

Complexity of Fermionic 2-SAT

Maarten Stroeks^{1,2} and Barbara M. Terhal^{1,2}

¹Delft Institute of Applied Mathematics, Delft University of Technology, The Netherlands

²QuTech, Delft University of Technology, The Netherlands

We introduce the fermionic satisfiability problem, FERMIONIC k -SAT: this is the problem of deciding whether there is a fermionic state in the null-space of a collection of fermionic, parity-conserving, projectors on n fermionic modes, where each fermionic projector involves at most k fermionic modes. We prove that this problem can be solved efficiently classically for $k = 2$. In addition, we show that deciding whether there exists a satisfying assignment with a given fixed particle number parity can also be done efficiently classically for FERMIONIC 2-SAT: this problem is a quantum-fermionic extension of asking whether a classical 2-SAT problem has a solution with a given Hamming weight parity. We also prove that deciding whether there exists a satisfying assignment for particle-number-conserving FERMIONIC 2-SAT for some given particle number is NP-complete. Complementary to this, we show that FERMIONIC 9-SAT is QMA₁-hard.

Contents

1	Introduction	2
1.1	Main results	2
1.2	Overview of paper and main ideas	3
1.2.1	Differences with Quantum 2-SAT	5
2	Preliminaries	5
2.1	Characterization of the Fermionic 2-SAT projectors	6
2.1.1	Higher-rank projectors	7
2.2	Particle-hole transformation	8
2.3	Quantum clusters	8
3	Characterizing the satisfying assignments of Fermionic 2-SAT	11
3.1	Properties of satisfying assignments	11
3.2	Cluster-product form of satisfying assignments	13
3.3	Excluding certain particle numbers or parities on quantum clusters	15
3.3.1	Quantum clusters: lines and loops	17
3.3.2	non-hPNC lines and loops	19
3.3.3	Efficient verification of allowed hidden particle number or parity on a quantum cluster	21
4	An efficient classical algorithm for Fermionic 2-SAT (with and without fixed parity)	21

Maarten Stroeks: m.e.h.m.stroeks@tudelft.nl

4.1	Solving Fermionic 2-SAT: proof of Theorem 1.4	21
4.2	Solving Fermionic 2-SAT with fixed parity: proof of Theorem 1.5	22
5	PNC Fermionic 2-SAT with fixed particle number is NP-complete	25
6	Complexity of Fermionic k-SAT and related problems	25
7	Discussion	26
A	Mathematical micro-facts for Section 2	28
A.1	Invariance of Fermionic 2-SAT clauses that exclude all-empty or all-filled states	28
A.2	Action of the particle-hole transformation K_S on $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$ clauses . . .	29
B	Example of a unique 4-fermion non-Gaussian, non-product satisfying assignment	29
C	Classical 2-SAT with fixed parity	31
D	Proof of Lemma 6.1	34

1 Introduction

Classical k -satisfiability (k -SAT) is the problem of deciding whether there exists a configuration of n Boolean variables which satisfies a collection of k -variable Boolean clauses. While 2-SAT is solvable in linear time [1], $(k \geq 3)$ -SAT is NP-complete. Bravyi introduced the problem QUANTUM k -SAT [2]: this is the problem of deciding whether there exists an n -qubit state which is in the null-space of some collection of k -qubit projectors. Bravyi showed that QUANTUM 2-SAT can be decided classically efficiently, and later it was shown it can even be done in linear time [3, 4]. For $k \geq 3$, it has been shown that QUANTUM k -SAT is QMA₁-complete [2, 5].

Here, we introduce the *fermionic* satisfiability problem. An n -mode fermionic system corresponds to set of annihilation operators $\{a_j\}_{j=1}^n$ (with Hermitian conjugate creation operators $\{a_j^\dagger\}_{j=1}^n$) which satisfy $\{a_i^\dagger, a_j\} := a_i^\dagger a_j + a_j a_i^\dagger = \delta_{ij}I$, $\{a_i, a_j\} = 0$. In addition, there is a state $|\text{vac}\rangle$, —called the vacuum, or empty, zero-particle state—, for which $a_i |\text{vac}\rangle = 0$, $\forall i$.

FERMIONIC k -SAT is the problem of deciding whether there exists an n -mode fermionic state which is in the null-space of a collection of k -mode fermionic projectors. Each k -mode projector is a polynomial in a_j and $a_j^\dagger$ with $j \in S$, where $S \subseteq [n]$ and $|S| = k$. Furthermore, the k -mode projector commutes with the parity operator $\hat{P} := (-1)^{\sum_{i=1}^n a_i^\dagger a_i}$. As a special case, we also study PARTICLE-NUMBER-CONSERVING (PNC) FERMIONIC k -SAT, where each projector additionally commutes with the particle-number operator $\hat{N} := \sum_{i=1}^n a_i^\dagger a_i$. In addition to studying the satisfiability of these problems, we investigate the complexity of PNC FERMIONIC 2-SAT and FERMIONIC 2-SAT when, respectively, the particle number is fixed to some given value $\hat{N} = N \in \{0, 1, \dots, n\}$, or the parity $\hat{P} = P \in \{+1, -1\}$ is fixed.

1.1 Main results

Problem 1.1 (FERMIONIC k -SAT). *Given is a set of projectors $\{\Pi_i\}$ on n fermionic modes, where each Π_i is a polynomial in the operators $a_j^\dagger$ and a_j with index j contained in*

a subset of at most k fermionic modes. Furthermore, $[\Pi_i, \hat{P}] = 0$ for all projectors $\{\Pi_i\}$ with $\hat{P} = (-1)^{\sum_{i=1}^n a_i^\dagger a_i}$. Decide whether there exists a $|\psi\rangle$ s.t. $\forall i, \Pi_i |\psi\rangle = 0$ (YES), or for all $|\psi\rangle$, $\sum_i \langle \psi | \Pi_i | \psi \rangle > \frac{1}{\text{poly}(n)}$ (NO).

To be precise in our definition of FERMIONIC k -SAT one should, in principle, provide the accuracy with which the projectors Π_i are specified, i.e. with respect to what elementary gate set, see e.g. [3, 6] for how this is handled for QUANTUM k -SAT. Since this is separate from our results, we leave such precise definition to future work.

We also consider FERMIONIC k -SAT with global additional constraints, i.e. one asks for an assignment with given particle number or parity:

Problem 1.2 (FERMIONIC k -SAT with fixed parity P). *Given n fermionic modes, $P \in \{-1, +1\}$, and a set of projectors $\{\Pi_i\}$ where each projector Π_i is a polynomial in the operators $a_j^\dagger$ and a_j with index j contained in a subset of at most k fermionic modes. Furthermore, $\forall i, [\Pi_i, \hat{P}] = 0$. Decide whether there exists $|\psi\rangle$ s.t. $\forall i, \Pi_i |\psi\rangle = 0$ and $\hat{P} |\psi\rangle = P |\psi\rangle$ (YES), or $\forall |\psi\rangle$ s.t. $\hat{P} |\psi\rangle = P |\psi\rangle$, $\sum_i \langle \psi | \Pi_i | \psi \rangle > \frac{1}{\text{poly}(n)}$ (NO).*

A special case of Problem 1.1 is the case when each projector commutes with the particle number, i.e. $\forall i, [\Pi_i, \hat{N}] = 0$, which we will refer to as PARTICLE-NUMBER-CONSERVING (PNC) FERMIONIC k -SAT. We also consider the following problem:

Problem 1.3 (PNC FERMIONIC k -SAT with fixed particle number N). *Given n fermionic modes, an integer $N \in \{0, 1, \dots, n\}$ and a set of projectors $\{\Pi_i\}$ where each projector Π_i is a polynomial in the operators $a_j^\dagger$ and a_j with index j contained in a subset of at most k fermionic modes. Furthermore, $\forall i, [\Pi_i, \hat{N}] = 0$. Decide whether there exists a state $|\psi\rangle$ s.t. $\forall i, \Pi_i |\psi\rangle = 0$ and $\hat{N} |\psi\rangle = N |\psi\rangle$ (YES), or $\forall |\psi\rangle$ s.t. $\hat{N} |\psi\rangle = N |\psi\rangle$, $\sum_i \langle \psi | \Pi_i | \psi \rangle > \frac{1}{\text{poly}(n)}$ (NO).*

Our main results are:

Theorem 1.4. FERMIONIC 2-SAT $\in \text{P}$, and can be decided in time $O(n + m)$, where $m = |\{\Pi_i\}|$ denotes the number of projectors.

Theorem 1.5. FERMIONIC 2-SAT with fixed parity $\in \text{P}$, and can be decided in time $O(nm)$, where $m = |\{\Pi_i\}|$ denotes the number of projectors.

Theorem 1.6. PNC FERMIONIC 2-SAT with fixed particle number is NP-complete.

As part of the proof of Theorem 1.5, we show that classical 2-SAT with fixed (Hamming weight) parity can be solved in $O(nm)$ time, see Theorem C.1 in Appendix C for details. To the best of our knowledge, this has not been proved before.

To provide further context for our results, we also provide a straightforward proof that FERMIONIC k -SAT is QMA₁-hard for $k = 9$ in Section 6.

1.2 Overview of paper and main ideas

FERMIONIC 2-SAT instances are defined on a graph G whose vertices are fermionic modes and whose edges correspond to projectors involving pairs of modes. These projectors are of rank at most three (if there is to be a satisfying assignment) in the two fermionic mode subspace, consisting of at most three rank-1 projectors. These rank-1 projectors — we will refer to them as *clauses* — can be either genuinely quantum clauses, or classical clauses which exclude a particular occupation of the two modes. When *restricted to a two-qubit*

space, a genuinely quantum clause can exclude support on a state such as $\alpha|00\rangle + \beta|11\rangle$ (with $|\alpha|, |\beta| \neq 0$) which we call a $\Pi_e^{02,q}$ clause, or exclude support on a state such as $\alpha|01\rangle + \beta|10\rangle$ (with $|\alpha|, |\beta| \neq 0$) which we call a $\Pi_e^{1,q}$ clause. Since projectors in FERMIONIC 2-SAT are parity conserving by definition, these are the only allowed quantum clauses besides the classical clauses. Note that we mention the caveat *when restricted to a two-qubit space* since fermionic operators when viewed as acting on an n -qubit space are generally non-local due to Jordan-Wigner strings.

Let us discuss the ideas behind the proofs of our main results of Section 1.1. Given the graph G , we first remove all classical clauses, leaving us possibly with a set of disconnected graphs which we call quantum clusters, see e.g. Figure 1 in Section 2.3. We will refer to modes that are not involved in any quantum clauses as classical modes. It turns out to be useful to distinguish at least two types of quantum clusters, namely ones that are so-called hidden particle-number-conserving (hPNC) and ones which are not (non-hPNC). The hPNC clusters can be brought to a form in which they only contain $\Pi_e^{1,q}$ quantum clauses by a particle-hole transformation defined in Section 2.2: after this transformation, every clause commutes with the cluster particle number operator, hence the terminology. In Section 2 we introduce and discuss these various concepts.

In Section 3 we start by examining the effect of ‘constraint propagation’ by $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$ quantum clauses, see Lemma 3.1 and its Corollary 3.2. This leads up to a characterization of the satisfying assignments: wlog, they are of — what we call — *cluster-product form*, see Proposition 3.3. Compared to QUANTUM 2-SAT which, except for on some rank-3 projectors, always has a product satisfying assignment (if any) [3], this form is more general, but also restricted in an interesting way. For assignments with or without overall fixed parity, we will show that one can limit oneself to cluster-product states which are products of operators (acting on the vacuum) which create (1) classical occupations on so-called classical modes, (2) Gaussian states on hPNC clusters, (3) Gaussian states on non-hPNC clusters which are lines or loops, and (4) non-Gaussian states on certain 4-fermion non-hPNC clusters. To obtain this result, we proceed in Section 3.3 to exclude certain quantum numbers, —these are hidden cluster particle number or cluster parity which uniquely fix a cluster satisfying assignment—, on quantum clusters. Crucial in this is Lemma 3.4 which provides a uniquely fermionic simplification, which leaves us to analyze quantum clusters which are lines and loops in Section 3.3.1. Fermions on a line do map directly onto qubits, hence we can use some proof techniques borrowed from QUANTUM 2-SAT investigations. Importantly, these results allow one to efficiently (with effort linear in the number of modes and clauses) verify which quantum numbers, i.e. hidden cluster particle number and cluster parity, are allowed on quantum clusters in the cluster-product form. This is expressed in Lemma 3.10.

To prove Theorems 1.4 and 1.5 in Section 3, we bring back all classical clauses in the graph G and reduce the problem of deciding FERMIONIC 2-SAT after several preprocessing steps to deciding a classical 2-SAT instance (thereby proving Theorem 1.4). For Theorem 1.5, there is the additional constraint of asking for a satisfying assignment of fixed parity. In Appendix C, we provide an efficient algorithm for deciding classical 2-SAT where one asks for a solution of fixed Hamming weight parity (Theorem C.1), which is used as a subroutine in the algorithm for the fermionic problem (Theorem 1.5). This result on parity-constrained classical 2-SAT may not be surprising, but we are unaware of any previous algorithm or proof. Crucial in the reductions in the proofs of Theorems 1.4 and 1.5 is the fact that non-classical assignments on quantum clusters heavily constrain how classical clauses which straddle the quantum cluster and a classical mode can be satisfied, namely the clause has to be true by only the choice of the classical mode, since the non-classical assignment

requires that each mode inside the cluster should be allowed to be both empty (0) *and* filled (1). The nature of the hPNC and non-hPNC clusters asks for a different treatment in the proofs of these theorems.

In Section 5, we examine the complexity of particle-number-conserving FERMIONIC 2-SAT where we ask for an assignment with given fixed particle number, and prove Theorem 1.6. Given the efficiently checkable cluster-product form, it may not be surprising that this problem is contained in NP. It can be proved to be NP-complete since the class of problems contains the weighted 2-SAT problem which asks for a 2-SAT satisfying assignment with fixed Hamming weight, as a special case. In Section 6, we prove a quantum hardness result for FERMIONIC k -SAT, still leaving a considerable gap between our efficient classical algorithms for $k = 2$ and this hardness result for $k = 9$. We discuss open problems in the Discussion Section 7.

1.2.1 Differences with Quantum 2-SAT

We take a paragraph to reflect on the difference between our FERMIONIC 2-SAT findings and QUANTUM 2-SAT.

For QUANTUM k -SAT and FERMIONIC k -SAT, the projectors act nontrivially on k qubits resp. involve k fermionic modes. Fermion-to-qubit mappings allow for the projectors of a given instance of FERMIONIC k -SAT to be mapped to qubit projectors. However, each fermionic projector involving k modes gets mapped—in a general setting—onto a q -qubit projector, where $q \geq k$ [7]. Therefore, one cannot deduce the complexity of FERMIONIC k -SAT directly from the complexity of QUANTUM k -SAT.

Central to the solution of QUANTUM 2-SAT is the fact that a satisfying assignment, if it exists, can be taken to be a product state (except for two-qubit projectors whose null-space is a single entangled state), see e.g. Theorem 4 in [3]. We do not have such a statement for FERMIONIC 2-SAT. In fact, we find that satisfying assignments on non-hPNC clusters can include small 4-fermion non-Gaussian states. Furthermore, for FERMIONIC 2-SAT, quantum clusters with vertices with degree at least 3 obey specific constraints which relate to uniquely fermionic signs, see Lemma 3.4 in Section 3.3, excluding many particle numbers (but leaving only Gaussian states for hPNC clusters). In contrast, for quantum 2-SAT, product state assignments are generally built from superpositions of all particle numbers. Indirectly, Lemma 3.4 also limits the use of any possible non-Gaussian assignment: non-Gaussian states are only *sometimes* needed on non-hPNC clusters of at most 4 fermionic modes (and any state of three or fewer fermionic modes is always Gaussian [8]). In Appendix B we work through this 4-fermion non-Gaussian example explicitly, showing how the departure from product state assignments for QUANTUM 2-SAT comes about.

2 Preliminaries

In this section we review fermionic language and concepts, and introduce several tools for the treatment of assignments on quantum clusters.

For a subset $S \subseteq [n]$ and $|S| = k$, we define the operator $a_S^\dagger := a_{i_1}^\dagger a_{i_2}^\dagger \dots a_{i_k}^\dagger$, where $i_1 < i_2 < \dots < i_k$. This ordering of the fermionic modes will later be conveniently chosen depending on the FERMIONIC 2-SAT instance in Section 2.3. We can write any n -mode fermionic state as

$$|\psi\rangle = \sum_{S \subseteq [n]} \alpha_S |S\rangle, \quad |S\rangle := a_S^\dagger |\text{vac}\rangle, \quad (1)$$

where $\alpha_S \in \mathbb{C}$ is the amplitude for state $|S\rangle$. Each state $|S\rangle$ corresponds to a configuration of mode *occupations* $\{x_i\}_{i=1}^n$, i.e. the bit $x_i = 1$ (resp. $x_i = 0$) signifies that mode i is occupied, so $i \in S$, (resp. unoccupied, so $i \notin S$). When we speak of a *classical assignment* in the remainder of this work, we mean a state of type $|S\rangle = a_S^\dagger |\text{vac}\rangle$.

One can map an (ordered) set of fermionic creation and annihilation operators via the Jordan-Wigner (JW) transformation onto a set of qubit operators:

$$a_i^\dagger \stackrel{\text{JW}}{=} Z_1 \dots Z_{i-1} \sigma_i^+, \quad (2)$$

and the fermionic vacuum state $|\text{vac}\rangle \rightarrow |00\dots 0\rangle$, using the definition $\sigma_i^+ = |1\rangle\langle 0|_i$ (and similarly $\sigma_i^- = |0\rangle\langle 1|_i$).

Gaussian fermionic states are thermal states and eigenstates of *quadratic* fermionic Hamiltonians — corresponding to so-called non-interacting fermions — and these states can be fully characterized by a $2n \times 2n$ correlation matrix

$$M_{ij} = \frac{i}{2} \text{Tr}([c_i, c_j] \rho), \quad c_{2i-1} = a_i + a_i^\dagger, \quad c_{2i} = i(a_i - a_i^\dagger), \quad \{c_i, c_j\} = 2\delta_{ij}I. \quad (3)$$

A state is a pure Gaussian state iff its correlation matrix obeys $M^T M = I$, see e.g. [8, 9]. The n -mode states with particle number $\hat{N} = 0$ (vacuum), $\hat{N} = 1$ (single-particle "Slater determinant" state), $\hat{N} = n - 1$ (single-hole state) and $\hat{N} = n$ (all-filled) are Gaussian states. The classical states $|S\rangle$ are also examples of Gaussian states. See more information on Gaussian states in [10].

2.1 Characterization of the Fermionic 2-SAT projectors

Let $G = (V, E)$ be a graph such that the vertices $v \in V$ label the modes ($|V| = n$) and each edge is associated with a projector $\{\Pi_e\}_{e \in E}$. Since the projectors in FERMIONIC 2-SAT are parity conserving, any projector $\Pi_{e=(j,k)}$ is a polynomial in $a_j, a_j^\dagger, a_k$ and $a_k^\dagger$ of even degree. Obviously, since each projector Π_e on modes j and k commutes with the overall parity operator $\hat{P}$, it also commutes with the edge parity $P_{jk} = (-1)^{a_j^\dagger a_j + a_k^\dagger a_k}$. Similarly, if Π_e commutes with $\hat{N}$, it commutes with the edge particle number $\hat{N}_{jk} = \hat{N}_j + \hat{N}_k$. Hence, if we consider a two-mode projector in isolation, it can be represented as a QUANTUM 2-SAT projector on two qubits which is parity-conserving (or, in special cases, also particle number-conserving). This means that if the projector, say of rank-1 in the 2-qubit space, projects onto a product state, the product state can only be a computational basis state. And if it projects onto an entangled state, then in the number-conserving case, it can only project onto states such like $\alpha |01\rangle + \beta |10\rangle$. In the merely parity-conserving case, it can also project onto a state $\alpha |00\rangle + \beta |11\rangle$ for some α, β . This picture is entirely preserved when we consider the action of the fermionic projectors onto n modes. In principle, one can apply a Jordan-Wigner transformation, see Eq. (2), and represent these fermionic projectors as 2-qubit projectors to which strings of Pauli Z 's are appended, so that the qubit projectors are not 2-local, but this does not affect the number- or parity-conserving properties. In some sense, we thus consider a simplified form of QUANTUM 2-SAT due to number- and parity-conserving constraints, but the fermionic nature of the problem, i.e. the effect of the Pauli Z -strings through the Jordan-Wigner transformation, makes the problem genuinely different from QUANTUM 2-SAT.

Let's now discuss the nature of rank-1 projectors¹ in fermionic language. With the

¹They have rank-1 only in the 2-mode subspace, not in the n -mode fermionic space, but for convenience we use this term throughout this paper. Same for statements on higher rank projectors which have a rank $r \leq 4$ in the 2-mode subspace.

two-qubit picture in mind, we have two different types of rank-1 projectors for an edge $e = (j, k)$ (with convention $j < k$). The only possible rank-1 projector onto an odd-parity state is

$$\begin{aligned}\Pi_e^1 &= |\beta_e|^2 a_j^\dagger a_j a_k a_k^\dagger + \beta_e \gamma_e^* a_j^\dagger a_k - \gamma_e \beta_e^* a_j a_k^\dagger + |\gamma_e|^2 a_j a_j^\dagger a_k^\dagger a_k \\ &\rightarrow \text{"project onto/exclude } \beta_e |10\rangle + \gamma_e |01\rangle",\end{aligned}\quad (4)$$

with $|\beta_e|^2 + |\gamma_e|^2 = 1$. When Π_e^1 does not project onto just "01" or "10", we call this a genuinely quantum clause and write $\Pi_e^{1,q}$, otherwise we refer to them as classical clauses $\Pi_e^{1,c}$. The only type of rank-1 projector onto an even-parity state is Π_e^{02} :

$$\begin{aligned}\Pi_e^{02} &= |\alpha_e|^2 a_j a_j^\dagger a_k a_k^\dagger - \alpha_e \delta_e^* a_j a_k + \delta_e \alpha_e^* a_j^\dagger a_k^\dagger + |\delta_e|^2 a_j^\dagger a_j a_k^\dagger a_k \\ &\rightarrow \text{"project onto/exclude } \alpha_e |00\rangle + \delta_e |11\rangle",\end{aligned}\quad (5)$$

with $|\alpha_e|^2 + |\delta_e|^2 = 1$. Again, when Π_e^{02} does not project onto just "00" or "11", it is a genuinely quantum clause and we write $\Pi_e^{02,q}$, otherwise it is a classical projector $\Pi_e^{02,c}$.

For convenience, we introduce the following shorthand for a given quantum clause on an edge e :

$$\begin{aligned}u_e &:= \gamma_e^* / \beta_e^* \in \mathbb{C}, \text{ for } \Pi_e^{1,q}, \\ v_e &:= \alpha_e^* / \delta_e^* \in \mathbb{C}, \text{ for } \Pi_e^{02,q},\end{aligned}\quad (6)$$

which we shall only use later on when $\beta_e, \delta_e \neq 0$, i.e. in Lemmas 3.1, 3.4, 3.7 and 3.8.

For PNC FERMIONIC 2-SAT, we only have projectors of type Π_e^1 in Eq. (4), and only classical projectors

$$\begin{aligned}\Pi_e^0 &= a_j a_j^\dagger a_k a_k^\dagger \rightarrow \text{"project onto/exclude } |00\rangle", \\ \Pi_e^2 &= a_j^\dagger a_j a_k^\dagger a_k \rightarrow \text{"project onto/exclude } |11\rangle".\end{aligned}\quad (7)$$

Proposition A.1 in Appendix A proves that these projectors are invariant under taking any unitary linear combination of the two modes j and k , but this understanding also follows from the two-qubit picture above.

2.1.1 Higher-rank projectors

The projectors appearing in FERMIONIC 2-SAT are of rank at most three and can be obtained by summing the unit-rank projectors above as long as these project onto orthogonal states. This is true since the 1-eigenvalue sub-space of a projector Π_e is spanned by eigenstates of the commuting edge parity operator $\hat{P}_{e=(j,k)}$, or, in the particle-number-conserving case, by eigenstates of the commuting edge number operator $\hat{N}_{e=(j,k)}$. By viewing the action of the rank-1 projectors in the two-qubit space, it is clear that there exist two orthogonal projectors in the $\{|01, 10\rangle\}$ space and two orthogonal projectors in the $\{|00\rangle, |11\rangle\}$ space.

Observe that a rank-2 projector which is the sum of two orthogonal $\Pi_e^{02,q}$ clauses is a rank-2 classical projector (a sum of two classical clauses, excluding both "00" and "11"), and similarly a rank-2 projector which is the sum of two orthogonal $\Pi_e^{1,q}$ clauses is classical (excluding both "01" and "10"). We will call an edge a *classical edge* ($\in E_c$) when its at-most-rank-3 projector is constructed from rank-1 classical clauses, otherwise we call it a quantum edge ($\in E_q$), thus the edgeset can be written as $E = E_c \cup E_q$. Hence, an edge in E_q consists of one — $\Pi_e^{02,q}$ or $\Pi_e^{1,q}$ — or two — $\Pi_e^{02,q}$ and $\Pi_e^{1,q}$ — quantum

clauses. Otherwise, the edge belongs to E_c . It is worthwhile to note that a rank-2 projector $\Pi_e^{1,q} + \Pi_e^{02,q}$ only consists of quadratic terms in creation and annihilation operators, i.e. the quartic contributions cancel as the quartic contribution $\propto a_j^\dagger a_j a_k^\dagger a_k$ of each quantum clause is independent of the choice of $\alpha_e, \beta_e, \gamma_e, \delta_e$ and opposite in sign for a $\Pi_e^{02,q}$ clause versus a $\Pi_e^{1,q}$ clause as can be seen from Eqs. 4 and 5. This fact will be used in Lemma 3.9.

2.2 Particle-hole transformation

We define the unitary particle-hole transformation K_S on a subset $S \subseteq [n]$ of the n modes. The transformation interchanges the creation of a particle in S with the annihilation of a particle, i.e. the creation of a hole or absence of a particle, in S :

$$\begin{aligned} \forall j \in S \quad &\rightarrow \quad K_S a_j K_S^{-1} = \left[\prod_{\substack{k < j \\ \text{s.t. } k \in S}} (-1) \right] a_j^\dagger, \quad K_S a_j^\dagger K_S^{-1} = \left[\prod_{\substack{k < j \\ \text{s.t. } k \in S}} (-1) \right] a_j, \\ \forall j \notin S \quad &\rightarrow \quad K_S a_j K_S^{-1} = \left[\prod_{\substack{k < j \\ \text{s.t. } k \in S}} (-1) \right] a_j, \quad K_S a_j^\dagger K_S^{-1} = \left[\prod_{\substack{k < j \\ \text{s.t. } k \in S}} (-1) \right] a_j^\dagger. \end{aligned} \quad (8)$$

K_S also transforms the vacuum state $|\text{vac}\rangle$, i.e. $|\text{vac}\rangle \rightarrow K_S |\text{vac}\rangle$, where the new vacuum state corresponds to the all-filled state for the modes in the set S , and the original vacuum for the modes which are not in the set S . The $\pm$ signs in the transformation K_S in Eq. (8) can be understood by performing a Jordan-Wigner transformation which leads to the manifestly-unitary transformation $K_S = \bigotimes_{k \in S} X_k$, with X_k denoting Pauli- X on qubit k . When $S = \emptyset$, $K_{S=\emptyset} = I$.

When we apply the particle-hole transformation K_S to the genuinely quantum projectors in Eqs. (4) and (5), it can interchange $\Pi_e^{1,q}$ clauses and $\Pi_e^{02,q}$ clauses, depending on the edge e and the subset S . In Appendix A we give the effect of the transformation in mathematical detail (which is not extremely insightful).

An additional property of the particle-hole transformation K_S is that it preserves Gaussianity of a state ρ . We can express the unitary transformation K_S on the Majorana operators $c_{2i-1} \stackrel{J.W.}{=} Z_1 \dots Z_{i-1} X_i, c_{2i} \stackrel{J.W.}{=} Z_1 \dots Z_{i-1} Y_i$ (consistent with Eqs. 2 and 3). Then K_S is some product of Majorana operators $\Pi_{i \in S_{\text{Maj}}} c_i$, using $X_k = [\prod_{1 \leq j < k} (-i c_{2j-1} c_{2j})] c_{2k-1}$. We have $K_S c_k K_S^{-1} = \pm c_k$, with the sign depending on whether $k \in S_{\text{Maj}}$ and $|S_{\text{Maj}}|$, inducing a (diagonal) orthogonal transformation in the $\{c_i\}_{i=1}^{2n}$ basis. Any orthogonal transformation $R \in O(2n)$ preserves the Gaussianity of a pure fermionic state as it preserves the property that the correlation matrix M still obeys $M^T M = I$ [8], hence K_S preserves Gaussianity. Note that $K_S \hat{P} K_S^{-1} = \pm \hat{P}$, switching the overall parity P when $|S_{\text{Maj}}|$ is odd.

2.3 Quantum clusters

The graph $G = (V, E)$ describing the fermionic 2-SAT problem (where an edge represents a rank-1, 2 or 3 projector) can contain zero, one or more *quantum clusters*, see Fig. 1, defined as follows.

Definition 2.1 (Quantum clusters). *Take the subgraph G_{sub} of G obtained by removing all classical clauses in E_c , and removing all vertices that are not touching any $e \in E_q$. We define the quantum clusters $G_q^i = (V_q^i, E_q^i)$ (for $i = 1, \dots, Q$, and with $n_q^i := |V_q^i|$) to be the connected components of G_{sub} . The cluster particle number equals $\hat{N}_q^i = \sum_{j \in V_q^i} a_j^\dagger a_j$ and the cluster parity is $\hat{P}_q^i = (-1)^{\sum_{j \in V_q^i} a_j^\dagger a_j}$.*




Figure 1: Several connected subgraphs G_q^i of G with genuinely quantum clauses (orange), the other clauses (black) are classical clauses. The graph G_{sub} is obtained by removing all classical edges and vertices not touching quantum edges (referred to as classical modes), leaving the disconnected quantum clusters G_q^i .

If the quantum cluster contains only $\Pi_e^{1,q}$ clauses, we can call the cluster particle-number-conserving, and states with a fixed cluster particle number span the set of possible satisfying assignments in the cluster. Some clusters are not directly particle-number conserving, but can be brought to this form by a unitary particle-hole transformation defined in Section 2.2. We call these clusters *hidden particle number conserving* (hPNC), see also Fig. 2.

Definition 2.2 (Hidden Particle Number Conserving Cluster (hPNC)). *A quantum cluster G_q on the set of vertices V_q is hidden particle number conserving (hPNC) if we can divide the set of vertices into two sets, $V_q = A \cup B$, and all quantum clauses connecting vertices from A with vertices from B are $\Pi_e^{02,q}$ clauses, while all quantum clauses connecting vertices in either A or B internally, are $\Pi_e^{1,q}$ clauses. A special case is when there are no $\Pi_e^{02,q}$ clauses whatsoever in the quantum cluster, $-A$ is the set of all vertices, and $B = \emptyset$, and G_q is particle number conserving (PNC).*

Observe that in a hPNC cluster G_q , one can assume wlog that all edges are single clauses (rank-1): if they are of rank-2 or higher, than either, they are classical edges (hence they had been removed from the graph G , hence cluster G_q), or they are a sum of a $\Pi_e^{02,q}$ clause and a $\Pi_e^{1,q}$ clause, and hence the cluster cannot be hPNC. Note that if we have the cluster parity operator $\hat{P}_q$ for n_q modes, then the particle-hole transformation in Eq. (8) for a subset $S \subseteq [n_q]$, transforms $K_S \hat{P}_q K_S^{-1} = \pm \hat{P}_q$: hence an eigenstate of cluster parity is still an eigenstate of cluster parity after a particle-hole transformation, but possibly with a flipped eigenvalue.

In order to efficiently determine whether G_q is hPNC and examine some other properties later, it will be useful to define a “maximal spanning” subgraph T_q which is hPNC as follows. The idea is that this (non-unique) subgraph T_q can be obtained via a greedy approach and can be used to efficiently decide whether a quantum cluster G_q is hPNC.

Definition 2.3 (Maximal spanning hPNC subgraph of a quantum cluster). *Given a quantum cluster G_q (with n_q vertices and $|E_q|$ quantum clauses/edges), the graph T_q is a graph with the same vertices as G_q but a subset of quantum clauses of G_q such that (1) T_q is hPNC, (2) The quantum clauses in T_q generate a spanning tree for the graph G_q (i.e. all vertices in G_q are connected by some quantum clauses in T_q) and (3) T_q is maximal, i.e. it*




Figure 2: Example of a hidden particle number conserving quantum cluster G_q as given in Definition 2.2 where the green quantum edges are $\Pi_e^{02,q}$ clauses and the pink quantum edges are $\Pi_e^{1,q}$ clauses. If one does a particle-hole transformation K_B on all the fermionic modes in B (or K_A on A), the green $\Pi_e^{02,q}$ clauses become $\tilde{\Pi}_e^{1,q}$ clauses, while the pink $\Pi_e^{1,q}$ clauses become $\tilde{\Pi}_e^{02,q}$ clauses, see Eq. (31), thus staying of particle-conserving type. Hence, after the particle-hole transformation, all clauses in G_q are particle-number-conserving, and the space of satisfying assignments is then spanned by eigenstates $|\psi\rangle$ of the so-called *hidden cluster particle number* $\hat{N}_q$, obeying $\hat{N}_q(K_B |\psi\rangle) = N_q(K_B |\psi\rangle)$. Note that for a hPNC cluster with $B \neq \emptyset$, a satisfying assignment $|\psi\rangle$ will generally not be an eigenstate of cluster particle number in the original basis, but an eigenstate of the cluster parity.

becomes non-hPNC when adding any more clauses from G_q . Such (generally non-unique) T_q can be constructed with effort $O(n_q + |E_q|)$ as follows.

→ Start with T_q containing all n_q vertices, but no clauses. Pick a vertex and assign it to A .

→ For each vertex adjacent to the current one, add it to A if the connecting clause is of type $\Pi_e^{1,q}$, and add it to B if it is of type $\Pi_e^{02,q}$ (if it is the sum of $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$, pick one of them at random). Add these connecting clauses to T_q and move to the neighbor vertex.

→ Repeat this procedure for all vertices adjacent to the previous vertices (if they have not yet been assigned to A or B) until all vertices are in the bipartition $A \cup B$.

→ Then, for all clauses in G_q which are not yet in T_q , do the following. Add the clause to T_q if it is a $\Pi_e^{02,q}$ clause and the vertices at its ends are in opposite sets (i.e., one in A and the other in B). Conversely, include the clause in T_q if it is a $\Pi_e^{1,q}$ clause and the vertices at its ends have equal assignments.

If after this efficient procedure, no clauses are left in G_q , then obviously $T_q = G_q$ and G_q is hPNC.

For later reference, we state a simple fact about a choice for picking the starting vertex in T_q :

Fact 2.4. Suppose a quantum cluster G_q contains a vertex with degree at least 3, then a maximal spanning hPNC subgraph T_q can be efficiently constructed such that it also contains a vertex of degree at least 3, simply by choosing the initial vertex in the greedy approach in Definition 2.3 to be the vertex in G_q of degree at least 3.

3 Characterizing the satisfying assignments of Fermionic 2-SAT

3.1 Properties of satisfying assignments

The following ‘Partner’ Lemma captures the idea of *propagation of a constraint* by genuinely quantum clauses: recall that in the qubit case, a rank-1 *entangled* projector is always ‘propagating’, see Definition 1 and Lemma 5 in [3]: the assignment of a product state for one qubit uniquely fixes the assignment of a product state for the other qubit, given that the state has to be orthogonal to the entangled projector. Here propagation is a more general concept as we can’t assume a product state assignment. The Lemma roughly means that if a satisfying assignment for the whole graph G has support on a certain n -mode state, it also needs to have support on a ‘partnered’ n -mode state.

Lemma 3.1 (Partner Lemma). *Given an edge $e = (j, k)$ (with $j < k$), with (1) clause $\Pi_e^{1,q}$ or (2) $\Pi_e^{02,q}$ (or both) in an instance of FERMIONIC 2-SAT. Suppose that $|\psi\rangle$ is an n -mode satisfying assignment.*

1. $\Pi_e^{1,q}$ clause $\Rightarrow$ If $|\psi\rangle$ has support on a $|S\rangle = a_S^\dagger |\text{vac}\rangle$ which is given by occupations $\{x_i\}_{i=1}^n$ such that $x_j \neq x_k$, then $|\psi\rangle$ must also have support on a ‘partnered’ state $|S'\rangle$ with occupations $\{x'_i\}_{i=1}^n$ with $x'_j = \overline{x_j}$, $x'_k = \overline{x_k}$, $x'_{i \neq j,k} = x_i$. In particular, we must have for the amplitudes α_S in the satisfying assignment:

$$\begin{aligned} \alpha_S &= -\text{sign}(k, S')/\text{sign}(j, S) u_e \alpha_{S'}, \text{ if } x_j = 1 \text{ and } x_k = 0, \\ \alpha_{S'} &= -\text{sign}(k, S)/\text{sign}(j, S') u_e \alpha_S, \text{ if } x_j = 0 \text{ and } x_k = 1, \end{aligned} \quad (9)$$

where $\text{sign}(l, T)$ for $l \in T$ denotes the sign arising from the reordering of $a_T^\dagger$ to $a_l^\dagger a_{T \setminus l}^\dagger$ and u_e as in Eq. (6) for $\Pi_e^{1,q}$.

2. $\Pi_e^{02,q}$ clause $\Rightarrow$ If $|\psi\rangle$ has support $|S\rangle$ with occupations $\{x_i\}_{i=1}^n$ such that $x_j = x_k$, then $|\psi\rangle$ must also have support on $|S'\rangle$ with $\{x'_i\}_{i=1}^n$ with $x'_j = \overline{x_j}$, $x'_k = \overline{x_k}$, $x'_{i \neq j,k} = x_i$. In particular, we must have for the amplitudes α_S in the satisfying assignment:

$$\begin{aligned} \alpha_S &= -\text{sign}(jk, S) v_e \alpha_{S'}, \text{ if } x_j = x_k = 1, \\ \alpha_{S'} &= -\text{sign}(jk, S') v_e \alpha_S, \text{ if } x_j = x_k = 0, \end{aligned} \quad (10)$$

where $\text{sign}(lm, T)$ for $l, m \in T$ denotes the sign arising from reordering of $a_T^\dagger$ to $a_l^\dagger a_m^\dagger a_{T \setminus \{l,m\}}^\dagger$ and v_e as in Eq. (6) for $\Pi_e^{02,q}$.

Proof. Consider case 1. Since $|\psi\rangle = \sum_{T \subseteq [n]} \alpha_T a_T^\dagger |\text{vac}\rangle$ is a satisfying assignment, we must have $\langle \psi | \Pi_e^{1,q} | \psi \rangle = 0$. Using Eq. (4) for $\Pi_e^{1,q}$ ($j < k$), this implies

$$\langle \psi | \Pi_e^{1,q} | \psi \rangle = \sum_{T \subseteq [n] \text{ s.t. } j \in T, k \notin T} |\text{sign}(j, T) \beta_e^* \alpha_T + \text{sign}(k, T') \gamma_e^* \alpha_{T'}|^2 = 0, \quad (11)$$

where for each subset $T = (x_1, \dots, x_j = 1, \dots, x_k = 0, \dots, x_n)$, the partnered subset T' is defined as $(x_1, \dots, x_j = 0, \dots, x_k = 1, \dots, x_n)$. The relation between α_S and $\alpha_{S'}$ follows from Eq. (11). Consider case 2. To ensure $\langle \psi | \Pi_e^{02,q} | \psi \rangle = 0$ with $\Pi_e^{02,q}$ in Eq. (5), we must have that

$$\langle \psi | \Pi_e^{02,q} | \psi \rangle = \sum_{T \subseteq [n] \text{ s.t. } j, k \in T} |\alpha_e^* \alpha_{T'} + \text{sign}(jk, T) \delta_e^* \alpha_T|^2 = 0, \quad (12)$$

where for each subset $T = (x_1, \dots, x_j = 1, \dots, x_k = 1, \dots, x_n)$, the partnered subset T' is defined as $(x_1, \dots, x_j = 0, \dots, x_k = 0, \dots, x_n)$. The relation between α_S and $\alpha_{S'}$ follows from Eq. (12). $\square$

Lemma 3.1 strongly restricts the possible satisfying assignments on quantum clusters. The idea is that for a hPNC cluster G_q repeatedly applying the case 1 partner rule in Lemma 3.1 makes particles propagate on the connected quantum cluster graph G_q , so that all occupations for a fixed hidden cluster particle number are generated in the cluster. Then given a fixed hidden cluster particle number, there is at most 1 unique satisfying assignment.

If the cluster is non-hPNC, one can argue that there will be some additional $\Pi_e^{02,q}$ clauses, creating and annihilating pairs of particles via case 2, and in that case there is at most 1 unique satisfying assignment given fixed parity (and no satisfying assignments with fixed hidden cluster particle number).

In these arguments, care must be taken not to assume anything about the structure of global assignment, i.e. we derive the necessary structure of such global assignment based on how quantum clauses on each cluster can be satisfied. We capture these ideas in the following Corollary:

Corollary 3.2. *Suppose a quantum cluster G_q is hPNC with bipartition $A \cup B$, and $|\psi\rangle$ is a satisfying assignment for the entire graph G containing G_q . If $K_B |\psi\rangle$ (with $K_B = I$ in the strictly PNC case) has support on a n -mode state $|S\rangle$ with hidden cluster particle number N_q on G_q , i.e. $\hat{N}_q |S\rangle = N_q |S\rangle$, then it must have support on all $\binom{n_q}{N_q}$ n -mode states $|S'\rangle$ with particle number N_q on G_q where $|S'\rangle$ has equal occupation as $|S\rangle$ on all other modes in G . Moreover, there is at most one satisfying assignment per hidden cluster particle number $N_q \in \{0, 1, \dots, n_q\}$ on G_q alone. Thus for a hPNC cluster, N_q is the cluster particle number of the particle-hole transformed cluster, which we will refer to as the hidden cluster particle number.*

Suppose a quantum cluster G_q is non-hPNC and $|\psi\rangle$ is a satisfying assignment for G . If $|\psi\rangle$ has support on a n -mode state $|S\rangle$ with cluster parity P_q on G_q , i.e. $\hat{P}_q |S\rangle = P_q |S\rangle$, then it must have support on all $2^{n_q}/2$ n -mode states $|S'\rangle$ with parity P_q on G_q where $|S'\rangle$ has equal occupation as $|S\rangle$ on all other modes in G . There is at most one satisfying assignment per parity $P \in \{-1, +1\}$ on G_q alone.

Proof. The first part of the Corollary follows directly from repeatedly applying ‘the propagation rule’ of case 1 of Lemma 3.1 to an N_q -particle n -mode state $|S\rangle$ on which $K_B |\psi\rangle$ is supported, propagating the particles (only for $N_q = 0$ (all empty) and $N_q = n_q$ there is nothing to propagate and case 1 does not apply). The coefficient relation Eq. (9) implies that there is at most one satisfying assignment per hidden cluster particle number $N_q \in \{0, 1, \dots, n_q\}$ on the vertices of G_q alone. As for the second part of the corollary, consider a non-hPNC cluster and construct the maximal spanning hPNC subgraph T_q in Definition 2.3 with bipartition $A \cup B$. If we were to apply the particle-hole transformation K_B on T_q , any satisfying assignment $|\psi\rangle$ for the projectors in this subgraph T_q would be such that $K_B |\psi\rangle$ is a unique state for fixed cluster particle number as per the first part of this corollary. Since G_q is non-hPNC, using this bipartition of T_q , there will be (a) at least one internal $\Pi_e^{02,q}$ clause inside either A or B , or (b) one $\Pi_e^{1,q}$ clause connecting A and B . If we apply the particle-hole transformation K_B , then in case (a) such $\Pi_e^{02,q}$ clause stays a $\tilde{\Pi}_e^{02,q}$ “pair-creating” clause, implying that any satisfying assignment $|\psi\rangle$ for the quantum clauses in G_q has the property that $K_B |\psi\rangle$ is a unique state given some cluster

parity: this follows from case 2 in Lemma 3.1 and states of any particle number with that cluster parity must occur in the superposition (working in the particle-hole transformed basis). If $K_B |\psi\rangle$ is a state which involves all particle numbers with given parity, then so is the satisfying assignment $|\psi\rangle$, since K_B maps any state $|S\rangle$ onto some other state $|S'\rangle$. In case (b) upon the particle-hole transformation K_B , the connecting clause $\Pi_e^{1,q}$ transforms to “pair-creating” clause $\tilde{\Pi}_e^{02,q}$ and then the same arguments apply as in case (a). $\square$

3.2 Cluster-product form of satisfying assignments

Let’s examine the consequences of Corollary 3.2 for satisfying assignments for the entire graph G . Let **hPNC** be the set of hPNC quantum clusters, let **non – hPNC** be the set of non-hPNC quantum clusters and let **Class** be the collection of remaining modes not contained in any quantum cluster, to which we have referred as classical modes. Let’s us fix an ordering of the fermionic modes in the graph G using these sets **hPNC**, **non – hPNC** and **Class** and some chosen internal ordering of modes inside the sets. Say, we first assign labels $1, 2, \dots, n_q^1$ to the modes in the first hPNC quantum cluster G_q^1 , then labels $n_q^1 + 1, n_q^1 + 2, \dots, n_q^1 + n_q^2$ to the modes in the second hPNC quantum cluster G_q^2 , etc., then do the same for the non-hPNC quantum clusters, and then finally for all the classical modes. What is relevant is that in this ordering, we order the modes in each cluster immediately following each other, so the satisfying assignment of the cluster itself can be used in the assignment for the whole graph in so-called cluster-product form. Given this ordering, Corollary 3.2 implies that *any* satisfying assignment $|\psi\rangle$ must be a superposition of states of ordered cluster-product form, i.e.

$$|\psi\rangle = \sum_{\{N_q^i\}, \{P_q^j\}, \{x_k\}} \beta_{\{N_q^i\}, \{P_q^j\}, \{x_k\}} \left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} O_i(N_q^i) \right) |\text{vac}\rangle, \quad (13)$$

Here the operator $O_i(N_q^i)$ creates the unique state with hidden cluster particle number N_q^i using the modes of hPNC quantum cluster i , the operator $O_j(P_q^j)$ creates the unique state with cluster parity $^2 P_q^j = \pm 1$ using the modes in the non-hPNC cluster j , and the operator $O_k(x_k = 0) = I$, $O_k(x_k = 1) = a_k^\dagger$ for a classical mode $k \in \text{Class}$. Note that the operators $O_k(x_k)$, $O_j(P_q^j)$, $O_i(N_q^i)$ do not generally commute. We note that for a hPNC cluster i , non-hPNC cluster j and $k \in \text{Class}$, some values for respectively N_q^i , P_q^j and x_k in Eq. (13) may be excluded in order to be a satisfying assignment. For example, this can be due to further constraints on parity or particle number which we will examine in the next Section 3.3. We only consider the constraints due to classical clauses in the proofs of Theorems 1.4 and 1.5 in Section 4.

Let us now argue that taking a superposition of cluster-product states as in Eq. (13) is unnecessary, i.e. individual cluster-product states suffice. For this, we first examine some basic properties of the cluster-product form. For any $i_0 \in \text{hPNC}$, the cluster-product form is an eigenstate of the hidden cluster particle number $\hat{N}_q^{i_0}$ (where $K_B^{i_0}$ applies the

²Remember that cluster parity is defined in the original basis while *hidden* cluster particle number is defined in the particle-holed transformed basis in Definition 2.1 and Fig. 2.




Figure 3: Satisfying assignments of FERMIONIC 2-SAT can be taken to be of cluster-product form: ordered products of operators which create satisfying assignments on quantum clusters and classical assignments on classical modes starting from the vacuum state, as in Eq. (17). The possibly-created states on a hPNC cluster $G_q^i \equiv i \in \text{hPNC}$ (orange clusters) are completely specified by (hidden) cluster particle numbers $N_q^i = 0, 1, \dots, n_q^i$. The possibly-created states on a non-hPNC cluster $j \in \text{non-hPNC}$ (blue clusters) are completely specified by the cluster parity $P_q^j = \pm 1$. The possibly-created states for a classical mode $k \in \text{Class}$ are completely specified by the occupied/unoccupied label $x_k = 0, 1$. The black edges correspond to classical clauses. Note that these classical clauses can also be internal to a quantum cluster or straddle a quantum cluster and a classical mode.

particle-hole transformation for cluster i_0), i.e.

$$\begin{aligned} & \hat{N}_q^{i_0} \left(K_B^{i_0} \left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} O_i(N_q^i) \right) |\text{vac}\rangle \right) = \\ & \left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} [\hat{N}_q^i K_B^i] \delta_{i,i_0} O_i(N_q^i) \right) |\text{vac}\rangle = \\ & N_q^{i_0} \left(K_B^{i_0} \left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} O_i(N_q^i) \right) |\text{vac}\rangle \right), \end{aligned} \quad (14)$$

using Eq. (8), i.e. we can commute these operators through the ordered product. Similarly, for any $j_0 \in \text{non-hPNC}$

$$\begin{aligned} & \hat{P}_q^{j_0} \left(\left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} O_i(N_q^i) \right) |\text{vac}\rangle \right) = \\ & P_q^{j_0} \left(\left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} O_i(N_q^i) \right) |\text{vac}\rangle \right), \end{aligned} \quad (15)$$

and for any mode $k_0 \in \text{Class}$, we have

$$\begin{aligned} & \hat{N}_q^{k_0} \left(\left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} O_i(N_q^i) \right) |\text{vac}\rangle \right) = \\ & N_q^{k_0} \left(\left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} O_i(N_q^i) \right) |\text{vac}\rangle \right), \end{aligned} \quad (16)$$

We will use the following shorthand for the cluster-product form:

$$|\phi(\{N_q^i\}, \{P_q^j\}, \{x_k\})\rangle \equiv \left(\prod_{k \in \text{Class}} O_k(x_k) \right) \left(\prod_{j \in \text{non-hPNC}} O_j(P_q^j) \right) \left(\prod_{i \in \text{hPNC}} O_i(N_q^i) \right) |\text{vac}\rangle. \quad (17)$$

These facts give us a final characterization of the satisfying assignments, see Fig. 3:

Proposition 3.3 (Cluster-product form of satisfying assignments). *If there exists a satisfying assignment of a FERMIONIC 2-SAT problem, possibly with fixed particle number or fixed parity, then one can assume, with loss of generality, that it is of cluster-product form $|\phi(\{N_q^i\}, \{P_q^j\}, \{x_k\})\rangle$ in Eq. (17).*

Proof. Given Corollary 3.2 the most general solution is of the form in Eq. (13). However, if any superposition state with different cluster parities $\{P_q^j\}$ for non-hPNC clusters is a satisfying assignment, then so is each individual term in the superposition since cross terms with distinct cluster parities must be zero as $\sum_i \Pi_i$ commutes with any cluster parity. Similarly, if any superposition state with different hidden particle numbers $\{N_q^i\}$ for hPNC clusters is a satisfying assignment, then so is each individual term in the superposition since (using manipulations in the particle-hole transformed basis) cross-terms must be 0. Similarly, if any superposition state with different values $\{x_k\}$ is a satisfying assignment, then so is each individual term in the superposition since cross-terms must be 0 as $\sum_i \Pi_i$ commutes with $\hat{N}_k$. Thus wlog we can restrict ourselves to assignments of cluster-product form $|\phi(\{N_q^i\}, \{P_q^j\}, \{x_k\})\rangle$. $\square$

3.3 Excluding certain particle numbers or parities on quantum clusters

In this section we further examine the structure of satisfying assignments, namely the allowed values for (hidden) cluster particle numbers and parities in the cluster-product form in Eq. (17) which will restrict the search for an assignment. The following Lemma provides an important simplification for quantum clusters with vertices with degree at least 3.

Lemma 3.4 (Fermionic Degree ≥ 3 Simplification). *Let $G_q = (V_q, E_q)$ be a quantum cluster on n_q modes (=vertices) with at least one vertex of degree ≥ 3 and let T_q be its maximal spanning hPNC subgraph with partition $A \cup B$ in Definition 2.3. Then, for any satisfying n_q -mode assignment $|\psi\rangle$ for G_q , we must have $\langle S | K_B | \psi \rangle = 0$, for any $|S\rangle = a_S^\dagger |\text{vac}\rangle$ with hidden cluster particle number $\hat{N}_q = 2, 3, \dots, n_q - 2$.*

Proof. Using Fact 2.4, the graph T_q contains a vertex j of degree ≥ 3 and let three of its neighbor vertices be labeled k, l, m . We may assume wlog that $j < k < l < m < i_1 < i_2 < \dots$, with $i_1, i_2, \dots$ labeling the other modes in the quantum cluster G_q . We write $e = (j, k)$, $e' = (j, l)$ and $e'' = (j, m)$. In the particle-hole transformed basis, all clauses adjacent to j are of type $\Pi_e^{1,q}$ by construction via Definition 2.3 and Fact 2.4. After the particle-hole transformation we write $\Pi_{e=(j,k)}^{1,q}$, $\Pi_{e'=(j,l)}^{1,q}$ and $\Pi_{e''=(j,m)}^{1,q}$ with non-zero coefficients u_e , $u_{e'}$ and $u_{e''}$ introduced in Eq. (6), uniquely characterizing the clauses.

Suppose that $K_B |\psi\rangle$ has support on a state $|S\rangle = a_S^\dagger |\text{vac}\rangle$ with particle number $N_q = 2, 3, \dots, n_q - 2$. Lemma 3.1 implies that $K_B |\psi\rangle$ must have support on *all* $\binom{4}{2}$ states $|S'\rangle = a_{S'}^\dagger |\text{vac}\rangle$ s.t. in $|S'\rangle$ two modes from $\{j, k, l, m\}$ are occupied, and $|S'\rangle$ has occupation equal to $|S\rangle$ on all modes $V_q \setminus \{j, k, l, m\}$. Let the occupied modes in $V_q \setminus \{j, k, l, m\}$ be denoted by subset T (which depends on S). Let $K_B |\psi\rangle = \sum_S \alpha_S |S\rangle$, using, say, the fermionic mode ordering established at the start of Section 3.2, and assume wlog that in this ordering $T = T_{\text{before}} \cup T_{\text{after}}$ where T_{before} is a set of modes which come before j, k, l, m in the chosen fermionic order and T_{after} is a set of modes which come after and let $a_T^\dagger = a_{T_{\text{before}}}^\dagger a_{T_{\text{after}}}^\dagger$. Then $K_B |\psi\rangle$ must have support on each of the six states

$$a_j^\dagger a_k^\dagger a_T^\dagger |\text{vac}\rangle, a_j^\dagger a_l^\dagger a_T^\dagger |\text{vac}\rangle, a_j^\dagger a_m^\dagger a_T^\dagger |\text{vac}\rangle, a_k^\dagger a_l^\dagger a_T^\dagger |\text{vac}\rangle, a_k^\dagger a_m^\dagger a_T^\dagger |\text{vac}\rangle, a_l^\dagger a_m^\dagger a_T^\dagger |\text{vac}\rangle, \quad (18)$$

with amplitudes $\alpha_{j,k,T}, \alpha_{j,l,T}, \alpha_{j,m,T}, \alpha_{k,l,T}, \alpha_{k,m,T}$ and $\alpha_{l,m,T}$. Lemma 3.1 implies the following relations between these amplitudes:

$$\begin{aligned} (1) \alpha_{j,k,T} &= +u_{e'} \alpha_{k,l,T}, \quad (2) \alpha_{j,l,T} = -u_e \alpha_{k,l,T}, \quad (3) \alpha_{j,l,T} = +u_{e''} \alpha_{l,m,T}, \\ (4) \alpha_{j,k,T} &= +u_{e''} \alpha_{k,m,T}, \quad (5) \alpha_{j,m,T} = -u_e \alpha_{k,m,T}, \quad (6) \alpha_{j,m,T} = -u_{e'} \alpha_{l,m,T}. \end{aligned} \quad (19)$$

However, conditions (1-6) imply $+1 = -1$. Hence, $K_B |\psi\rangle$ cannot have support on n -mode states $|S\rangle = a_S^\dagger |\text{vac}\rangle$ with cluster particle number $\hat{N}_q = 2, 3, \dots, n_q - 2$. $\square$

We note that the previous Lemma is intrinsically fermionic, i.e. the anti-commutation of creation operators plays a role in the signs of Eq. (19) and there is no QUANTUM 2-SAT counterpart. The following corollary summarizes the consequences.

Corollary 3.5. *Consider an assignment of cluster-product form*

$$\left| \phi \left(\{N_q^i\}_{i \in \text{hPNC}}, \{P_q^j\}_{j \in \text{non-hPNC}}, \{x_k\}_{k \in \text{Class}} \right) \right\rangle, \quad (20)$$

in Eq. (17), which, per Proposition 3.3, exists if there is a satisfying assignment for G . We summarize some consequences for the quantum numbers N_q^i and P_q^i in the cluster-product form due to the quantum clauses on clusters.

1. For degree ≥ 3 hPNC clusters $i \in \text{hPNC}$, the only hidden cluster particle numbers can be $N_q^i = 0, 1, n_q^i - 1, n_q^i$ which are Gaussian states, remaining Gaussian when undoing the particle-hole transformation as argued in Section 2.2.
2. For any hPNC cluster $i \in \text{hPNC}$ with n_q^i modes, the hidden cluster particle number can always be $N_q^i = 0$ or $N_q^i = n_q^i$, since Lemma 3.1 imposes no constraints on the coefficients, i.e. particles don't propagate since there are either no particles or all modes are filled.
3. For a hPNC cluster i with n_q^i modes, let there be a cluster product assignment with hidden cluster particle number $N_q^i \neq 0, N_q^i \neq n_q^i$. Then the assignment will have support on states $|S\rangle$ in which any mode j in the cluster is empty ($x_j = 0$) and states for which it is filled ($x_j = 1$). Hence in this case any classical clause straddling the cluster on mode j and a classical mode $k \in \text{Class}$ or mode in another quantum cluster, will have to be satisfied (if it can) no matter the value for $x_j = 0, 1$. When $n_q^i > 2$, if the state with hidden cluster particle number $N_q^i \neq 0, N_q^i \neq n_q^i$ obeys any additional classical clause constraints inside the cluster, then either the all-empty $N_q^i = 0$ or all-filled assignment $N_q^i = n_q^i$ (or both) should also obey these constraints since on any pair of modes, at least 3 out of 4 occupations on these two modes will be allowed (thus always including either 00 or 11).
4. If G contains a degree ≥ 3 non-hPNC cluster $j \in \text{non-hPNC}$ with a number of modes $n_q^j \geq 5$, then G has no satisfying assignment. Via Corollary 3.2 the satisfying assignment is characterized by total parity P and has support on states with any hidden cluster particle number with this parity. However, Lemma 3.4 implies that the assignment cannot have support on odd hidden cluster particle number $N_q = 3$ (in case of odd parity) or $N_q = 2$ (in case of even parity) when $n_q \geq 5$, hence there cannot be any satisfying assignment.
5. For a degree ≥ 3 non-hPNC cluster $G_q \equiv j \in \text{non-hPNC}$ with $n_q^j = 4$, $K_B |\psi\rangle$ (with the set B defined through the graph T_q of the cluster G_q), can only have support on

states with odd cluster parity. Undoing the particle hole transformation, this excludes one choice of parity P for this cluster in Eq. (20) (which one depends on whether $|B|$ is odd or even). In Appendix B we show that the satisfying assignment for such cluster of four fermionic modes is a non-Gaussian fermionic state.

6. For a non-hPNC G_q cluster with n_q modes, let there be an cluster product assignment with some cluster parity P_q . Then this assignment will have support on states $|S\rangle$ in which any mode j in the cluster is empty and states for which it is filled: it has support on all occupations x_j for any j in the cluster, with the given fixed parity by Corollary 3.2. Hence any classical clause straddling the cluster on mode j and a classical mode $k \in \text{Class}$ or mode in another quantum cluster, will have to be satisfied (if it can) no matter the value for $x_j = 0, 1$.

The corollary means that if a vertex in the cluster is of degree at least 3, then, if the cluster is non-hPNC, we must have $n_q \leq 4$, and hence it is easy to verify whether an assignment exists. Otherwise, if it is hPNC and has a degree ≥ 3 vertex, its assignments are simple and Gaussian as they are very restricted in particle number (in the particle-hole transformed basis), collapsing the search space to being linearly dependent on the number of modes in the cluster. We capture this efficiency more precisely in Lemma 3.10 in the next section.

If all vertices in the cluster G_q are of degree less than 3, the graph G_q is a line or a loop. For a line, one can apply the Jordan-Wigner transformation and map the problem onto quantum 2-SAT on a line which can fully be solved, using existing methods. For a loop there is a small adaptation due to the Jordan-Wigner Pauli-Z strings. The cases of G_q being a line or a loop are dealt with in Section 3.3.1 below.

Corollary 3.5 also makes clear that the assignments of hidden particle number and parity on the cluster lead to strong restrictions on whether and how one can satisfy the remaining classical clauses in E_c , and this will be used to prove our main Theorems 1.4, 1.5 and 1.6.

3.3.1 Quantum clusters: lines and loops

To handle lines and loops, it is convenient to switch to qubit language as FERMIONIC 2-SAT projectors on a line map directly onto QUANTUM 2-SAT projectors on qubits via the Jordan-Wigner transformation in Eq. (2).

Let us first consider the taxonomy of graphs G_q that are lines or loops, shown in Figure 4. Remember that there is at most a single $\Pi_e^{02,q}$ or single $\Pi_e^{1,q}$ per edge e , otherwise one can reduce the pair to a classical edge, see the arguments in Section 2.1.1, and we do not yet consider any classical clauses in this section yet. If G_q is a line, its maximal spanning hPNC subgraph T_q (see Definition 2.3) is also a line, and G_q is either hPNC or non-hPNC as in Figure 4(a) and (b). If G_q is a loop, then T_q is either a loop, see Figure 4(c) and (d), or a line, Figure 4(e) and (f).




Figure 4: Different types of line and loop clusters in the particle-hole-transformed basis. The black edges make up T_q (representing $\Pi_e^{1,q}$ -type clauses) and the blue edges are $\Pi_e^{02,q}$ -type clauses.

First, we consider the case of G_q being a line. The following Lemma is to be applied to T_q in the particle-holed transformed basis, thus containing only $\Pi^{1,q}$ -type clauses. It is essentially a rephrasing of some results in [11].

Lemma 3.6. *Given a line with n_q fermionic modes and (rank-1) clauses $\Pi_{(i,i+1)}^{1,q}$ between modes i and $i+1$ for $i \in \{1, 2, \dots, n_q - 1\}$. For each particle number $N_q \in \{0, 1, \dots, n_q\}$, there is a unique assignment satisfying the clauses $\Pi_{(i,i+1)}^{1,q}$.*

Proof. Via the Jordan-Wigner transformation in Eq. (2), each 2-mode rank-1 projector $\Pi_{(i,i+1)}^{1,q}$ becomes a projector onto a two-qubit quantum state $\Pi_{i,i+1}^{\text{qubit}} = |\psi\rangle\langle\psi|_{i,i+1}$ with $|\psi\rangle_{(i,i+1)} = \alpha|01\rangle + \beta|10\rangle$ (where $|\alpha|^2 + |\beta|^2 = 1$). The entangled state $|\psi\rangle_{(i,i+1)}$ can generally be expressed as $B_i \otimes I |\Psi^-\rangle$ with a two-qubit singlet state $|\Psi^-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$ and some diagonal, invertible B_i .

Let $H = I \otimes A_2 \dots \otimes A_{n_q} H_{\text{FM-Heis}} I \otimes A_2^\dagger \dots \otimes A_{n_q}^\dagger \geq 0$, with $H_{\text{FM-Heis}} = \sum_i |\Psi^-\rangle\langle\Psi^-|_{i,i+1}$, the 1D ferromagnetic Heisenberg model. Here A_2 is uniquely chosen such that $I \otimes A_2 |\Psi^-\rangle_{1,2} = B_1 \otimes I |\Psi^-\rangle_{1,2} = |\psi\rangle_{1,2}$. Given A_2 , A_3 is chosen such that $A_2 \otimes A_3 |\Psi^-\rangle_{2,3} = B_2 \otimes I |\Psi^-\rangle_{2,3} = |\psi\rangle_{2,3}$ etc. Importantly, the satisfying assignments of $\{\Pi_{(i,i+1)}^{\text{qubit}}\}$ are the zero-eigenvalue states of H . A zero-eigenvalue state with particle (or "excitation") number N_q can thus be obtained by applying $I \otimes A_2^{\dagger-1} \otimes \dots \otimes A_{n_q}^{\dagger-1}$ to any state $|\psi_{\text{FM-Heis}}^{N_q}\rangle$ in the null-space of $H_{\text{FM-Heis}}$, with equal particle number N_q as the $A_i^{\dagger-1}$ matrices are diagonal. Hence, given a null-space solution with excitation number N_q for $H_{\text{FM-Heis}}$, one obtains a satisfying assignment for the fermionic problem on a line.

Due to the $U^{\otimes n_q}$ -invariance of $H_{\text{FM-Heis}}$, the eigenspaces must be spanned by states of the form $|\chi^{\otimes n_q}\rangle$, hence symmetric subspaces of n_q qubits [12]. There is only one such state with a given particle number N_q , namely the fully permutation-symmetric state, which is in the null-space of $H_{\text{FM-Heis}}$ since $|\Psi^-\rangle\langle\Psi^-|_{i,i+1} |\psi_{\text{FM-Heis}}^{N_q}\rangle = |\Psi^-\rangle\langle\Psi^-|_{i,i+1} \pi(i, i+1) |\psi_{\text{FM-Heis}}^{N_q}\rangle = -|\Psi^-\rangle\langle\Psi^-|_{i,i+1} |\psi_{\text{FM-Heis}}^{N_q}\rangle = 0$ where $\pi(i, i+1)$ is the permutation between qubits i and $i+1$. Note that the unique satisfying assignment for the quantum 2-SAT clauses is thus $I \otimes A_2^{\dagger-1} \otimes \dots \otimes A_{n_q}^{\dagger-1} |\psi_{\text{FM-Heis}}^{N_q}\rangle$. $\square$

Let us now consider a hPNC cluster for which both G_q and T_q are a loop as in Fig. 4(c). Again the following Lemma is to be applied to T_q in the particle-hole-transformed basis in which it consists of $\Pi_e^{1,q}$ clauses only. For there to be any satisfying assignments on T_q , the clause closing the loop needs to be consistent with the other clauses. For FERMIONIC 2-SAT clauses this consistency condition also turns out to depend on the cluster parity:

Lemma 3.7. *Given a loop consisting of n_q fermionic modes with clauses $\Pi_{(i,i+1 \bmod n_q)}^{1,q}$ for $i \in \{1, 2, \dots, n_q\}$, with coefficients $u_{(i,i+1 \bmod n_q)}$ defined in Eq. (6). Besides the satisfying assignments at $N_q = 0$ (all-empty) and $N_q = n_q$ (all-filled), there are satisfying assignments at a given N_q iff $(-1)^{N_q+n_q-1} [1/u_{(1,n_q)} \prod_{i=1}^{n_q-1} u_{(i,i+1)}] = 1$. So, if $|1/u_{(1,n_q)} \prod_{i=1}^{n_q-1} u_{(i,i+1)}| = 1$, there are only satisfying assignments at all even N_q (when n_q is odd) or odd N_q (when n_q is even).*

Proof. The proof of Lemma 3.7 is similar to the proof of Lemma 3.6. The only difference is that in the case of a loop, the clause $\Pi_{(1,n_q)}^{1,q}$, when Jordan-Wigner transformed, becomes

$$2\Pi_{(1,n_q)}^{\text{qubit}} = I_1 \otimes B_{n_q} \left(|1\rangle_1 |0\rangle_{n_q} - Z_2 \dots Z_{n_q-1} |0\rangle_1 |1\rangle_{n_q} \right) \times \left(\langle 1|_1 \langle 0|_{n_q} - Z_2 \dots Z_{n_q-1} \langle 0|_1 \langle 1|_{n_q} \right) I_1 \otimes B_{n_q}^\dagger, \quad (21)$$

with B_{n_q} diagonal and $(B_{n_q})_{11}/(B_{n_q})_{22} = -\beta_{(1,n_q)}/\gamma_{(1,n_q)} = -1/u_{(1,n_q)}^*$ (see Eq. (6)). The satisfying assignments of the loop problem correspond to zero-energy states of H in the proof of Lemma 3.6 that are also projected to zero by $\Pi_{(1,n_q)}^{\text{qubit}}$. We have that $(B_i)_{11}/(B_i)_{22} = -\gamma_{(i,i+1)}/\beta_{(i,i+1)} = -u_{(i,i+1)}^*$ for $i \in \{1, 2, \dots, n_q - 1\}$, from which the entries of $A_2, \dots, A_{n_q}$ follow (see the proof of Lemma 3.6). The relation $\Pi_{(1,n_q)}^{\text{qubit}} I \otimes A_2^{\dagger-1} \otimes \dots \otimes A_{n_q}^{\dagger-1} \left| \psi_{\text{FMHeis}}^{N_q} \right\rangle = 0$ holds iff $u_{(1,n_q)} = (-1)^{N_q} (A_{n_q})_{22}^*/(A_{n_q})_{11}^*$ is obeyed (where we have used $(B_{n_q})_{11}/(B_{n_q})_{22} = -1/u_{(1,n_q)}^*$). Using that the entries of $A_2, \dots, A_{n_q}$ can be expressed in terms of $\{u_{(i,i+1)}\}_{i=1}^{n_q-1}$, this condition is equivalent to the condition $(-1)^{N_q+n_q-1} [1/u_{(1,n_q)} \prod_{i=1}^{n_q-1} u_{(i,i+1)}] = 1$ from the lemma statement. $\square$

3.3.2 non-hPNC lines and loops

We have thus far considered the hPNC line or loop graphs. For non-hPNC lines and loops, however, we must consider the effect of the $\Pi_e^{02,q}$ clauses (in the particle-hole-transformed basis) on the satisfiability of these clusters as in Fig. 4(b),(d), (e) and (f). These $\Pi_e^{02,q}$ clauses can be of two types: they can either be part of a rank-2 projector as in Fig. 4(d), or they can close a loop as in Figure 4(e). We will refer to the former as a *non-loop-closing* $\Pi_e^{02,q}$ clause and to the latter as a *loop-closing* $\Pi_e^{02,q}$ clause. Of course, any loop G_q contains at most one loop-closing $\Pi_e^{02,q}$ clause. If G_q is a line, then all $\Pi_e^{02,q}$ clauses are non-loop-closing as in Fig. 4(b). If G_q and its maximal spanning hPNC subgraph T_q are loops, all $\Pi_e^{02,q}$ clauses are also non-loop-closing (Fig. 4(d)).

Note that one can always apply a different particle-hole transformation (with a different T_q) such that the instance in Figure 4(f) becomes an instance in Figure 4(d), by ensuring that the loop is closed on a rank-2 edge. Therefore, if G_q contains a loop-closing $\Pi_e^{02,q}$ clause, then wlog it is the only $\Pi_e^{02,q}$ clause in G_q .

Lemma 3.8 below provides a detailed account of conditions under which a parity P_q is allowed on a quantum cluster G_q , provided that P_q is allowed on T_q . The important take-away is that for any non-hPNC line or loop, these conditions come down to straightforward relations between the u_e and v_e coefficients of respectively $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$ clauses along the line or loop.

Lemma 3.8. *Given a non-hPNC cluster G_q which is a line or a loop on n_q modes. The following statements hold in the particle-hole-transformed basis in which T_q only has $\Pi_e^{1,q}$ clauses and P_q refers to the parity of the particle-hole-transformed assignment. Assume that T_q allows for satisfying assignments of parity P_q . If G_q contains a single $\Pi_e^{02,q}$ clause, then the same parities of T_q are allowed as satisfying assignments for G_q . If G_q contains at least two $\Pi_e^{02,q}$ clauses, then suppose:*

1. G_q is a line (Figure 4(b)) then any P_q of T_q is allowed for G_q iff the following consistency condition is satisfied. For each pair of subsequent non-loop-closing $\Pi_e^{02,q}$ clauses on edges $(j, j+1)$ and $(k, k+1)$ ($j < k$),

$$(\star) \quad v_{(j,j+1)} = \left[\prod_{i=j}^{k-1} u_{(i,i+1)} u_{(i+1,i+2)} \right] v_{(k,k+1)} \quad (22)$$

must hold.

2. G_q is a loop and T_q a line (Figure 4(e)). P_q is always allowed on G_q since there is only one $\Pi_e^{02,q}$ clause in G_q .

3. Both G_q and T_q are loops (Figure 4(d)). Any P_q of T_q is allowed for G_q , if condition $(\star)$ holds on $\Pi_e^{02,q}$ clauses along the loop, with the following caveat. If one of the non-loop-closing clauses is on edge $(1, n_q)$, then the one condition (on subsequent pairs) involving $\Pi_{(1, n_q)}^{02,q}$ has $v_{(1, n_q)}$ replaced by $-1/u_{(1, n_q)}^2 v_{(1, n_q)}$ on the RHS of $(\star)$.

Proof. By assumption, P_q is allowed on T_q . According to Corollary 3.2, if P_q is allowed on G_q , then a satisfying assignment has support on all states $|S\rangle = a_S^\dagger |\text{vac}\rangle$ with parity P_q . For P_q to be allowed on G_q , the $\Pi_e^{02,q}$ clauses on G_q should not lead to inconsistencies when connecting these states of parity P_q via Lemma 3.1. Such inconsistencies can come up when, starting with any $|S\rangle$, one repeatedly applies Lemma 3.1 to end up at state $|S\rangle$ again. If the associated coefficient relations in Lemma 3.1 are not consistent, P_q is excluded on G_q . For G_q lines and loops containing a single $\Pi_e^{02,q}$ clause, these inconsistencies are always avoided. For non-hPNC lines and loops with multiple $\Pi_e^{02,q}$ clauses, the consistency of these relations can be checked rather straightforwardly. To see this, we distinguish between cases (1), (2) and (3) in the lemma statement.

For case (1), all $\Pi_e^{02,q}$ clauses are non-loop-closing. The following process, realized by repeated application of Lemma 3.1, can lead to inconsistency. Start off with $|S\rangle$ with equal occupation on an edge with a $\Pi_e^{02,q}$ clause. One creates a pair of particles/holes on edge e , migrates them to a clause $\Pi_{e'}^{02,q}$ at edge e' via $\Pi_e^{1,q}$ clauses, and undoes the creation of the particle/hole pair through clause $\Pi_{e'}^{02,q}$ to end up at $|S\rangle$. Provided that the condition in the lemma statement holds, this can never lead to inconsistency. Ensuring that the condition holds on *subsequent* $\Pi_e^{02,q}$ clauses along G_q is sufficient to ensure the consistency between non-subsequent $\Pi_e^{02,q}$ clauses.

For case (2), the $\Pi_e^{02,q}$ clauses cannot lead to inconsistency, since there is only one such clause in G_q .

For case (3), the same process as for case (1) determines the consistency conditions. However, there can be a non-loop-closing $\Pi_{(1, n_q)}^{02,q}$ clause (i.e., on edge $(1, n_q)$). Again using Lemma 3.1, it can be shown straightforwardly that the altered $(\star)$ condition implies consistency between the $\Pi_{(1, n_q)}^{02,q}$ clause and its subsequent $\Pi_e^{02,q}$ clause along the loop. $\square$

Importantly, we can show that satisfying assignments on non-hPNC clusters which are lines or loops are always fermionic Gaussian states of fixed parity. This is argued using the fact that a $\Pi_e^{1,q} + \Pi_e^{02,q}$ projector is quadratic, see Section 2.1.1, and arguing that one can always add extra clauses onto edges in the cluster so that all edges become of this form $\Pi_e^{1,q} + \Pi_e^{02,q}$. Satisfying assignments must then lie in the null-space of a quadratic fermionic (positive semi-definite) Hamiltonian, hence be a Gaussian state.

Lemma 3.9. *Given a non-hPNC cluster G_q which is a line or loop. The satisfying assignments (if they exist) on this cluster are Gaussian states.*

Proof. Let us first consider line or loop clusters (in the particle-hole transformed basis) for which the only $\Pi_e^{02,q}$ clauses are non-loop-closing. There can be multiple such clauses (i.e., cases (1) and (3) in Lemma 3.8), or just a single one. In the former case, we assume Eq. (22) is satisfied (otherwise, the cluster would not be satisfiable). Let us add a non-loop-closing $\Pi_e^{02,q}$ clause to an edge not yet containing one. If this $\Pi_e^{02,q}$ clause also obeys Eq. (22), then the satisfying assignments of the original problem are still satisfying assignments. One can keep adding $\Pi_e^{02,q}$ clauses in this manner until *all* edges in the cluster are of type $\Pi_e^{1,q} + \Pi_e^{02,q}$. In that case, all cluster projectors are quadratic and hence the satisfying assignments are Gaussian states as per observations in Section 2.1.

What is left is to consider clusters whose only $\Pi_e^{02,q}$ clause is a loop-closing $\Pi_e^{02,q}$ clause (i.e., Figure 4(e) and case (2) in Lemma 3.8). Importantly, a quantum cluster with exactly the same satisfying assignments can be obtained by deleting the loop-closing $\Pi_e^{02,q}$ clause, and adding an appropriate non-loop-closing $\Pi_e^{02,q}$ clause anywhere along T_q . Hence we have obtained a cluster of type (1) in Lemma 3.8, for which we have already shown that the lemma statement holds. $\square$

3.3.3 Efficient verification of allowed hidden particle number or parity on a quantum cluster

The following Lemma captures that verifying what cluster particle numbers and parities are allowed in the cluster-product form assignment, Eq. (20), already restricted by the results in the previous section, is classically efficient.

Lemma 3.10 (Cluster checks). *Given a quantum cluster $G_q = (V_q, E_q)$ consisting of n_q modes and m_q quantum clauses. If G_q is an hPNC cluster, one can check in $O(n_q + m_q)$ time which hidden cluster particle numbers N_q (being the cluster particle number of $K_B |\psi\rangle$, with $|\psi\rangle$ a satisfying assignment) correspond to satisfying assignments. If G_q is a non-hPNC cluster, one can check in $O(n_q + m_q)$ time which cluster parities correspond to satisfying assignments.*

Proof. First, suppose G_q is a hPNC cluster. By Corollary 3.5 point 2, hidden cluster particle numbers $N_q = 0$ and $N_q = n_q$ are *always* permitted.

If the hPNC cluster G_q has a vertex of degree at least 3, then by Corollary 3.5 point 1, to see if a single-particle state with $N_q = 1$ and/or single-hole state with $N_q = n_q - 1$ are permitted, one simply lists all m_q partnering constraints (Lemma 3.1) between n_q coefficients in the potential satisfying assignment and checks the consistency of these constraints in time $O(n_q + m_q)$. On the other hand, suppose the hPNC cluster G_q has no vertex of degree at least 3. If G_q is a line, all hidden cluster particle numbers are allowed via Lemma 3.6, and the corresponding cluster particle states are unique. In case G_q is a loop, then all cluster particle numbers for which the condition in Lemma 3.7 is obeyed, are permitted, with the caveat that $N_q = 0$ and $N_q = n_q$ are permitted in any case.

Now suppose that G_q is a non-hPNC cluster. If G_q has a vertex of degree at least 3, then by Corollary 3.5 any assignment is excluded for $n_q \geq 5$, one parity is excluded for $n_q = 4$, and obviously n_q cannot be ≤ 3 . Clearly, one can efficiently verify the existence of an assignment for $n_q = 4$. If the non-hPNC cluster G_q does not have a vertex of degree at least 3, then G_q can be a line or a loop and which (if any) parity is allowed for those clusters can be checked in $O(n_q + m_q)$ time using the conditions in Lemmas 3.7 and 3.8. $\square$

4 An efficient classical algorithm for Fermionic 2-SAT (with and without fixed parity)

4.1 Solving Fermionic 2-SAT: proof of Theorem 1.4

Let us now prove Theorem 1.4, which states that FERMIONIC 2-SAT can be solved in time $O(n + m)$. Solving FERMIONIC 2-SAT comes down to performing certain checks on properties of quantum clusters, and running a (classical) 2-SAT algorithm based on these constraints and the remaining classical clauses:

Proof of Theorem 1.4. Due to Proposition 3.3 we can restrict the assignment to cluster-product form. We will be constructing a classical 2-SAT instance, which will be used

for deciding whether the FERMIONIC 2-SAT instance is satisfiable. This classical 2-SAT instance is denoted by $\mathcal{C}$. First, let us consider only the constraints from quantum edges.

1. We first consider all hPNC clusters. For these clusters, Corollary 3.5 point 2 implies that there are *always* two classical assignments that satisfy all quantum clauses in the cluster and that are each other's negations (see Figure 5 for an illustration). Now consider the set of classical clauses, either internal to the cluster or straddling between the cluster and some classical modes or between the cluster and some other quantum cluster. By Corollary 3.5 point 3 for $n_q > 2$, obeying any internal classical clauses on the cluster in any assignment with hidden particle number $N_q \neq 0, N_q \neq n_q$ would also allow for two classical assignments $N_q = 0, n_q$. Hence wlog we can assume that we have to choose from these classical assignments for any $n_q > 2$ cluster when constructing a global satisfying assignment. One can impose this using a set of 2-SAT clauses whose two unique satisfying assignments are these two classical assignments (which are each other's negations): note that we need to apply the particle-hole transformation K_B to obtain these 2-SAT clauses. These 2-SAT clauses, the classical clauses internal to the cluster and any straddling clauses are included in $\mathcal{C}$. The case $n_q = 2$ is dealt with separately for convenience. For $n_q = 2$, if there is a single clause and no internal classical clause, we do the same as for $n_q > 2$. If $n_q = 2$ with one $\Pi_e^{1,q}$ clause and one $\Pi_e^{02,q}$ clause, the cluster is non-hPNC and we deal with it in the next point 2. If $n_q = 2$ and the projector is rank-3, then it is hPNC, with 2 classical internal clauses and there is at most one satisfying Gaussian, non-classical assignment. For example, the two classical clauses exclude the assignments 00 and 11, allowing only the state $(\alpha a_1^\dagger + \beta a_2^\dagger) |\text{vac}\rangle$ orthogonal to some $\Pi_{(1,2)}^{1,q}$ clause. Hence, in this case only, one has to choose this non-classical (Gaussian) assignment and handle the possible classical clauses straddling this cluster via Corollary 3.5 point 3: a straddling clause on mode 1 in the cluster, i.e. $(x_1 \vee u)$ or $(\bar{x}_1 \vee u)$, gets replaced by clause (u) with literal u on some classical mode (and similarly for mode 2 in the cluster). These latter clauses are included in $\mathcal{C}$.
2. For all quantum clusters that are non-hPNC, we first check via Lemma 3.10 in $O(n + m)$ time whether there exists a satisfying assignment. Corollary 3.5 point 6 says that if there exists a satisfying assignment, there can be no additional internal classical clauses (as the state has support on all occupations with given parity). We handle any straddling classical clauses by replacing $(x_m \vee u_i)$ or $(\bar{x}_m \vee u_i)$, with mode m inside the non-hPNC cluster and classical literal u_i , by clauses (u_i) . These latter clauses are included in $\mathcal{C}$. The case when the classical clause is straddling between two quantum clusters can be dealt with similarly.

After these preprocessing steps in which we may find that there can't be a satisfying assignment (then output NO), let the collection of classical clauses obtained via the preprocessing, as well as the remaining clauses on solely classical modes, be the classical 2-SAT instance $\mathcal{C}$. We then solve instance $\mathcal{C}$ in time $O(n + m)$ [1]. $\square$

4.2 Solving Fermionic 2-SAT with fixed parity: proof of Theorem 1.5

Let us now address the question of whether fixing the parity changes the complexity of the FERMIONIC 2-SAT problem. We prove Theorem 1.5 in this section, which states that this problem can still be solved efficiently, in time $O(nm)$. Our algorithm for solving FERMIONIC 2-SAT with fixed parity $P \in \{-1, +1\}$ consists of running certain efficient checks on the allowed satisfying assignments on quantum clusters as a preprocessing step




Figure 5: Satisfying assignments on non-hPNC quantum clusters are such that each mode has to be both empty *and* occupied, see Corollary 3.5 point 6. On hPNC clusters, one only has to assign either of two classical states to check whether the instance is satisfiable (apart from a special case when the cluster only has two modes, $n_q = 2$), when we do not aim to find an assignment of given parity. The black vertices and edges respectively denote modes that are not in any quantum cluster (i.e., classical modes) and classical clauses.

to construct a (classical) 2-SAT with fixed parity instance, on which one runs the $O(nm)$ time algorithm developed in Theorem C.1 in Appendix C.

Importantly, non-hPNC clusters are treated differently than hPNC clusters in this next proof. Namely, any assignment in a non-hPNC cluster satisfies the classical cluster-straddling clauses (or not) *independent* of its parity, while for hPNC clusters, assignments of different parity can lead to different ways of satisfying the straddling classical clauses.

Proof of Theorem 1.5. By Proposition 3.3 an assignment of fixed parity is of cluster-product form. We proceed in three steps.

1. We run the $O(nm)$ time algorithm (developed in Appendix C) for solving the (classical) 2-SAT instance $\mathcal{C}$ obtained in the proof of Theorem 1.4 via items 1 and 2, with fixed parity. If the algorithm outputs an assignment of the right parity, we are done and output YES. If there is no assignment, output NO. If there is an assignment but with the wrong parity (which the algorithm will tell you), let's call the assignment X_0 , and proceed to the next step.
2. We consider whether it is possible to flip the parity of any non-hPNC cluster in X_0 . No variables internal to a non-hPNC cluster enter the instance $\mathcal{C}$ due to how we handle the straddling clauses in item 2 in the proof of Theorem 1.4. Therefore, if we can flip the parity of a single cluster $G_q^j \in \text{non-hPNC}$, the assignment on $G \setminus G_q^j$ of X_0 is still valid, and hence we have full assignment with the correct parity. According to Corollary 3.2, each non-hPNC cluster has at most one satisfying assignment per parity. Which parities (one or both) are allowed for each non-hPNC cluster can be checked in time $O(n+m)$ via Lemma 3.10. So we iterate over all non-hPNC clusters and if there is one whose parity can be flipped, we are done (output YES). If not, we proceed to the next step.
3. The next step pertains to all hPNC clusters, except the $n_q = 2$ clusters with a rank-3 edge which were treated differently in obtaining $\mathcal{C}$ in item 1 of the proof of Theorem 1.4: we will not attempt flipping parity on these clusters as there is only one possible assignment.

Due to the failure of obtaining a flipped parity assignment in the previous two steps, we know that it is not possible to change the parity by flipping classical modes and/or flipping the hidden particle number $N_q = 0$ (all-empty) to $N_q = n_q$ (all-filled) assignment on any hPNC cluster (which are the possible assignments inside the hPNC clusters in $\mathcal{C}$ which are examined in step 1). This has the following consequence. We define the classical 2-SAT instances $\mathcal{C}^{+,M}$ which is a superset of the classical clauses of $\mathcal{C}$ and some classical clauses which pertain to a subset of non-hPNC clusters $M \subseteq \text{hPNC}$. That is, $\mathcal{C}^{+,M}$ is obtained by adding to the collection of clauses in $\mathcal{C}$, for each cluster $G_q^j \in M$, the following: replace each straddling clause between a mode $i \in G_q^j$ and classical mode, of the form $(x_i \vee u)$ or $(\bar{x}_i \vee u)$ by clause (u) with u a literal of the classical mode. One does the same for any straddling clause with another hPNC quantum cluster provided it is not in the set M (if it is non-hPNC or in the set M , there would be no satisfying assignment for $\mathcal{C}^{+,M}$). Due to failure in the previous two steps, either an instance $\mathcal{C}^{+,M}$ will have no satisfying assignment, or an assignment of the wrong parity, i.e. the same as X_0 , for any M , since any satisfying assignment of $\mathcal{C}^{+,M}$ is also a satisfying assignment for $\mathcal{C}$. Hence, the only way to still obtain a flipped parity assignment as compared to X_0 is to possibly flip the parity of some individual hPNC cluster by *not choosing* the hidden particle number $N_q = 0$ or $N_q = n_q$ assignment on it, as we did for $\mathcal{C}$. Note that if we do not constrain the overall parity, as in Theorem 1.4, there was never a reason to consider other possible assignments (see also Fig. 5).

Hence, we first consider for each hPNC cluster G_q^i whether one can flip its parity, and then if so, we decide whether an altered classical 2-SAT instance $\mathcal{C}_{G_q^i}$ has a satisfying assignment using a 2-SAT solver. To check whether one can flip its parity, observe that no internal classical clauses must be present as any assignment with hidden particle numbers $N_q \neq 0, N_q \neq n_q$ has the property that each mode j in the cluster can be both 0 and 1. If classical clauses are present and/or the parity cannot be flipped for other reasons, we move onto the next hPNC cluster. As for checking on a hPNC cluster whether an assignment of flipped parity is allowed, given its quantum clauses, Lemma 3.10 shows that this can be done efficiently, in $O(n + m)$ time. For a hPNC cluster which is a line or a loop, in principle non-Gaussian assignments with hidden particle number $N_q \neq 0, 1, n_q - 1, n_q$ may be allowed. However, one can observe that, if a flipped parity assignment is possible, one can take this assignment actually to be a Gaussian state, i.e. one only with hidden particle number $N_q = 0, 1, n_q - 1, n_q$.

If the parity can be flipped for cluster G_q^i , we construct $\mathcal{C}_{G_q^i}$. To get from $\mathcal{C}$ to $\mathcal{C}_{G_q^i}$ we (1) replace all straddling clauses between a mode j in cluster G_q^i and a classical mode $(x_j \vee u)$ and $(\bar{x}_j \vee u)$ by clause (u) . One does the same for a straddling clause with a mode in a hPNC quantum cluster (if there is a straddling clause with a non-hPNC quantum cluster, then there won't be a satisfying assignment when flipping cluster G_q^i). From $\mathcal{C}$ we remove all classical clauses which constrain the assignment on G_q^i to have hidden particle number $N_q = 0$ or $N_q = n_q$ as described in item 1 in the proof of Theorem 1.4: there are no more variables in $\mathcal{C}_{G_q^i}$ which pertain to $\mathcal{C}_{G_q^i}$. If $\mathcal{C}_{G_q^i}$ has an assignment, then it must now be an assignment of the right parity and we output YES. If not, for all hPNC clusters that we iterate over, we output NO, and this process is clearly efficient.

□

5 PNC Fermionic 2-SAT with fixed particle number is NP-complete

Let us now consider the complexity of solving Problem 1.3, which corresponds to deciding PNC FERMIONIC 2-SAT with particle number fixed to some $N \in \{0, 1, \dots, n\}$. In contrast to fixing the parity, we show that fixing the particle number to N makes FERMIONIC 2-SAT NP-complete (Theorem 1.6).

Proof of Theorem 1.6. Proposition 3.3 and Corollary 3.5 show that satisfying assignments of PNC FERMIONIC 2-SAT are of cluster-product form. In the PNC version of FERMIONIC 2-SAT, quantum clauses are only of type $\Pi_e^{1,q}$. As a consequence, all quantum clusters are PNC, even without performing a particle-hole transformation. Therefore, we can characterize satisfying assignments on quantum clusters by the particle number on that cluster. A witness W thus consists of a list of $O(n)$ integers $W = (\{N_q^i\}_{i \in \text{hPNC}}, \{x_k\}_{k \in \text{Class}})$. The verifier algorithm verifies for each quantum cluster i , using Lemma 3.10 whether N_q^i is indeed an allowed particle number on that cluster. Note that the verifier does not need to construct or have the actual fermionic state to do this. Using the collection of particle numbers on quantum clusters N_q^i and Corollary 3.5, she infers for each mode in a quantum cluster whether it is empty, occupied or both. Combined with $\{x_k\}_{k \in \text{Class}}$, it is then verified whether all classical clauses $e \in E_c$ are indeed satisfied. As a final step, the verifier checks if the total particle number of the witness W is indeed the given N . This verification algorithm is efficient. To show NP-hardness, let us consider an instance of FERMIONIC 2-SAT consisting of monotone classical clauses only, i.e. only clauses of type $(x_i \vee x_j)$. Deciding this problem with fixed Hamming weight N is equivalent to deciding whether there is vertex cover with N vertices, which is NP-complete. Indeed, it is known that classical 2-SAT with fixed Hamming weight, i.e. *the weighted 2-SAT problem* (W2SAT) is NP-complete, see <https://en.wikipedia.org/wiki/2-satisfiability>. $\square$

Remark (Classical Assignments). *Part 1 of the proof of Theorem 1.4 for hPNC clusters and classical modes shows that PNC FERMIONIC 2-SAT (which only has $\Pi_e^{1,q}$ quantum clauses and thus hPNC quantum clusters) without any particle number constraint always has a classical satisfying assignment, i.e., a satisfying assignment of the form $a_S^\dagger |\text{vac}\rangle$, with the exception of 2-mode quantum clusters with rank-3 projectors in Gaussian states. Note that this is similar to quantum 2-SAT always having a product assignment with the exception of rank-3 edges [3]. In contrast, when fixing particle number to some integer N as in Problem 1.3, there are instances with only non-Gaussian satisfying assignments, like the satisfying assignments on a line for $1 < N_q < n_q - 1$ in Lemma 3.6.*

6 Complexity of Fermionic k-SAT and related problems

Let us provide some background and perspective on the (known) complexity of related fermionic problems. Using qubit-to-fermion mappings, such as the one in Ref. [13], we can straightforwardly argue the following, see Appendix D for the proof.

Lemma 6.1. *FERMIONIC k -SAT $\in$ QMA and FERMIONIC k -SAT is QMA₁-hard for $k = 9$.*

This Lemma mirrors the results for QMA₁-hardness of QUANTUM k -SAT for $k \geq 3$ [5, 2].

One may expect that FERMIONIC k -SAT $\in$ QMA₁ for all variants of the problem (particle or non-particle number conserving, with a fixed particle number or not), but one has to be cautious about what basic gates are used in the class QMA₁ (usually H , CNOT

and T) and whether this set of gates allows for a zero-error acceptance in the YES case when the fermionic problem and its specification is mapped to qubits.

It may be possible to reduce $k = 9$ in Lemma 6.1 to a lower k by adapting the space-time circuit-to-Hamiltonian construction in Section 3.3 in [14] which proves QMA-completeness of a fermionic circuit Hamiltonian with projectors which involve at most 4 fermionic modes ($k = 4$), *under the restriction* that there is 1 particle per track on the 2D lattice. Thus we don't have an overall particle constraint, but several fixed particle sectors, which can possibly be shown to be equivalent to an overall constraint.

Let us also mention a result that is related to this work. It is known that the 'FERMIONIC MAX-2-SAT' problem *with* particle number constraint is QMA-complete, see the next Theorem 6.2. Note that this problem is fundamentally different from FERMIONIC 2-SAT, since the former is a ground-state energy estimation problem and not a question of whether a given Hamiltonian is frustrated or not.

Theorem 6.2 (Theorem 2 in [15]). *Determining the ground state energy with $1/\text{poly}(n)$ precision for a class of Fermi-Hubbard Hamiltonians with n fermionic modes, with particle number $N = n/2$ (half-filling) is QMA-complete. The particle-number-conserving Fermi-Hubbard Hamiltonian on a graph $G = (V, E)$ is*

$$H_{\text{FH}} = U \sum_{i \in V} n_{i,+} n_{i,-} + \sum_{(i,j) \in E} \sum_{\sigma = \pm} t_{i,j} (a_{i,\sigma}^\dagger a_{j,\sigma} + a_{j,\sigma}^\dagger a_{i,\sigma}), \quad (23)$$

with $n_{i,+} = a_{i,+}^\dagger a_{i,+}$, $U, t_{i,j} \in \mathbb{R}$. For QMA-completeness, bounds are specified on the parameters U and $\{t_{i,j}\}$.

7 Discussion

Interestingly, it is not clear whether there is an efficient classical algorithm to solve instances of QUANTUM 2-SAT (with only parity-conserving projectors) when we ask for an assignment with fixed parity: in this case there can be non-product satisfying assignments. In addition, we don't have the "natively-fermionic" Lemma 3.4 for QUANTUM 2-SAT. We conjecture that this problem is of different complexity than FERMIONIC 2-SAT with fixed parity which we have proved to be efficiently solvable in Theorem 1.5. This would mesh elegantly with the point of view that parity-conserving interactions are fundamental for fermions.

We have seen that the satisfying assignments of FERMIONIC 2-SAT are of cluster-product form with some modes with purely classical occupations, some modes in a possibly-large clusters in Gaussian states of fixed parity or hidden fixed particle number, and some modes in a 4-fermion cluster in a non-Gaussian state. It will be interesting to explore how we can use these results for FERMIONIC 2-SAT to develop approximation algorithms or heuristic strategies, quantum or classical, to solve FERMIONIC k -SAT for $k > 2$. For FERMIONIC ($k > 2$)-SAT one may expect that satisfying assignments can be genuinely many-mode non-Gaussian states: due to its QMA₁-hardness, satisfying assignments for general FERMIONIC k -SAT problems are not expected to be classically efficiently describable. Thus for such problems, one can seek quantum heuristic strategies, i.e. a quantum equivalent of classical heuristic SAT solvers, which aim at constructing a satisfying state on a quantum computer: such strategies could build on the nature of satisfying assignments for the (FERMIONIC) 2-SAT problem. It is an open question whether there are interesting classical mathematical problems which can be formulated as a question about

the existence of a satisfying assignment to a finite-size, quantum or FERMIONIC k -SAT problem,—perhaps the results in [16] can be useful here—. Such construction would be a quantum counterpart to a classical computer-assisted proof obtained through the use of a classical heuristic SAT solver [17].

As a general question, it might be interesting to consider FERMIONIC ($k > 2$)-SAT problems with additional fermionic symmetry, for example consider whether there are complexity-theoretic consequences of time-reversal, spatial-parity and charge-conjugation symmetries as used in the classification of non-interacting fermionic models [18].

Acknowledgements

We thank Y. Herasymenko and M.H. Shaw for insightful discussions. This work is supported by QuTech NWO funding 2020-2024 – Part I “Fundamental Research”, project number 601.QT.001-1, financed by the Dutch Research Council (NWO).

References

- [1] Bengt Aspvall, Michael F. Plass, and Robert Tarjan. “A linear-time algorithm for testing the truth of certain quantified Boolean formulas”. *Information Processing Letters* **8**, 121–123 (1979).
- [2] Sergey Bravyi. “Efficient algorithm for a quantum analogue of 2-SAT”. In *Contemporary Mathematics*. Volume 536. American Mathematical Society (2011). [arXiv:quant-ph/0602108](#).
- [3] Itai Arad, Miklos Santha, Aarthi Sundaram, and Shengyu Zhang. “Linear-time algorithm for Quantum 2-SAT”. *Theory of Computing* **14**, 1–27 (2018).
- [4] Niel de Beaudrap and Sevag Gharibian. “A Linear Time Algorithm for Quantum 2-SAT”. In *31st Conference on Computational Complexity (CCC 2016)*. *Volume 50 of Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:21. Dagstuhl, Germany (2016). Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [5] David Gosset and Daniel Nagaj. “Quantum 3-SAT is QMA_1 -Complete”. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. *Page 756–765*. FOCS ’13USA (2013). IEEE Computer Society.
- [6] Dorian Rudolph. “Towards a universal gateset for QMA_1 ” (2024). [arXiv:2411.02681](#).
- [7] Sergey B. Bravyi and Alexei Yu. Kitaev. “Fermionic quantum computation”. *Annals of Physics* **298**, 210–226 (2002).
- [8] Fernando de Melo, Piotr Cwiklinski, and Barbara M Terhal. “The power of noisy fermionic quantum computation”. *New Journal of Physics* **15**, 013015 (2013).
- [9] Christina V. Kraus, Michael M. Wolf, J. Ignacio Cirac, and Géza Giedke. “Pairing in fermionic systems: A quantum-information perspective”. *Physical Review A* **79** (2009).
- [10] Jacopo Surace and Luca Tagliacozzo. “Fermionic gaussian states: an introduction to numerical approaches”. *SciPost Physics Lecture Notes* (2022).
- [11] Zhengfeng Ji, Zhaohui Wei, and Bei Zeng. “Complete characterization of the ground-space structure of two-body frustration-free Hamiltonians for qubits”. *Phys. Rev. A* **84**, 042338 (2011).
- [12] Aram W. Harrow. “The church of the symmetric subspace” (2013). [arXiv:1308.6595](#).

- [13] Yaroslav Herasymenko, Anurag Anshu, Barbara M. Terhal, and Jonas Helsen. “Fermionic Hamiltonians without trivial low-energy states”. *Phys. Rev. A* **109**, 052431 (2024).
- [14] Nikolas P Breuckmann and Barbara M Terhal. “Space-time circuit-to-Hamiltonian construction and its applications”. *Journal of Physics A: Mathematical and Theoretical* **47**, 195304 (2014).
- [15] Bryan O’Gorman, Sandy Irani, James Whitfield, and Bill Fefferman. “Intractability of electronic structure in a fixed basis”. *PRX Quantum* **3**, 020322 (2022).
- [16] Robbie King and Tamara Kohler. “Gapped Clique Homology on Weighted Graphs is QMA1-Hard and Contained in QMA”. In 2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS). Pages 493–504. Los Alamitos, CA, USA (2024). IEEE Computer Society.
- [17] Marijn J. H. Heule, Oliver Kullmann, and Victor W. Marek. “Solving and verifying the Boolean Pythagorean triples problem via cube-and-conquer”. In Nadia Creignou and Daniel Le Berre, editors, Theory and Applications of Satisfiability Testing – SAT 2016. Pages 228–245. Cham (2016). Springer International Publishing.
- [18] Alexei Kitaev. “Periodic table for topological insulators and superconductors”. *AIP Conference Proceedings* **1134**, 22–30 (2009).
- [19] Robert Tarjan. “Depth-first search and linear graph algorithms”. *SIAM Journal on Computing* **1**, 146–160 (1972).
- [20] Robert Tarjan. “Edge-disjoint spanning trees and depth-first search”. *Acta Informatica* **6**, 171–185 (1976).
- [21] Tomás Feder. “Network flow and 2-satisfiability”. *Algorithmica* **11**, 291–319 (1994).
- [22] Tomoyuki Morimae, Daniel Nagaj, and Norbert Schuch. “Quantum proofs can be verified using only single-qubit measurements”. *Physical Review A* **93** (2016).

A Mathematical micro-facts for Section 2

A.1 Invariance of Fermionic 2-SAT clauses that exclude all-empty or all-filled states

For an edge $e = (j, k)$ one can define rotated annihilation operators:

$$\begin{pmatrix} \tilde{a}_j \\ \tilde{a}_k \end{pmatrix} = \begin{pmatrix} U_{11} & U_{12} \\ U_{21} & U_{22} \end{pmatrix} \begin{pmatrix} a_j \\ a_k \end{pmatrix}, \quad (24)$$

with U a unitary matrix.

We can define

$$\Pi_e^0 = (I - \tilde{a}_j^\dagger \tilde{a}_j)(I - \tilde{a}_k^\dagger \tilde{a}_k), \quad (25)$$

and

$$\Pi_e^2 = \tilde{a}_j^\dagger \tilde{a}_j \tilde{a}_k^\dagger \tilde{a}_k. \quad (26)$$

The following holds

Proposition A.1. *Projectors Π_e^0 and Π_e^2 in Eqs. (25) and (26) are invariant under any transformation U , as in Eq. (24), and can thus be viewed as classical clauses in the $\{a_j, a_k\}$ mode basis, with the classical Boolean variables corresponding to occupation numbers of the $\{a_j, a_k\}$ modes. That is, Π_e^2 becomes the clause $(\bar{x}_j \vee \bar{x}_k)$ (excluding 11) and Π_e^0 becomes $(x_j \vee x_k)$ (excluding 00) as in Eq. (7).*

Proof. Inserting the transformation U into Π_e^0 with $e = (j, k)$ gives

$$\begin{aligned}\Pi_e^0 &= \tilde{a}_j \tilde{a}_j^\dagger \tilde{a}_k \tilde{a}_k^\dagger \\ &= [I - (U_{11}^* a_j^\dagger + U_{12}^* a_k^\dagger)(U_{11} a_j + U_{12} a_k)] [I - (U_{21}^* a_j^\dagger + U_{22}^* a_k^\dagger)(U_{21} a_j + U_{22} a_k)] \\ &= I - a_j^\dagger a_j - a_k^\dagger a_k + |\det U|^2 a_j^\dagger a_j a_k^\dagger a_k = a_j a_j^\dagger a_k a_k^\dagger,\end{aligned}\quad (27)$$

where we have used unitarity of U : $U_{21} U_{22}^* = -U_{11} U_{12}^*$ and $|\det U| = 1$. Similarly, since $\Pi_e^2 = I - \tilde{a}_j^\dagger \tilde{a}_j - \tilde{a}_k^\dagger \tilde{a}_k + \Pi_e^0$, and $\tilde{a}_j^\dagger \tilde{a}_j + \tilde{a}_k^\dagger \tilde{a}_k$ is invariant under U , Π_e^2 is also invariant under U : $\Pi_e^2 = \tilde{a}_j^\dagger \tilde{a}_j \tilde{a}_k^\dagger \tilde{a}_k = a_j^\dagger a_j a_k^\dagger a_k$. $\square$

A.2 Action of the particle-hole transformation K_S on $\Pi_e^{1,q}$ and $\Pi_e^{02,q}$ clauses

The transformation rules are

$$j \in S, k \notin S: \begin{cases} K_S \Pi_e^{1,q} K_S^{-1} = \tilde{\Pi}_e^{02,q} \text{ with } \tilde{\alpha}_e = - \left[\prod_{\substack{j \leq i < k \\ \text{s.t. } i \in S}} (-1) \right] \beta_e \text{ and } \tilde{\delta}_e = \gamma_e \\ K_S \Pi_e^{02,q} K_S^{-1} = \tilde{\Pi}_e^{1,q} \text{ with } \tilde{\beta}_e = - \left[\prod_{\substack{j \leq i < k \\ \text{s.t. } i \in S}} (-1) \right] \alpha_e \text{ and } \tilde{\gamma}_e = \delta_e. \end{cases} \quad (28)$$

$$j \notin S, k \in S: \begin{cases} K_S \Pi_e^{1,q} K_S^{-1} = \tilde{\Pi}_e^{02,q} \text{ with } \tilde{\alpha}_e = \left[\prod_{\substack{j \leq i < k \\ \text{s.t. } i \in S}} (-1) \right] \gamma_e \text{ and } \tilde{\delta}_e = \beta_e. \\ K_S \Pi_e^{02,q} K_S^{-1} = \tilde{\Pi}_e^{1,q} \text{ with } \tilde{\beta}_e = \left[\prod_{\substack{j \leq i < k \\ \text{s.t. } i \in S}} (-1) \right] \delta_e \text{ and } \tilde{\gamma}_e = \alpha_e. \end{cases} \quad (29)$$

$$j \in S, k \in S: \begin{cases} K_S \Pi_e^{1,q} K_S^{-1} = \tilde{\Pi}_e^{1,q} \text{ with } \tilde{\beta}_e = - \left[\prod_{\substack{j \leq i < k \\ \text{s.t. } i \in S}} (-1) \right] \gamma_e \text{ and } \tilde{\gamma}_e = \beta_e. \\ K_S \Pi_e^{02,q} K_S^{-1} = \tilde{\Pi}_e^{02,q} \text{ with } \tilde{\alpha}_e = - \left[\prod_{\substack{j \leq i < k \\ \text{s.t. } i \in S}} (-1) \right] \delta_e \text{ and } \tilde{\delta}_e = \alpha_e. \end{cases} \quad (30)$$

$$j \notin S, k \notin S: \begin{cases} K_S \Pi_e^{1,q} K_S^{-1} = \tilde{\Pi}_e^{1,q} \text{ with } \tilde{\beta}_e = \left[\prod_{\substack{j \leq i < k \\ \text{s.t. } i \in S}} (-1) \right] \beta_e \text{ and } \tilde{\gamma}_e = \gamma_e. \\ K_S \Pi_e^{02,q} K_S^{-1} = \tilde{\Pi}_e^{02,q} \text{ with } \tilde{\alpha}_e = \left[\prod_{\substack{j \leq i < k \\ \text{s.t. } i \in S}} (-1) \right] \alpha_e \text{ and } \tilde{\delta}_e = \delta_e. \end{cases} \quad (31)$$

B Example of a unique 4-fermion non-Gaussian, non-product satisfying assignment

We work through explicitly how the satisfying assignments depend on fermionic parity for a simple illustrative 4-fermionic problem where one has a line of three $\Pi_e^{1,q}$ clauses on modes 1,2,3 and 4 as in Lemma 3.6, and one adds a single $\Pi_{(2,4)}^{02,q}$ clause between modes 2 and 4. This is an example of a non-hPNC cluster with vertex 2 having degree 3, i.e. an example of Fig. 4(d).

This example gives insight into why product state assignments suffice for QUANTUM 2-SAT (with the exception of cases involving rank-3 projectors) but 4-fermion non-product,

non-Gaussian states are needed for FERMIONIC 2-SAT. On a separate note: from Lemma 3.6 it is clear that if we ask for fixed particle number for QUANTUM 2-SAT, one *can* have a unique satisfying assignment which is not a product state, i.e. a fully permutation-symmetric state of n_q qubits with a fixed number N_q of 1s (with N_q unequal to 0 or the maximum n_q).

The instance we consider here consists of 4 fermionic modes with projectors $\Pi_{(1,2)}^{1,q}$, $\Pi_{(2,3)}^{1,q}$, $\Pi_{(3,4)}^{1,q}$ which we give in qubit language. Specifically, we take $\Pi_{(1,2)}^{\text{qubit}} = |\Psi^-\rangle \langle \Psi^-|_{12}$, $\Pi_{(2,3)}^{\text{qubit}} = |\Psi^-\rangle \langle \Psi^-|_{23}$, $\Pi_{(3,4)}^{\text{qubit}} = |\Psi^-\rangle \langle \Psi^-|_{34}$, with $|\Psi^-\rangle$ denoting a singlet state. For a fixed particle number $0 \leq N \leq 4$, the permutation symmetric state with N excitations is the unique satisfying assignment (i.e., ground state of the ferromagnetic Heisenberg model) via Lemma 3.6. For each N , there is thus a satisfying assignment, i.e.

$$\begin{aligned} |\psi_{N=0}\rangle &= |0000\rangle \\ |\psi_{N=1}\rangle &= |0001\rangle + |0010\rangle + |0100\rangle + |1000\rangle, \\ |\psi_{N=2}\rangle &= |1100\rangle + |0011\rangle + |1010\rangle + |1001\rangle + |0110\rangle + |0101\rangle \\ |\psi_{N=3}\rangle &= |0111\rangle + |1011\rangle + |1101\rangle + |1110\rangle, \\ |\psi_{N=4}\rangle &= |1111\rangle. \end{aligned} \quad (32)$$

Now imagine we wish to add a fourth fermionic $\Pi_{(2,4)}^{02,q}$ -type projector with amplitudes $\alpha_{(2,4)} = \delta_{(2,4)} = 1/\sqrt{2}$ which, in qubit language —note the additional Z_3 due to the Jordan-Wigner transformation in Eq. (2)—, equals

$$\begin{aligned} \Pi_{(2,4)}^{\text{qubit}} &= \frac{1}{2} \left(\sigma_2^- \sigma_2^+ \sigma_4^- \sigma_4^+ + \sigma_2^- Z_3 \sigma_4^- + \sigma_2^+ Z_3 \sigma_4^+ + \sigma_2^+ \sigma_2^- \sigma_4^+ \sigma_4^- \right) \\ &= \frac{1}{2} \left(|00\rangle \langle 00|_{24} + Z_3[|00\rangle \langle 11|_{24} + |11\rangle \langle 00|_{24}] + |11\rangle \langle 11|_{24} \right). \end{aligned} \quad (33)$$

From the expression for $\Pi_{(2,4)}^{\text{qubit}}$, it is clear that an assignment $|\psi\rangle$ can only be a satisfying assignment if it is of the form $|\psi_{\text{even}}\rangle = a_0 |\psi_{N=0}\rangle + a_2 |\psi_{N=2}\rangle + a_4 |\psi_{N=4}\rangle$ (even parity) or $|\psi_{\text{odd}}\rangle = a_1 |\psi_{N=1}\rangle + a_3 |\psi_{N=3}\rangle$ (odd parity). Let us see which (if any) of these states is projected to zero by $\Pi_{(2,4)}^{\text{qubit}}$.

$$\begin{aligned} \Pi_{(2,4)}^{\text{qubit}} |\psi_{\text{even}}\rangle &= \frac{1}{2} \left(a_0(|0000\rangle + |0101\rangle) + a_2(|1010\rangle - |1111\rangle + |0000\rangle + |0101\rangle) \right. \\ &\quad \left. + a_4(-|1010\rangle + |1111\rangle) \right), \\ \Pi_{(2,4)}^{\text{qubit}} |\psi_{\text{odd}}\rangle &= \frac{1}{2} \left(a_1(|0010\rangle - |0111\rangle + |1000\rangle + |1101\rangle) \right. \\ &\quad \left. + a_3(-|0010\rangle + |0111\rangle + |1000\rangle + |1101\rangle) \right), \end{aligned} \quad (34)$$

where the different signs are caused by Z_3 in $\Pi_{(2,4)}^{\text{qubit}}$. Clearly, $\Pi_{(2,4)}^{\text{qubit}} |\psi_{\text{even}}\rangle = 0$ for $a_0 = -a_2$ and $a_2 = a_4$, and $\Pi_{(2,4)}^{\text{qubit}} |\psi_{\text{odd}}\rangle$ cannot be zero. So, $|\psi_{\text{even}}\rangle$ with $a_0 = -a_2$ and $a_2 = a_4$ is the unique satisfying assignment of $\{\Pi_{(1,2)}^{\text{qubit}}, \Pi_{(2,3)}^{\text{qubit}}, \Pi_{(3,4)}^{\text{qubit}}, \Pi_{(2,4)}^{\text{qubit}}\}$, which is clearly not a product state. Let us set $a_0 = 1/2\sqrt{2}$, $a_2 = -1/2\sqrt{2}$ and $a_4 = -1/2\sqrt{2}$ wlog (so that $|\psi_{\text{even}}\rangle$ is normalized).

Back in fermionic language, we thus have a satisfying assignment

$$|\psi_{\text{even},f}\rangle = \frac{1}{2\sqrt{2}} \left(I - (a_1^\dagger a_2^\dagger + a_3^\dagger a_4^\dagger + a_2^\dagger a_3^\dagger + a_1^\dagger a_4^\dagger + a_2^\dagger a_4^\dagger + a_1^\dagger a_3^\dagger) - a_1^\dagger a_2^\dagger a_3^\dagger a_4^\dagger \right) |\text{vac}\rangle. \quad (35)$$

To prove that this is a non-Gaussian state, we argue as follows. A pure Gaussian state has a correlation matrix $M \in \mathbb{R}^{8 \times 8}$ (with entries in Eq. (3)) with orthonormal columns as $M^T M = I$. To show that this does not hold for our state, let's evaluate the first column of M ($M_{j,1}$ for $j = 1, 2, \dots, 8$).

By definition, we have that $M_{1,1} = 0$. Furthermore, $M_{2,1} = 2 \langle \psi_{\text{even},f} | a_1^\dagger a_1 | \psi_{\text{even},f} \rangle - 1$, which can be simply seen to equal 0. For the other entries, let us distinguish between

$$\begin{aligned} M_{2j-1,1} &= i \langle \psi_{\text{even},f} | (a_j + a_j^\dagger)(a_1 + a_1^\dagger) | \psi_{\text{even},f} \rangle, \text{ with odd index } 2j-1, \\ M_{2j,1} &= i \langle \psi_{\text{even},f} | i(a_j - a_j^\dagger)(a_1 + a_1^\dagger) | \psi_{\text{even},f} \rangle, \text{ with even index } 2j, \end{aligned} \quad (36)$$

for $j = 2, 3, 4$. For either type of entry, we have to evaluate the following expectation values, which can be done straightforwardly.

$$\begin{aligned} \langle \psi_{\text{even},f} | a_j a_1^\dagger | \psi_{\text{even},f} \rangle &= \begin{cases} -1/4 & \text{if } j = 2, \\ 0 & \text{if } j = 3, \\ 1/4 & \text{if } j = 4. \end{cases} & \langle \psi_{\text{even},f} | a_j^\dagger a_1 | \psi_{\text{even},f} \rangle &= \begin{cases} 1/4 & \text{if } j = 2, \\ 0 & \text{if } j = 3, \\ -1/4 & \text{if } j = 4. \end{cases} \\ \langle \psi_{\text{even},f} | a_j^\dagger a_1^\dagger | \psi_{\text{even},f} \rangle &= \begin{cases} 0 & \text{if } j = 2, \\ 1/4 & \text{if } j = 3, \\ 0 & \text{if } j = 4. \end{cases} & \langle \psi_{\text{even},f} | a_j a_1 | \psi_{\text{even},f} \rangle &= \begin{cases} 0 & \text{if } j = 2, \\ -1/4 & \text{if } j = 3, \\ 0 & \text{if } j = 4. \end{cases} \end{aligned} \quad (37)$$

Using these expressions, we conclude the following.

$$M_{2j-1,1} = 0 \text{ for } j = 2, 3, 4 \quad \text{and} \quad M_{2j,1} = \begin{cases} 1/2 & \text{if } j = 2, \\ 1/2 & \text{if } j = 3, \\ -1/2 & \text{if } j = 4. \end{cases} \quad (38)$$

Thus the first column of M is $(0, 0, 0, 1/2, 0, 1/2, 0, -1/2)^T$, which has 2-norm $\sqrt{3/4} < 1$, so $|\psi_{\text{even},f}\rangle$ is not Gaussian.

Let's compare our findings briefly with an equivalent set of QUANTUM 2-SAT constraints: Will there be a product assignment here? The QUANTUM 2-SAT equivalent of the FERMIONIC 2-SAT instance considered here is $\{\Pi_{(1,2)}^{\text{qubit}} = |\Psi^-\rangle \langle \Psi^-|_{12}, \Pi_{(2,3)}^{\text{qubit}} = |\Psi^-\rangle \langle \Psi^-|_{23}, \Pi_{(3,4)}^{\text{qubit}} = |\Psi^-\rangle \langle \Psi^-|_{34}, \Pi_{(2,4)}^{\text{qubit}} = \frac{1}{2}(|00\rangle + |11\rangle)(\langle 00| + \langle 11|)_{24}\}$ (note the absence of the Jordan-Wigner Z_3). This QUANTUM 2-SAT instance indeed has a product state satisfying assignment, namely $\bigotimes_{j=1}^4 \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)_j$.

Our pedestrian findings here are captured by the statement for $n_q = 4$ in Corollary 3.5 point 5. The degree-3 vertex in this case is vertex 2, the particle-hole-transformation acts on modes 1, 3, 4 and changes the parity of the non-Gaussian even parity state in Eq. (35) to an odd parity state, while an even parity state in the particle-hole transformed basis (hence odd particle here) is disallowed. The exclusion of one of the two parities for this FERMIONIC 2-SAT instance is essentially what causes the unique satisfying assignment to be non-product and non-Gaussian.

C Classical 2-SAT with fixed parity

In this section, we prove that classical 2-SAT can be solved efficiently, even when constraining the Hamming weight parity of the assignment. This is not necessarily a trivial

problem as they are, for example, simple examples of 2-SAT instances with exponentially many solutions, *all* with the same parity. Consider for instance

$$\bigwedge_{j=0}^{n/4-1} (x_{j+1} \vee \bar{x}_{j+2}) \wedge (x_{j+2} \vee \bar{x}_{j+3}) \wedge (x_{j+3} \vee \bar{x}_{j+4}) \wedge (\bar{x}_{j+1} \vee x_{j+4}), \quad (39)$$

with n a multiple of four. This instance corresponds to disjoint units of four variables, where the only satisfying assignments on each unit are the all-zeros or the all-ones assignment. Clearly, there are exponentially many satisfying assignments, and they are all of even parity. In fact, there are even simple examples of *connected* 2-SAT instances with exponentially many satisfying assignments, all with the same parity. Thus even if one is guaranteed that a 2-SAT instance has (exponentially) many satisfying assignments, the question of whether there is a satisfying assignment with a given parity is a nontrivial one. We prove the following

Theorem C.1. *Given an instance of classical 2-SAT on n variables with m clauses, and a parity $P \in \{-1, +1\}$. Decide whether there exists a satisfying assignment x with Hamming weight parity P (YES), or there is no satisfying assignment with Hamming weight parity P (NO). This problem can be decided in time $O(nm)$.*

Proof. First, we find a satisfying assignment in time $O(n + m)$ [1] for the 2-SAT instance, if it is satisfiable. If this solution has Hamming weight parity P , then we are done (output YES), else we proceed. For convenience, let us redefine the 2-SAT instance such that the obtained solution is the all-zeros assignment. Then what is left is to check whether there exists an odd Hamming weight satisfying assignment for the redefined instance. The redefined instance wlog consists of clauses $(x_i \vee \bar{x}_j)$ (with $i < j$ or $i > j$) and $(\bar{x}_i \vee \bar{x}_j)$ (with $i < j$ wlog). Note that there can be multiple clauses per pair of variables i, j .

Let us consider the sub-graph G_{sub} consisting of just the clauses of type $(x_i \vee \bar{x}_j)$ (with $i < j$ or $i > j$). We associate a directed edge $i \rightarrow j$ with each clause $(x_i \vee \bar{x}_j)$ (i.e., edges point from “variable” i to “negated variable” j). Note that two variables i, j can simultaneously be connected by a $(x_i \vee \bar{x}_j)$ clause and a $(x_j \vee \bar{x}_i)$ clause, resulting in a $i \rightarrow j$ edge and a $j \rightarrow i$ edge in G_{sub} . Next, we identify the strongly-connected components (SCC’s) of the directed graph G_{sub} in time $O(n + m)$ [19]. The only satisfying assignments on these SCC’s are the all-zeros and the all-ones assignments, see Lemma C.2 below. Hence we compress each SCC into a single Boolean variable. We label the new collection of variables by $\{y_j\}_{j=1}^{\leq n}$, where some were previously single variables and others are compressed SCC’s. Since each directed cycle is (part of) an SCC, the compressed problem cannot contain any directed cycles, i.e., it is a directed acyclic graph. The graph with $(y_i \vee \bar{y}_j)$ -type edges can be topologically sorted in time $O(n + m)$ [20]. The topologically-sorted graph now only includes edges of type $(y_i \vee \bar{y}_j)$ for which $i < j$. Next, let us add the $(\bar{x}_i \vee \bar{x}_j)$ -type clauses, which after compression have become either $(\bar{y}_i \vee \bar{y}_j)$ clauses or self-edges $(\bar{y}_i \vee \bar{y}_i)$. The former can be added s.t. $i < j$ in the topologically sorted graph (since, obviously, $(\bar{y}_i \vee \bar{y}_j)$ and $(\bar{y}_j \vee \bar{y}_i)$ are equivalent) to construct the new graph $\tilde{G}$, see Fig. C.2. The self-edges of type $(\bar{y}_i \vee \bar{y}_i)$ exclude $y_i = 1$. These variables with self-edges can be flagged in time $O(n + m)$.

With each variable j in the compressed problem, we associate a weight $w_j \in \{0, 1\}$. The weight w_j corresponds to the number (modulo 2) of original variables in the SCC that has been compressed into variable y_j . Note that for any uncompressed variable, we assign weight 1. To obtain a satisfying assignment of odd Hamming weight (if it exists), we need to flip an odd number of odd-weight variables. In order to do so, we might have to simultaneously flip some even-weight variables.

Importantly, if we flip a variable j (with some weight w_j) to $y_j = 1$, then, by construction, all clauses to variables in *later* layers of the topological sort are still satisfied. Clauses from variable j to variables in *earlier* layers can now become unsatisfied and thus some variables i s.t. $i < j$ might also have to be set to $y_i = 1$. Similarly, flipping such a variable i can only lead to unsatisfied clauses that connect i to variables in earlier layers of the topological sort, not to variables in later layers.

The algorithm runs by first identifying the earliest odd-weight variable j in the topological sort. For some instances, this *might* already be in the first layer. We set $y_j = 1$ and see whether the propagation to earlier layers of the topological sort does not lead to any contradiction. If it does not, then we have found a satisfying assignment of flipped parity (output YES), since all variables that are flipped in earlier layers due to propagation are even-weight variables by definition. If the propagation leads to contradiction, then clearly there is *no* satisfying assignment for which $y_j = 1$. Next, we reset to the all-zeros assignment and identify the next odd-weight variable k in the topological sort (which is in the same or in a later layer). Again, one checks whether setting $y_k = 1$ leads to contradiction in earlier layers. If it does not, then we have constructed a satisfying assignment with flipped parity (output YES), since y_k is the *only* odd-weight variable set to 1 in that assignment. Indeed, if y_j would also have to be set to 1, then there would be a contradiction. If setting $y_k = 1$ does lead to contradiction, then there is *no* satisfying assignment for which $y_k = 1$. We proceed again by resetting to the all-zeros assignment and identifying the next odd-weight variable in the topological sort. We iterate over all odd-weight variables in this manner and either obtain a satisfying assignment of flipped parity or conclude that no odd-weight variable can be set to 1 consistently and hence there is no satisfying assignment with flipped parity (output NO). Iterating over all odd-weight variables and checking for the consistency of flipping them takes time $O(nm)$. $\square$




Figure 6: First three layers of a compressed topologically sorted graph $\tilde{G}$ derived from G , where each variable j has been given a weight w_j . The red clauses are of type $(y_i \vee y_j)$ (with $i < j$) and the blue clauses are of type $(\bar{y}_i \vee \bar{y}_j)$.

Lemma C.2. *Given an instance of classical 2-SAT with $(x_i \vee \bar{x}_j)$ -type clauses only. Let us consider the directed graph with edges $i \rightarrow j$ for each clause $(x_i \vee \bar{x}_j)$ (i.e., edges pointing from “variable” i to “negated variable” j). If this graph is strongly connected, then only the all-zeros and all-ones assignments are satisfying assignments.*

Proof. If the directed graph is strongly connected, then for *each* pair of variables (i, j) there is path from i to j and a path from j to i . Let us set $x_i = 0$. Then clauses along the path of directed edges from variable i to variable j can only be satisfied if $x_j = 0$. Conversely, let us set $x_i = 1$ and let us now travel “upstream” to variable j . The clauses

along the path of directed edges from variable j to variable i can only be satisfied if $x_j = 1$. Hence the lemma statement follows. $\square$

Remark. *The algorithm in Theorem C.1 cannot be used to solve the NP-complete problem of finding an assignment with fixed Hamming weight (particle number) since one can generally not appropriately change the particle number by flipping a polynomial set of designated variables. Similarly, the algorithm above cannot be used to count the number of solutions to a 2-SAT instance (which is a #P-complete problem), although there is a way of enumerating all solutions with effort growing with the number of solutions, which uses similar techniques as our algorithm [21].*

D Proof of Lemma 6.1

Proof. Containment in QMA can be obtained by simply mapping the n -mode fermionic problem onto a n -qubit problem using e.g. the Jordan-Wigner transformation in Eq. (2). This means that each projector Π_i using k fermionic modes is mapped to a projector which is a sum of terms, each of which has a k -local part to which a string of Pauli Z s of some length is appended. Hence the FERMIONIC k -SAT problem maps onto a local Hamiltonian problem but some Pauli strings in the Hamiltonian have weight larger than k . However, one can still apply standard QMA-arguments for the proof verification as in Section IV in [22] since this method only relies on the fact that the number of Pauli terms in the Hamiltonian is $\text{poly}(n)$.

To prove hardness, we can map the QMA₁-complete problem in [5] onto a FERMIONIC $k = 9$ -SAT problem using the unitary “assimilation mapping” described in Lemma 15 in [13]. We take the 3-local circuit Hamiltonian, with its 3-local projectors, on some n qubits in the QMA₁-proof and view it as a Hamiltonian on n qubits plus $n/2$ fermionic modes which gets mapped on a space of $3n/2$ fermionic modes and the Pauli operators of one qubit get mapped on products of two (Majorana) fermion operators out of three fermionic modes. Thus each 3-local projector becomes a fermionic projector Π_i involving at most 9 fermionic modes which conserves fermionic parity. In the YES case, if there is a n -qubit satisfying assignment for the QUANTUM 3-SAT problem, it can be rewritten as a $3n/2$ -fermionic state which is a satisfying assignment for the FERMIONIC 9-SAT problem. In the NO case, if there is no approximate satisfying assignment for QUANTUM 3-SAT, imagine that there is a fermionic state $|\Phi\rangle$ for the FERMIONIC 9-SAT problem for which $\sum_i \langle \Phi | \Pi_i | \Phi \rangle = \beta$. Then performing the inverse mapping and tracing over the additional fermionic registers in $|\Phi\rangle \langle \Phi|$ to get the qubit reduced density matrix ρ^{qubit} , this implies that $\sum_i \text{Tr} \Pi_i^{\text{qubit}} \rho^{\text{qubit}} = \beta$ which by assumption implies that $\beta \geq 1/\text{poly}(n)$. $\square$