S. Bravyi and B. Terhal, "Complexity of stoquastic frustration-free Hamiltonians," *SIAM J. Comput.* **39** no. 4, (2009) 1642, [arXiv:0806.1746](#).
- [10] M. B. Hastings and M. H. Freedman, "Obstructions To Classically Simulating The Quantum Adiabatic Algorithm," *Quantum Information and Computation* **13** no.11/12, (2013) 1038-1076 [arXiv:1302.5733](#).
- [11] J. Bringewatt, W. Dorland, S. P. Jordan, and A. Mink, "Diffusion Monte Carlo approach versus adiabatic computation for local Hamiltonians," *Phys. Rev. A* **97** no. 2, (Feb., 2018) 022323, [arXiv:1709.03971 \[quant-ph\]](#).

- [12] D. Kafri, C. Quintana, Y. Chen, A. Shabani, J. M. Martinis, and H. Neven, “Tunable inductive coupling of superconducting qubits in the strongly nonlinear regime,” *Physical Review A* **95** no. 5, (May, 2017) 052333, [arXiv:1606.08382](#).
- [13] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, “Nonstoquastic Hamiltonians and quantum annealing of an Ising spin glass,” *Phys. Rev. B* **95** no. 18, (May, 2017) 184416, [arXiv:1609.06558](#) [quant-ph].
- [14] G. Samach, “Tunable XX-Coupling Between High Coherence Flux Qubits,” in *APS March Meeting 2018*. 2018. <https://meetings.aps.org/Meeting/MAR18/Session/L33.13>.
- [15] V. I. Iglovikov, E. Khatami, and R. T. Scalettar, “Geometry dependence of the sign problem in quantum Monte Carlo simulations,” *Phys. Rev. B* **92** (Jul, 2015) 045110. <https://link.aps.org/doi/10.1103/PhysRevB.92.045110>.
- [16] C. Wu and S.-C. Zhang, “Sufficient condition for absence of the sign problem in the fermionic quantum monte carlo algorithm,” *Phys. Rev. B* **71** (Apr, 2005) 155115.
- [17] Z.-X. Li and H. Yao, “Sign-Problem-Free Fermionic Quantum Monte Carlo: Developments and Applications,” *arXiv e-prints* (May, 2018) [arXiv:1805.08219](#), [arXiv:1805.08219](#) [cond-mat.str-el].
- [18] R. F. Bishop and D. J. J. Farnell, “Marshall-Peierls sign rules, the quantum monte carlo method, and frustration,” *International Journal of Modern Physics B* **15** no. 10n11, (2001) 1736–1739.
- [19] M. Marvian, D. A. Lidar, and I. Hen, “On the Computational Complexity of Curing non-stoquastic Hamiltonians,” *Nature Communications* **10** no. 1, (2019) 1571, [arXiv:1802.03408](#).
- [20] B. M. Terhal, “The Power and Use of Stoquastic Hamiltonians,” in *Adiabatic Quantum Computing Conference*. 2017. <https://www.youtube.com/watch?v=4dK30QExF4M>.
- [21] T. Cubitt, A. Montanaro, and S. Piddock, “Universal Quantum Hamiltonians,” *National Academy of Sciences* **115** no. 38 (2018) 9497-9502, [arXiv:1701.05182](#).
- [22] S. Bravyi and M. Hastings, “On complexity of the quantum Ising model,” *Communications in Mathematical Physics* **349** no. 1 (2017) 1-45, [arXiv:1410.0703](#) [quant-ph].
- [23] D. Grier and L. Schaeffer, “The Classification of Stabilizer Operations over Qubits,” *ArXiv e-prints* (Mar., 2016), [arXiv:1603.03999](#) [quant-ph].
- [24] Y. Makhlin, “Nonlocal properties of two-qubit gates and mixed states and optimization of quantum computations,” *Quantum Information Processing* **1** no. 4, (2002) 243–252, [arXiv:0002045](#) [quant-ph].
- [25] N. Linden, S. Popescu, and A. Sudbery, “Nonlocal Parameters for Multiparticle Density Matrices,” *Physical Review Letters* **83** no. 2, (1999) 243–247, [arXiv:9801076](#) [quant-ph].
- [26] M. Grassl, M. Rötteler, and T. Beth, “Computing local invariants of quantum-bit systems,” *Phys. Rev. A* **58** (Sept., 1998) 1833–1839, [quant-ph/9712040](#).
- [27] R. A. Bertlmann and P. Krammer, “Bloch vectors for qudits,” *Journal of Physics A: Mathematical and Theoretical* **41** no. 23, (2008), [arXiv:0806.1174](#).
- [28] T. F. Gonzalez, “Clustering to minimize the maximum intercluster distance,” *Theoretical Computer Science* **38** (1985) 293–306.

A Symmetric Z-Matrix Cone in \mathbb{C}^{d^2}

It is understood that the set of symmetric Z-matrices forms a polyhedral cone. For convenience and concreteness we present the structure of this cone for operators acting on \mathbb{C}^{d^2} in terms of generalizations of the Pauli matrices.

Any Hermitian operator acting on $\mathbb{C}^d \otimes \mathbb{C}^d$ with basis elements $\{|i\rangle \mid i \in 0, \dots, d-1\}$, can be written as :

$$H = \sum_{x,y \in D,A,S} \sum_{i,j,m,n=0}^{d-1} a_{ij;mn}^{xy} \lambda_{ij}^x \otimes \lambda_{mn}^y,$$

with $a_{ij;mn}^{xy} \in \mathbb{R}$. Here the $\lambda_{ij}^{D,A,S}$ matrices with labels S (symmetric), A (anti-symmetric) and D (diagonal) are a modified set of generalized Gell-Mann matrices [27] given by:

$$\lambda_{ij}^S = \begin{cases} \sqrt{d/2} (|i\rangle\langle j| + |j\rangle\langle i|) & : i < j \\ 0 & : \text{otherwise} \end{cases} \quad (18)$$

$$\lambda_{ij}^A = \begin{cases} -i\sqrt{d/2} (|i\rangle\langle j| - |j\rangle\langle i|) & : i < j \\ 0 & : \text{otherwise} \end{cases} \quad (19)$$

$$\lambda_{ij}^D = \begin{cases} \frac{\sqrt{d}}{\sqrt{d-1}} e_0 + \sqrt{d} e_i & : i = j > 0 \\ \mathbb{I} & : i = j = 0 \\ 0 & : \text{otherwise} \end{cases} \quad (20)$$

with $e_i = |i\rangle\langle i| - \frac{1}{d}\mathbb{I}$. Note that $\text{Tr}[\lambda_{ij}^x \lambda_{mn}^y] = \delta_{im} \delta_{jn} \delta_{xy} d$, so that the λ matrices can be thought of as spanning vectors of the Hilbert-Schmidt vector space. We can thus employ the vector inner product of two bipartite matrices A and B acting on $\mathbb{C}^d \otimes \mathbb{C}^d$ as: $\langle A, B \rangle = \text{Tr}[A^\dagger B]/d^2$.

The necessary and sufficient conditions for H to be a symmetric Z -matrix (see Definition 1) are (1) $\forall i, j, m, n$ $\langle im|H|jn\rangle \in \mathbb{R}$ and (2) $\langle im|H|jn\rangle \leq 0$ when either $i \neq j$ or $m \neq n$ (or both).

The λ -basis is convenient because of the following properties. The diagonal part of H only has support on vectors of the form $\lambda^D \otimes \lambda^D$ while the off-diagonal part has no contribution from $\lambda^D \otimes \lambda^D$. Furthermore, the imaginary part of H only has support on vectors of the form $\lambda^A \otimes \lambda^S$, $\lambda^A \otimes \lambda^D$, $\lambda^S \otimes \lambda^A$ and $\lambda^D \otimes \lambda^A$, while both the real and the diagonal part of H have no contribution from these vectors. Thus the vector H has support on three orthogonal subspaces, the diagonal subspace D spanned by vectors $\lambda^D \otimes \lambda^D$, the imaginary subspace I , spanned by vectors $\lambda^A \otimes \lambda^S$, $\lambda^A \otimes \lambda^D$, $\lambda^S \otimes \lambda^A$ and $\lambda^D \otimes \lambda^A$, and the off-diagonal real subspace R spanned by the remaining vectors.

We may write $H = H_D + H_R + H_I$ where H_D, H_R, H_I are the vectors in these three subspaces.

The first (“realness”) condition thus corresponds to demanding that $H_I = 0$. By the linear independence of the λ -basis it implies that

$$\forall i, j, m, n, a_{ij;mn}^{AS} = a_{ij;mn}^{SA} = a_{ij;mn}^{AD} = a_{ij;mn}^{DA} = 0. \quad (21)$$

The second (“negativity”) condition then says that for all $i \neq j$ or $m \neq n$ (or both):

$$\langle im|H|jn\rangle = \text{Tr}[|jn\rangle\langle im|H] = d^2 \langle |jn\rangle\langle im|, H_R \rangle \leq 0.$$

This condition thus specifies some of the facets of the symmetric Z -matrix cone. When a vector $|im\rangle\langle jn|$ is interpreted as the normal vector of an oriented hyperplane in Hilbert-Schmidt space, the negativity condition can be thought of as the requirement that H_R lies within the union of half spaces defined by these oriented hyperplanes.

To write down the inequalities describing these facets, we can use that

$$|j\rangle\langle i| = \begin{cases} \frac{1}{2}(\lambda_{ij}^S - i\lambda_{ij}^A) & : i < j, \\ \frac{1}{2}(\lambda_{ji}^S + i\lambda_{ji}^A) & : i > j, \\ e_i + \frac{1}{d}\mathbb{I} & : i = j, \end{cases}$$

with

$$e_0 = \frac{1}{d} \sum_{m=1}^{d-1} \lambda_m^D, \quad e_{i \neq 0} = \frac{1}{\sqrt{d}} \lambda_i^D - \frac{1}{d(\sqrt{d}-1)} \sum_{m=1}^{d-1} \lambda_m^D.$$

One can then write down three types of inequalities:

- case 1: $i \neq j$ and $m \neq n$

$$\begin{aligned} & \langle |jn\rangle\langle im|, H_R \rangle \\ &= \begin{cases} \frac{1}{4} (a_{ij;mn}^{SS} - a_{ij;mn}^{AA}) \leq 0 & : \begin{matrix} (i < j, m < n) \\ (i > j, m > n) \end{matrix} \\ \frac{1}{4} (a_{ij;mn}^{SS} + a_{ij;mn}^{AA}) \leq 0 & : \begin{matrix} (i > j, m < n) \\ (i < j, m > n) \end{matrix} \end{cases} \end{aligned} \quad (22)$$

- case 2: $i = j$ and $m \neq n$

$$\begin{aligned} & \langle |in\rangle\langle im|, H_R \rangle = \frac{1}{2} \langle (e_i + \mathbb{I}/d) \otimes \lambda_{mn}^S, H_R \rangle \\ &= \frac{1}{2} (\langle e_i \otimes \lambda_{mn}^S, H_R \rangle + a_{00;mn}^{DS}/d) \leq 0 \end{aligned} \quad (23)$$

- case 3: $i \neq j$ and $m = n$

$$\begin{aligned} & \langle |jm\rangle\langle im|, H_R \rangle = \frac{1}{2} \langle \lambda_{ij}^S \otimes (e_m + \mathbb{I}/d), H_R \rangle \\ &= \frac{1}{2} (\langle \lambda_{ij}^S \otimes e_m, H_R \rangle + a_{ij;00}^{SD}/d) \leq 0 \end{aligned} \quad (24)$$

There is an interesting geometric fact here. Consider for instance Inequality (23). For a fixed n and $m \neq n$ Inequality (23) is only concerned with the elements of the vector H_R with support on basis elements of the form $\lambda^D \otimes \lambda_{mn}^S$, which do not appear in any other inequalities. So for a fixed n and $m \neq n$ we may think of the set of inequalities as living in a d -dimensional subspace of Hilbert-Schmidt space orthogonal to all other sets of inequalities. The vectors $|i\rangle\langle i| \otimes |n\rangle\langle m|_R = (e_i + \frac{\mathbb{I}}{d}) \otimes \lambda_{mn}^S$ are the normal vectors of the supporting hyperplanes of our convex set in this space, and the vectors $\{e_m\}$ correspond to the extremal points of a regular simplex. Fixing a value for $\langle \frac{\mathbb{I}}{d} \otimes \lambda_{mn}^S, H_R \rangle = \frac{a_{ij;00}^{SD}}{d}$ we can take a $d-1$ -dimensional slice of our convex set and see that the vectors $e_i \otimes \lambda_{mn}^S$ form the normals of a new supporting hyperplane whose distances from the origin of the slice is proportional to our value of $\frac{a_{ij;00}^{SD}}{d}$. Thus inequalities (23) and (24) give the polyhedral cone a simplicial structure.

This is illustrated in figure 3 for $d = 2, 3$.

In the case of two qubits $\lambda_{01}^S = X$, $\lambda_{01}^A = Y$, $\lambda_{01}^D = Z$, $e_0 = Z/2$ and $e_1 = -Z/2$. For the realness condition we require that $a_{IY} = a_{YI} = a_{YZ} = a_{ZY} = a_{YX} = a_{XY} = 0$. For the negativity condition we retrieve the inequalities:

$$\begin{aligned} a_{XX} &\leq a_{YY} & a_{XX} &\leq -a_{YY} \\ a_{IX} &\leq -a_{ZX} & a_{IX} &\leq a_{ZX} \\ a_{XI} &\leq -a_{XZ} & a_{XI} &\leq a_{XZ}. \end{aligned}$$

Together these (in)equalities are given in condensed form in Proposition 6.




Figure 3: An illustration of the simplicial polyhedral cones generated by the inequalities (23) or (24) for $d = 2$ and $d = 3$. For example in the case of Inequality (23), fixing n and $m \neq n$ gives solutions lying inside a cone bounded by the set of hyperplanes defined by the normal vectors $(|i\rangle\langle i| \otimes |n\rangle\langle m|)_R$ (labelled $|i\rangle\langle i|$ in the figures). See for example in the $d = 3$ figure the blue plane defined by the normal vector $(|0\rangle\langle 0| \otimes |n\rangle\langle m|)_R$. Note that the cone lies in the lower half of the space, for example the region where $\langle \frac{1}{d} \otimes \lambda_{mn}^S, H_R \rangle \leq 0$ in the case of Inequality (23).

B Proof of Theorem 7

Recall the definitions of

$$\Gamma_L = \{S, \beta\beta^T S, [\beta\beta^T]^2 S, \beta P, \beta\beta^T \beta P\},$$

and

$$\Gamma_R = \{P, \beta^T \beta P, [\beta^T \beta]^2 P, \beta^T S, \beta^T \beta\beta^T S\}.$$

When the triple product invariants are zero, Γ_L and Γ_R each contain coplanar vectors. By Proposition 30 (see below) we know that S is in the span of at most two left singular vectors of β and P is in the span of at most two right singular vectors of β .

If S , $\beta\beta^T S$ and $[\beta\beta^T]^2 S$ are linearly independent, then βP must be in the same span as S . If they are not linearly independent then S is a left singular vector of β , with unit vector \hat{e}_S . It follows that $\beta P \in \text{span}(\hat{e}_S, \hat{e}_{S_\perp})$, and $\beta\beta^T \beta P \in \text{span}(\hat{e}_S, \hat{e}_{S_\perp})$. This implies that $\beta\beta^T \hat{e}_{S_\perp} \in \text{span}(\hat{e}_S, \hat{e}_{S_\perp})$. Since $\hat{e}_S^T \beta\beta^T \hat{e}_{S_\perp} = 0$ it follows that $\beta\beta^T \hat{e}_{S_\perp} \in \text{span}(\hat{e}_{S_\perp})$ and \hat{e}_{S_\perp} is a left singular vector of β . Therefore there always exists a pair of left singular vectors of β such that S and βP are in their span. By Proposition 31 (see below) this implies that S is spanned by at most two left singular vectors of β and P is spanned by the corresponding right singular vectors. By Proposition 32 (see below) we see that H is real under local unitary rotations. \square

Proposition 30. *If S , $\beta\beta^T S$ and $[\beta\beta^T]^2 S$ are coplanar, then $\exists \hat{e}_x^L, \hat{e}_y^L$ s.t. $S \in \text{span}(\hat{e}_x^L, \hat{e}_y^L)$ with \hat{e}_x^L, \hat{e}_y^L left singular vectors of β . Similarly for P .*

Proof. Consider the vectors S , $\beta\beta^T S$ and $[\beta\beta^T]^2 S$. Consider the expansion of S into an orthonormal set

of left singular vectors of β : $S = \sum_i^3 S_i \hat{e}_i^L$. If $\beta\beta^T$ is not degenerate, then S , $\beta\beta^T S$ and $[\beta\beta^T]^2 S$ are coplanar only if $\exists i$ s.t. $S_i = 0$, which implies that $S \in \text{span}(\hat{e}_j^L, \hat{e}_k^L)$. If $\beta\beta^T$ is two fold degenerate then $\exists j, k$ s.t. $S_j \hat{e}_j^L + S_k \hat{e}_k^L$ is a left singular vector of $\beta\beta^T$, which implies that $S \in \text{span}(\hat{e}_i^L, \hat{e}_{jk}^L)$ with \hat{e}_{jk}^L a left singular vector of β . Finally, if $\beta\beta^T$ is three fold degenerate, then S is a left singular vector of β . Therefore, regardless of β , $\exists \hat{e}_x^L, \hat{e}_y^L$ s.t. $S \in \text{span}(\hat{e}_x^L, \hat{e}_y^L)$ with \hat{e}_x^L, \hat{e}_y^L left singular vectors of β . A symmetric argument holds for P . \square

Proposition 31. *If Γ_L and Γ_R each contain coplanar vectors, then if there exist a pair of left singular vectors \hat{e}_x^L, \hat{e}_y^L such that $\beta P, S \in \text{span}(\hat{e}_x^L, \hat{e}_y^L)$ then the right singular vectors which span P correspond to the left singular vectors which span S .*

Proof. Let $S \in \text{span}(\hat{e}_x^L, \hat{e}_y^L)$. Assume βP is in the same span as S , then $\beta^T \beta P, [\beta^T \beta]^2 P \in \text{span}(\hat{e}_x^R, \hat{e}_y^R)$. If $\beta^T \beta P$ and $[\beta^T \beta]^2 P$ are not proportional to one another then $P \in \text{span}(\beta^T \hat{e}_x^L, \beta^T \hat{e}_y^L)$. If $\beta^T \beta P$ and $[\beta^T \beta]^2 P$ are proportional to one another then they are proportional to a right singular vector of β , \hat{e}_P^R . If β is full rank then $\beta^T \beta$ is invertible and $P \propto \hat{e}_P^R \in \text{span}(\hat{e}_x^R, \hat{e}_y^R)$. If β is not full rank and $P \not\propto \hat{e}_P^R$ then $P \in \text{span}(\hat{e}_P^R, \hat{e}_0^R)$ with $\beta \hat{e}_0^R = 0$. Consider the left singular vector expansion of S into an orthonormal basis: $S = S_P \hat{e}_P^L + S_0 \hat{e}_0^L + S_\perp \hat{e}_\perp^L$. We require the following three vectors to be coplanar

$$\beta^T S \in \text{span}(\hat{e}_P^R, \hat{e}_0^R)$$

$$P \in \text{span}(\hat{e}_P^R, \hat{e}_0^R)$$

$$\beta^T \beta P \in \text{span}(\hat{e}_P^R).$$

The only way these three vectors can be coplanar is if $S_\perp = 0$ or $\beta^T \hat{e}_\perp^L = 0$. If β is rank 1 and $\beta^T \hat{e}_\perp^L = 0$ then $S_0 \hat{e}_\perp^L + S_\perp \hat{e}_\perp^L$ constitutes a left singular vector corresponding to \hat{e}_0^R , since $\beta \hat{e}_0^R = 0 (S_0 \hat{e}_\perp^L + S_\perp \hat{e}_\perp^L)$ and $\beta^T (S_0 \hat{e}_\perp^L + S_\perp \hat{e}_\perp^L) = 0 \hat{e}_0^R$ and so S and P share a common pair of singular vectors. If β is rank 2 then $S_\perp = 0$ and S and P share common left right singular vectors \hat{e}_P and \hat{e}_0 . \square

Proposition 32. *Given a two-qubit Hamiltonian H parametrized by β , S and P . H is real under local unitary rotations iff S is in the span of at most two left singular vectors of β and P is in the span of the corresponding right singular vectors.*

Proof. If H is real under local unitary rotations, then there exists a pair of rotations $O_L, O_R \in SO(3)$ s.t.

$$O_L \beta O_R^T = \beta' = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix},$$

$$O_L S = S' = \begin{pmatrix} S'_1 \\ 0 \\ S'_3 \end{pmatrix},$$

$$O_R P = P' = \begin{pmatrix} P'_1 \\ 0 \\ P'_3 \end{pmatrix}.$$

Expressed in the left and right basis $\{\hat{e}_i^L\}$ and $\{\hat{e}_i^R\}$ respectively. The left singular vectors of β are $\{O_L^T \hat{e}_i^L\}$ and the corresponding right singular vectors are $\{\text{sign}(\lambda_i) O_R^T \hat{e}_i^R\}$.

It is easy to see that S and P are spanned by at most two left singular vectors of β

$$S = S'_1 O_L^T \hat{e}_1^L + S'_3 O_L^T \hat{e}_3^L$$

$$P = \text{sign}(\lambda_1) P'_1 (\text{sign}(\lambda_1) O_R^T \hat{e}_1^R) + \text{sign}(\lambda_3) P'_3 (\text{sign}(\lambda_3) O_R^T \hat{e}_3^R)$$

That concludes the proof for one direction of the bi-conditional.

Suppose S and P are spanned by two or fewer left and right (respectively) singular vectors of β which share the same singular values. Then S and P are expressible as

$$S = S_1 \hat{e}_1^L + S_3 \hat{e}_3^L$$

$$P = P_1 \hat{e}_1^R + P_3 \hat{e}_3^R$$

with

$$\beta^T \hat{e}_i^L = |\lambda_i| \hat{e}_i^R$$

and

$$\beta \hat{e}_i^R = |\lambda_i| \hat{e}_i^L$$

There exists a pair of rotations $O_L, O_R \in SO(3)$ s.t.

$$O_L \beta O_R^T = \beta' = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}$$

with

$$S' = O_L S = \begin{pmatrix} S_1 \\ 0 \\ S_3 \end{pmatrix}$$

and

$$P' = O_R P = \begin{pmatrix} \text{sign}(\lambda_1) P_1 \\ 0 \\ \text{sign}(\lambda_3) P_3 \end{pmatrix}$$

which corresponds to a real Hamiltonian. Thus H is real under local unitary rotations. \square

C Polynomial Inequalities for Stoquasticity of a 2-qubit Hamiltonian

Given the standard form described in section 2.4, rotated by $SO(3)$ rotations given by equation 10, the rotated two-qubit Hamiltonian is described by:

$$\beta' = \begin{pmatrix} a'_{XX} & 0 & a'_{XZ} \\ 0 & a_{YY} & 0 \\ a'_{ZX} & 0 & a'_{ZZ} \end{pmatrix} \quad (25)$$

with

$$\begin{aligned} a'_{ZZ} &= \gamma_L \gamma_R (\cos(\theta_L) \cos(\theta_R) + a_{XX} \sin(\theta_L) \sin(\theta_R)) \\ a'_{ZX} &= \gamma_L (\cos(\theta_L) \sin(\theta_R) - a_{XX} \cos(\theta_R) \sin(\theta_L)) \\ a'_{XZ} &= \gamma_R (\cos(\theta_R) \sin(\theta_L) - a_{XX} \cos(\theta_L) \sin(\theta_R)) \\ a'_{XX} &= \sin(\theta_L) \sin(\theta_R) + a_{XX} \cos(\theta_L) \cos(\theta_R) \\ a'_{YY} &= a_{YY} \gamma_R \gamma_L \end{aligned}$$

and

$$S' = \begin{pmatrix} a_{XI} \cos(\theta_L) + a_{ZI} \sin(\theta_L) \\ 0 \\ a_{ZI} \gamma_L \cos(\theta_L) - a_{XI} \gamma_L \sin(\theta_L) \end{pmatrix} \quad (26)$$

$$P' = \begin{pmatrix} a_{IX} \cos(\theta_R) + a_{IZ} \sin(\theta_R) \\ 0 \\ a_{IZ} \gamma_R \cos(\theta_R) - a_{IX} \gamma_R \sin(\theta_R) \end{pmatrix} \quad (27)$$

Three inequalities must be satisfied in order for H to be a symmetric Z-matrix:

$$\sin(\theta_L) \sin(\theta_R) + a_{XX} \cos(\theta_L) \cos(\theta_R) \leq -|a_{YY}| \quad (28)$$

$$\begin{aligned} a_{IX} \cos(\theta_R) + a_{IZ} \sin(\theta_R) \leq \\ -|\cos(\theta_L) \sin(\theta_R) - a_{XX} \cos(\theta_R) \sin(\theta_L)| \end{aligned} \quad (29)$$

$$\begin{aligned} a_{XI} \cos(\theta_L) + a_{ZI} \sin(\theta_L) \leq \\ -|\cos(\theta_R) \sin(\theta_L) - a_{XX} \cos(\theta_L) \sin(\theta_R)| \end{aligned} \quad (30)$$

Note that the parameters γ_1 and γ_2 have fallen out of these inequalities.

If the above inequalities can be satisfied, then the Hamiltonian is stoquastic. If they cannot be satisfied then the Hamiltonian is not stoquastic, barring a single case as follows. If either $\{a_{XI}, a_{IX}\} = \{0, 0\}$ or $\{a_{ZI}, a_{IZ}\} = \{0, 0\}$, then an additional test of the same form must be applied after first performing a permutation such that the pair of zero coefficients are now in the Y position. For example, if $\{a_{XI}, a_{IX}\} = \{0, 0\}$ then one must also perform the above test on the diagonal form:

$$\beta = \begin{pmatrix} a_{YY} & 0 & 0 \\ 0 & a_{XX} & 0 \\ 0 & 0 & a_{ZZ} \end{pmatrix}$$

$$S = \begin{pmatrix} 0 \\ 0 \\ a_{ZI} \end{pmatrix} \quad P = \begin{pmatrix} 0 \\ 0 \\ a_{IZ} \end{pmatrix}$$

The inequalities (28), (29) and (30) can be broken up into six inequalities in order to ignore the absolute values:

- (1a) $\sin(\theta_L)\sin(\theta_R) + a_{XX}\cos(\theta_L)\cos(\theta_R) \leq 0$
- (1b) $(\sin(\theta_L)\sin(\theta_R) + a_{XX}\cos(\theta_L)\cos(\theta_R))^2 \geq a_{YY}^2$
- (2a) $a_{IX}\cos(\theta_R) + a_{IZ}\sin(\theta_R) \leq 0$
- (2b) $(a_{IX}\cos(\theta_R) + a_{IZ}\sin(\theta_R))^2 \geq (\cos(\theta_L)\sin(\theta_R) - a_{XX}\cos(\theta_R)\sin(\theta_L))^2$
- (3a) $a_{XI}\cos(\theta_L) + a_{ZI}\sin(\theta_L) \leq 0$
- (3b) $(a_{XI}\cos(\theta_L) + a_{ZI}\sin(\theta_L))^2 \geq (\cos(\theta_R)\sin(\theta_L) - a_{XX}\cos(\theta_L)\sin(\theta_R))^2$

We wish to map these trigonometric inequalities into polynomial inequalities of two variables by defining $x_1 = \tan(\theta_L)$ and $x_2 = \tan(\theta_R)$. However the mapping will depend on the value of $\cos(\theta_1)$ and $\cos(\theta_2)$, and so breaks up into several cases.

Case 1: $\cos(\theta_L), \cos(\theta_R) \neq 0$

Let $\delta_L = \text{Sign}(\cos(\theta_L))$ and $\delta_R = \text{Sign}(\cos(\theta_R))$ then

- (1a) $\delta_L\delta_R(x_1x_2 + a_{XX}) \leq 0$
- (1b) $(x_1x_2 + a_{XX})^2 \geq a_{YY}^2(1+x_1^2)(1+x_2^2)$
- (2a) $\delta_R(a_{IX} + a_{IZ}x_2) \leq 0$
- (2b) $(a_{IX} + a_{IZ}x_2)^2(1+x_1^2) \geq (x_2 - a_{XX}x_1)^2$
- (3a) $\delta_L(a_{XI} + a_{ZI}x_1) \leq 0$
- (3b) $(a_{XI} + a_{ZI}x_1)^2(1+x_2^2) \geq (x_1 - a_{XX}x_2)^2$

Case 2: $\cos(\theta_L) = 0, \cos(\theta_R) \neq 0$

Let $\delta_L = \text{Sign}(\sin(\theta_L))$ and $\delta_R = \text{Sign}(\cos(\theta_R))$ then

- (1a) $\delta_L\delta_Rx_2 \leq 0$
- (1b) $x_2^2/(1+x_2^2) \geq a_{YY}^2$
- (2a) $\delta_R(a_{IX} + a_{IZ}x_2) \leq 0$
- (2b) $(a_{IX} + a_{IZ}x_2)^2 \geq a_{XX}^2$
- (3a) $\delta_La_{ZI} \leq 0$
- (3b) $(1+x_2^2)a_{ZI}^2 \geq 1$

Case 3: $\cos(\theta_L) \neq 0, \cos(\theta_R) = 0$

Let $\delta_L = \text{Sign}(\cos(\theta_L))$ and $\delta_R = \text{Sign}(\sin(\theta_R))$ then

- (1a) $\delta_L\delta_Rx_1 \leq 0$
- (1b) $x_1^2/(1+x_1^2) \geq a_{YY}^2$
- (2a) $\delta_Ra_{IZ} \leq 0$
- (2b) $(1+x_1^2)a_{IZ}^2 \geq 1$
- (3a) $\delta_L(a_{XI} + a_{ZI}x_1) \leq 0$
- (3b) $(a_{XI} + a_{ZI}x_1)^2 \geq a_{XX}^2$

Case 4: $\cos(\theta_L) = 0, \cos(\theta_R) = 0$

Let $\delta_L = \text{Sign}(\sin(\theta_L))$ and $\delta_R = \text{Sign}(\sin(\theta_R))$ then

- (1a) $\delta_L\delta_R \leq 0$
- (1b) $1 \geq a_{YY}^2$
- (2a) $\delta_Ra_{IZ} \leq 0$
- (3a) $\delta_La_{ZI} \leq 0$

All of these inequalities are at most quadratic in any single variable x_1 or x_2 , so analytic solutions to their roots can be constructed and solutions can be found using graphical methods. If there exist values of x_1 and x_2 such that for some choice of $\delta_L, \delta_R \in \{-1, 1\}$ at least one of the above cases is satisfied, then H is stoquastic. If not, then we may say that H is not stoquastic save the exception described above.

This concludes the complete characterization of stoquasticity of a two-qubit Hamiltonian H .

D Proof of Lemma 21

Consider a set of 3×3 matrices $\{\beta^{uv}\}$ indexed by the pair uv which are diagonal in some fixed basis $\{\hat{e}_1, \hat{e}_2, \hat{e}_3\}$. Here we establish that whenever there exists a set of $SO(3)$ rotations $\{O_u\}$ mapping to a new set of diagonal matrices $\{O_u\beta^{uv}O_u^T\}$, there also exists a set of signed permutations $\{\Pi_u\}$, each with determinant 1, performing the same effective mapping: $\{O_u\beta^{uv}O_u^T\} = \{\Pi_u\beta^{uv}\Pi_u^T\}$, where by signed permutations we mean matrices of the form $\Pi_u = \Pi_u R_u$, with Π_u a permutation and R_u a diagonal matrix with diagonal elements ± 1 . Therefore whenever considering the class of sets of $SO(3)$ transformations which preserve the diagonality of the matrices $\{\beta^{uv}\}$, it suffices to consider only the signed permutations. Of course this is trivially true when considering a single diagonal matrix, but is less obvious when considering a set of matrices.

The proof is achieved by introducing a consistent mapping from a given $SO(3)$ matrix O to a signed permutation $\Pi(O)$, which we call the Π -reduction. This

mapping will have the property that for a diagonal matrix β , if $O_L \beta O_R^T$ is diagonal, then $O_L \beta O_R^T = \Pi(O_L) \beta \Pi(O_R)^T$. Since the mapping is consistent, it also holds when considering the action of a set of $SO(3)$ rotations $\{O_u\}$ on a set of diagonal matrices $\{\beta^{uv}\}$.

Let us first define the Π -reduction. Consider the determinant of an $SO(3)$ rotation in a particular fixed basis $\{\hat{e}_m\}$:

$$|O| = \left| \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} \right| = aei - afh - bdi + bfg + cdh - ceg$$

We can associate each of the monomials in the right hand side of the above equation with a matrix which is zero for all terms in the matrix which do not contribute to that monomial, and $\text{sign}(x)$ for all terms x which do contribute to that monomial. So for example:

$$afh \leftrightarrow \begin{pmatrix} \text{sign}(a) & 0 & 0 \\ 0 & 0 & \text{sign}(f) \\ 0 & \text{sign}(h) & 0 \end{pmatrix}$$

We may also impose a lexicographical ordering on the matrices, so that the matrix associated with bdi is lower in the ordering than the matrix associated with bfg . This is to ensure that the D -reduction has a consistent definition, and the particular choice of ordering is not important. We may then define the Π -reduction as follows:

Definition 33. Given a fixed basis $\{\hat{e}_m\}$, the Π -reduction $\Pi(O)$ of an $SO(3)$ operator O is a signed permutation associated with the lowest monomial, by lexicographical ordering, in the determinant polynomial of O . For all entries in O which do not contribute to the monomial, the corresponding entries in $\Pi(O)$ are 0, and for all entries x in O which do contribute to the monomial, the corresponding entries in $\Pi(O)$ are $\text{sign}(x)$.

Note that the determinant of O is $+1$, so there always exists a non-zero monomial. Finally, note that the Π -reduction is always a signed permutation, but not always an $SO(3)$ matrix.

Let $\beta = \sum_m \lambda_m \hat{e}_m \hat{e}_m^T$ and $\tilde{\beta} = \sum_m \tilde{\lambda}_m \hat{e}_m \hat{e}_m^T$.

Lemma 34. If $O_L \beta O_R^T = \tilde{\beta}$ and $|\lambda_m| \neq |\tilde{\lambda}_n| \rightarrow [O_L]_{nm} = [O_R]_{nm} = 0$.

Proof.

$$\hat{e}_n^T \tilde{\beta}^2 O_R \hat{e}_m = \tilde{\lambda}_n^2 [O_R]_{nm}$$

$$\hat{e}_n^T O_L \beta^2 \hat{e}_m = \lambda_m^2 [O_R]_{nm}$$

Since $\tilde{\beta} = \tilde{\beta}^T$ and $\beta = \beta^T$ we have:

$$\hat{e}_n^T \tilde{\beta}^2 O_R \hat{e}_m = \hat{e}_n^T O_L \beta^2 \hat{e}_m.$$

Therefore

$$\tilde{\lambda}_n^2 [O_R]_{nm} = \lambda_m^2 [O_R]_{nm}.$$

A symmetric argument can be given to show that

$$\tilde{\lambda}_n^2 [O_L]_{nm} = \lambda_m^2 [O_L]_{nm}.$$

Therefore if $|\lambda_m| \neq |\tilde{\lambda}_n|$ then $[O_L]_{nm} = [O_R]_{nm} = 0$. \square

Lemma 35. If $O_L \beta O_R^T = \tilde{\beta}$ and $\lambda_m \neq 0 \rightarrow \forall n [O_L]_{nm} = \text{sign}(\lambda_m) \text{sign}(\tilde{\lambda}_n) [O_R]_{nm}$.

Proof.

$$\hat{e}_n^T \tilde{\beta} O_R^T \hat{e}_m = \tilde{\lambda}_n [O_R]_{nm}$$

$$\hat{e}_n^T O_L \tilde{\beta} \hat{e}_m = \lambda_m [O_L]_{nm}$$

$$\hat{e}_n^T \tilde{\beta} O_R^T \hat{e}_m = \hat{e}_n^T O_L \tilde{\beta} \hat{e}_m$$

therefore

$$\tilde{\lambda}_n [O_R]_{nm} = \lambda_m [O_L]_{nm}.$$

If $[O_R]_{nm} = 0$ and $\lambda_m \neq 0$ then $[O_L]_{nm} = 0$ and

$$[O_L]_{nm} = \text{sign}(\lambda_m) \text{sign}(\tilde{\lambda}_n) [O_R]_{nm}.$$

Otherwise, if $[O_R]_{nm} \neq 0$ then by lemma 34 $|\lambda_m| = |\tilde{\lambda}_n|$ and if $\lambda_m \neq 0$ then

$$[O_L]_{nm} = \text{sign}(\lambda_m) \text{sign}(\tilde{\lambda}_n) [O_R]_{nm}.$$

\square

Lemma 36. Let $\Pi_L = \Pi(O_L)$ and $\Pi_R = \Pi(O_R)$. If $O_L \beta O_R^T = \tilde{\beta}$ and $\lambda_m \neq 0$ then $[\Pi_L]_{nm} = \text{sign}(\lambda_m) \text{sign}(\tilde{\lambda}_n) [\Pi_R]_{nm}$.

Proof. If $\lambda_m \neq 0$ then by lemma 35 $[O_L]_{nm} = \text{sign}(\lambda_m) \text{sign}(\tilde{\lambda}_n) [O_R]_{nm}$. It follows that if $[\Pi_L]_{nm} \neq 0$ and $[\Pi_R]_{nm} \neq 0$ then $[\Pi_L]_{nm} = \text{sign}(\lambda_m) \text{sign}(\tilde{\lambda}_n) [\Pi_R]_{nm}$. Furthermore, if $[\Pi_L]_{nm} = 0$ and $[\Pi_R]_{nm} = 0$ then $[\Pi_L]_{nm} = \text{sign}(\lambda_m) \text{sign}(\tilde{\lambda}_n) [\Pi_R]_{nm}$. So it suffices to show that if $\lambda_m \neq 0$ then $[\Pi_L]_{nm} \neq 0$ if and only if $[\Pi_R]_{nm} \neq 0$.

Given a $\lambda_m \neq 0$, we may consider three cases, depending on the number of non-zero entries in the m th column of O_L . Note that every column of an $SO(3)$ matrix has at least one non-zero element, since otherwise it would have determinant 0.

Case 1: The m th column of O_L has one non-zero entry. So there must exist a single $n \in \{1, 2, 3\}$ such that $[O_L]_{nm} \neq 0$ and, by lemma 35, $[O_R]_{nm} \neq 0$. Furthermore, if a column or row of O_L has a single non-zero entry, then $\Pi(O_L)$ must have a non-zero entry at that same position. This can be seen by noting that every monomial in the determinant polynomial of O_L must contain the non-zero entry $[O_L]_{nm}$, and so any matrix

returned by $\Pi(O_u)$ must have a non-zero entry at that same position. Therefore $[\Pi_L]_{nm} \neq 0$ and $[\Pi_R]_{nm} \neq 0$.

Case 2: The m th column of O_L has two non-zero entries. So there must exist $n_1, n_2 \in \{1, 2, 3\}$ such that $[O_L]_{n_1 m} \neq 0, [O_L]_{n_2 m} \neq 0$. By lemma 35, $[O_R]_{n_1 m} \neq 0, [O_R]_{n_2 m} \neq 0$. Furthermore, there exists $n_3 \in \{1, 2, 3\}$ such that $[O_L]_{n_3 m} = 0$ and, by lemma 35, $[O_R]_{n_3 m} = 0$. In other words O_L and O_R have the same block form, one 2×2 block, and one 1×1 block. All elements in the blocks must be non-zero. This implies O_L and O_R have the same non-zero monomials in their determinant polynomial, and thus have the same smallest non-zero monomial. Therefore $[\Pi_L]_{nm} \neq 0$ if and only if $[\Pi_R]_{nm} \neq 0$.

Case 3: The m th column of O_L has three non-zero entries. By lemma 34 $\lambda_n \neq 0$ for all n . Therefore by lemma 35 $[O_L]_{nm} = \text{sign}(\lambda_m)\text{sign}(\tilde{\lambda}_n)[O_R]_{nm}$ for all n and m . This implies that O_L and O_R will always have the same non-zero monomials in their determinant polynomials. Therefore $[\Pi_L]_{nm} \neq 0$ if and only if $[\Pi_R]_{nm} \neq 0$. \square

Lemma 37. *If $O_L \beta O_R^T = \tilde{\beta}$ then $\Pi(O_L)\beta\Pi(O_R)^T = \tilde{\beta}$ Proof.* Let $\Pi_x = \Pi(O_x)$.

$$\Pi_L \beta \Pi_R^T = \sum_{nn'} \sum_m \lambda_m [\Pi_L]_{nm} [\Pi_R]_{n'm} \hat{e}_n \hat{e}_n^T$$

By lemma 36, if $\lambda_m \neq 0$ and $[\Pi_L]_{nm} \neq 0$ then $[\Pi_R]_{nm} \neq 0$. Since for any m $[\Pi_R]_{nm}$ may only be nonzero for one n it follows that if $\lambda_m \neq 0$ and $[\Pi_L]_{nm} \neq 0$ then $[\Pi_R]_{n'm} = [\Pi_R]_{nm} \delta_{nn'} = S(\lambda_m) S(\tilde{\lambda}_n) [\Pi_L]_{nm} \delta_{nn'}$. Therefore

$$\Pi_L \beta \Pi_R^T = \sum_n \sum_m \text{sign}(\lambda_m) \text{sign}(\tilde{\lambda}_n) \lambda_m ([\Pi_L]_{nm})^2 \hat{e}_n \hat{e}_n^T$$

$$\Pi_L \beta \Pi_R^T = \sum_n \sum_m \text{sign}(\tilde{\lambda}_n) |\lambda_m| ([\Pi_L]_{nm})^2 \hat{e}_n \hat{e}_n^T$$

If $[\Pi_L]_{nm} \neq 0$ then $[O_L]_{nm} \neq 0$ and so $|\lambda_m| = |\tilde{\lambda}_n|$. Therefore

$$[\Pi_L] \beta \Pi_R^T = \sum_n \sum_m \text{sign}(\tilde{\lambda}_n) |\tilde{\lambda}_n| ([\Pi_L]_{nm})^2 \hat{e}_n \hat{e}_n^T$$

$$\Pi_L \beta \Pi_R^T = \sum_n \tilde{\lambda}_n \sum_m ([\Pi_L]_{nm})^2 \hat{e}_n \hat{e}_n^T$$

Since every column of Π_L has only one element which is proportional to 1, $\sum_m ([\Pi_L]_{nm})^2 = 1$ for all n . Therefore

$$\Pi_L \beta \Pi_R^T = \sum_n \tilde{\lambda}_n \hat{e}_n \hat{e}_n^T = O_L \beta O_R^T$$

\square

E Proof of Theorem 5

Clearly, the realness-under-Clifford rotations problem is in NP, since a prover can give the verifier the single-qubit Clifford rotations which make the terms in H real. To prove the problem to be NP-hard, we show that if one can efficiently solve the realness-under-Cliffords problem, then one can also efficiently solve an NP-complete subclass of the exact cover problem, called restricted exact cover by three-sets.

Definition 38. Restricted Exact Cover by Three-Sets (RXC3): *Given a finite set of $3N$ elements $E = \{e_1, e_2, \dots, e_{3N}\}$, and a collection $R = \{S_1, S_2, \dots\}$ of 3-element subsets of E , i.e. $|S_i| = 3$, with the restriction that every element in E appears in **exactly** three sets $S_i \in R$. Find a subcollection $R' \subset R$ such that every element in E occurs in exactly one member of R' .*

It has been proven in [28] that RXC3 is NP-complete.

Consider an instance of an RXC3 problem $A = (E, R)$. We can construct a corresponding Hamiltonian H_A as follows. Let each element in E correspond to a qubit. We can construct H_A iteratively, starting with $H_A = 0$. For every $e_i \in E$, define the pool of available Pauli operators $P_i = (X_i, Y_i, Z_i)$, where X_i is a Pauli- X on the i th qubit, and identity on all others. For every subset $(e_i, e_j, e_k) \in R$, take Paulis σ_i, σ_j , and σ_k from the respective pools P_i, P_j , and P_k , reducing the pool sizes by one. Then let $H_A \rightarrow H_A + \sigma_i \sigma_j + \sigma_j \sigma_k + \sigma_i \sigma_k$. Because of the restriction in RXC3 that every element $e_i \in E$ appears in exactly 3 subsets in R , we can be assured that our pools P_i will always be able to supply a Pauli operator throughout this procedure, and will be exhausted by the end. Unless a particular ordering is imposed on R and P_i , the definition of the above procedure is ambiguous. Any choice of ordering fixes the mapping. Let us suppose we have fixed a consistent ordering procedure for R and every P_i .

For the sake of clarity, we include an example of an RX3C problem, and its corresponding Hamiltonian, with the pools ordered as $\{X, Y, Z\}$.

RXC3 form:

$$\begin{aligned} E &= \{e_1, e_2, e_3, e_4, e_5, e_6\} \\ R &= \begin{aligned} &(e_1, e_2, e_3), \\ &(e_3, e_4, e_5), \\ &(e_2, e_3, e_5), \\ &(e_1, e_4, e_6), \\ &(e_2, e_5, e_6), \\ &(e_1, e_4, e_6) \end{aligned} \end{aligned} \quad (31)$$

Hamiltonian form:

$$\begin{aligned}
H_A = & X_1X_2 & +X_2X_3 & +X_1X_3 \\
& +Y_3X_4 & +X_4X_5 & +Y_3X_5 \\
& +Y_2Z_3 & +Z_3Y_5 & +Y_2Y_5 \\
& +Y_1Y_4 & +Y_4X_6 & +Y_1X_6 \\
& +Z_2Z_5 & +Z_5Y_6 & +Z_2Y_6 \\
& +Z_1Z_4 & +Z_4Z_6 & +Z_1Z_6
\end{aligned} \tag{32}$$

Given an RXC3 problem A , we can use the above procedure to produce a Hamiltonian H_A . We will now show that if we can efficiently determine whether or not H_A can be made real under Clifford operations, then we can efficiently determine the answer to A .

Suppose H_A can be made real under Clifford operations, then there exists some permutations on the qubits such that the only terms in the new Hamiltonian H'_A which contain Y operators are of the form Y_iY_j . Recall that H_A is composed exclusively of triplets of terms of the form $\sigma_i\sigma_j + \sigma_j\sigma_k + \sigma_i\sigma_k$, with σ_i, σ_j , and σ_k not appearing in any other terms. Therefore it must be the case that any of the terms in H'_A of the form Y_iY_j can be grouped into triplets of the form $T_{ijk} = Y_iY_j + Y_jY_k + Y_iY_k$, with Y_i, Y_j , and Y_k not appearing in any other terms. Every T_{ijk} corresponds exactly with a 3-element subset $(e_i, e_j, e_k) \in R$. Since Y_i, Y_j , and Y_k do not appear in any other terms, it follows that the set of 3-element subsets in R associated with all such triplets T_{ijk} form the exact covering R' . So if H_A can be made real under single qubit Clifford rotations, then A has an exact covering.

Suppose A has an exact covering R' . For every 3-element subset $(e_i, e_j, e_k) \in R'$ there exists a term $\sigma_i\sigma_j + \sigma_j\sigma_k + \sigma_i\sigma_k$ in H_A , with σ_i, σ_j , and σ_k not appearing in any other terms. Therefore we may choose to apply a sequence of permutations on the i, j and k qubits in H_A so that $\sigma_i \rightarrow Y_i, \sigma_j \rightarrow Y_j$ and $\sigma_k \rightarrow Y_k$ with the guarantee that Y_i, Y_j , and Y_k will then not appear in any other terms in our new Hamiltonian. We may continue in this fashion for every 3-element subset in R' without disturbing any of the previous qubits we have acted upon, and ultimately acting on every qubit. Therefore at the end of the procedure we will have produced a Hamiltonian H'_A for which the only terms containing Y operators are those of the form Y_iY_j , and therefore is real. So if A has an exact covering then H_A can be made real by single qubit Clifford rotations.