

Complexity of Supersymmetric Systems and the (co)Homology Problem

Chris Cade¹ and P. Marcos Crichigno²

¹QuSoft & CWI, Amsterdam, the Netherlands.

²Blackett Laboratory, Imperial College Prince Consort Rd., London, SW7 2AZ, U.K.

We consider the complexity of the local Hamiltonian problem in the context of fermionic Hamiltonians with $\mathcal{N} = 2$ supersymmetry and show that the problem remains QMA-complete. Our main motivation for studying this is the well-known fact that the ground state energy of a supersymmetric system is exactly zero if and only if a certain cohomology group is nontrivial. This opens the door to bringing the tools of Hamiltonian complexity to study the computational complexity of a large number of algorithmic problems that arise in homological algebra, including problems in algebraic topology, algebraic geometry, and group theory. We take the first steps in this direction by introducing the k -LOCAL COHOMOLOGY problem and showing that it is QMA₁-hard and, for a large class of instances, is contained in QMA. We then consider the complexity of estimating normalized Betti numbers and show that this problem is hard for the quantum complexity class DQC1, and for a large class of instances is contained in BQP. In light of these results, we argue that it is natural to frame many of these homological problems in terms of finding ground states of supersymmetric fermionic systems. As an illustration of this perspective we discuss in some detail the model of Fendley, Schoutens, and de Boer consisting of hard-core fermions on a graph, whose ground state structure encodes l -dimensional holes in the independence complex of the graph. This offers a new perspective on existing quantum algorithms for topological data analysis and suggests new ones.

Contents

1	Introduction	2
1.1	Summary of results	6
2	Supersymmetric Quantum Mechanics	8
2.1	Definition, algebra, and spectrum	9
2.2	Supersymmetric ground states and cohomology	11
2.3	Supersymmetry in lattice systems	14
3	The SUSY Local Hamiltonian Problem	16
3.1	Definitions	16
3.2	The SUSY local Hamiltonian problem is QMA-complete	17
4	The k-local Cohomology Problem	20
4.1	The k -local cohomology problem	24
4.2	The complexity of the k -local cohomology problem	27
5	The Betti Number Problem	30
5.1	QUASI-BETTI NUMBER ESTIMATION is DQC1-hard	35
5.2	QUASI-BETTI NUMBER ESTIMATION is in BQP	37
6	Betti Numbers in Topological Data Analysis	39
6.1	TDA and the fermionic hard-core model	40
6.2	Existing quantum algorithms for TDA	42
6.3	Alternative quantum and quantum-inspired approaches to TDA	43
7	Discussion and Outlook	44
A	Complexity Classes	47
B	DQC1 Algorithm for Estimating Quasi Betti Numbers	48

1 Introduction

The development of the theory of quantum computation has greatly enriched our understanding of computational complexity. Contemplating the ways in which quantum mechanics can be harnessed to process information has not only led to the discovery of new (quantum) complexity classes, but also to novel tools for establishing properties of known complexity classes, and revealed that certain computational problems which (most likely) cannot be efficiently solved by classical computers can be efficiently solved by quantum computers [80].¹ Determining the complexity of a given computational problem, however, can be a difficult task and the techniques for doing so typically depend on the sort of problem being considered, e.g., whether it is a problem in linear algebra, topology, or group theory. Some problems are manifestly quantum mechanical (e.g., finding the ground state of a quantum Hamiltonian [54, 52, 53]) but others may conceal their quantum mechanical nature. A “prime” example is that of factoring which, although at face value has little to do with quantum mechanics, was famously shown by Shor [80] to be efficiently solvable by a quantum computer, while it is believed that it cannot be solved efficiently by a classical computer.

In this paper, which is a follow-up to [17], we point out that there is a large class of computational problems which at first may not seem quantum mechanical in nature, but in fact they are. These problems fall under the mathematical field of “homological algebra” [85, 15], in particular the “(co)homology problem” described below. Although this can be a rather abstract field, the advantage is that it provides a unified framework encompassing a large number of disparate mathematical problems, including problems of practical interest (see, e.g., [50]).

Homological algebra has its origins in the study of topology, in particular in the problem of describing “ l -dimensional holes” in simplicial complexes, but exploded in its applications and reach when it was realized that the formalism can be applied to algebraic systems (see [86] for a history of the subject). Homological algebra has now become a powerful tool used in several branches of mathematics, including algebraic topology, algebraic geometry, and group theory, as well as in physics, including condensed matter theory, quantum gauge theories, and string theory.

The cohomology problem is described as follows. One is given a vector space, V , a “coboundary” operator $d : V \rightarrow V$ squaring to zero, $d^2 = 0$, and is asked to determine the elements $v \in V$ such that $dv = 0$, modulo elements of the form $v = du$ for some u . Such elements are, by definition, in the “cohomology” of d .² This includes problems as diverse as finding l -dimensional holes in a simplicial complex (via simplicial cohomology), studying fixed points under the action of a group (via group cohomology), and problems in number theory (via Galois cohomology).

One of our main goals here is to bring the tools of Hamiltonian complexity to bear on the complexity

¹See [84] for a review on quantum complexity theory and [20] for a survey on the application of quantum techniques to classical problems.

²One can similarly define the “homology” problem, defined by a “boundary” operator ∂ . These are closely related, with the cohomology problem being the “dual” of the homology problem. We will often refer only to the cohomology problem. The precise formulation of these problems is given in Section 4.




Figure 1: The intersection of supersymmetry, many-body physics, and computational complexity theory considered in this paper. A set of problems in homological algebra arises at this intersection, whose complexity is naturally captured by various quantum complexity classes. This includes, for instance, the problem of finding l -dimensional holes in a simplicial complex.

of problems in homological algebra, in particular the (co)homology problem just described. The field of Hamiltonian complexity lies at the intersection of many-body physics and computational complexity. Starting with Kitaev’s celebrated result that determining the ground state energy of a quantum local Hamiltonian is QMA-complete [54, 52, 53], it has led to a number of important insights and tools in quantum complexity theory.

The bridge between Hamiltonian complexity and the cohomology problem is provided by the study of Hamiltonians describing *supersymmetric systems*, a class of quantum mechanical systems which are invariant under the exchange of fermionic and bosonic states [66, 87–89]. By definition, the Hamiltonian for these quantum mechanical systems can be written as $H = \{Q, Q^\dagger\}$, where Q is known as the “supercharge” of the system, an operator sending bosonic states into fermionic states (and vice versa) and squaring to zero, $Q^2 = 0$.³ An important role is played by states with exactly zero energy, $E = 0$, called *supersymmetric ground states*. As noted by Witten [88, 89] these organize themselves into certain cohomology groups and thus the question of whether the cohomology group is nontrivial is translated into the physical question of whether the spectrum contains any states with $E = 0$. In this way, the complexity of determining cohomology groups can be studied using the tools of Hamiltonian complexity, applied to supersymmetric systems.

A new set of questions thus arises at the intersection of many-body physics, the theory of computational complexity, and supersymmetry (see Fig. 1) in what may be called “quantum computational (co)homology.”

To make this concrete, let us illustrate the interplay at this intersection with the beautiful fermion

³Technically, this corresponds to an $\mathcal{N} = 2$ model; see Section 2 for details. The symbol $\{\cdot, \cdot\}$ denotes the anticommutator: $\{a, b\} := ab + ba$.

hard-core Hamiltonian of Fendley, Schoutens, and de Boer [26]. The Hamiltonian describes a set of fermionic modes on the vertices i of a simple graph, $G = (V, E)$, created and annihilated by standard fermionic operators $a_i^\dagger$ and a_i , respectively.⁴ A fermion is allowed to occupy an empty vertex provided none of its neighbors is occupied, hence the term “hard-core.” The Hamiltonian is given by

$$H = \sum_{(i,j) \in E} P_i a_i^\dagger a_j P_j + \sum_{i \in V} P_i, \quad (1.1)$$

where $P_i = \prod_{j|(i,j) \in E} (1 - a_j^\dagger a_j)$ is a projector which vanishes if any vertex adjacent to i is occupied. The first term in (1.1) is a hopping term among adjacent sites (as long as the hard-core condition is satisfied). The second term is required for supersymmetry; the supercharge is given by $Q = \sum_i a_i^\dagger P_i$ as can be easily checked. Given a graph G one can define an associated simplicial complex, known as the independence complex $I(G)$ (see Section 5). It turns out that supersymmetric ground states $|\Omega_{l+1}\rangle$ of the Hamiltonian (1.1) with fermion number $F = l + 1$ (i.e., with $l + 1$ vertices occupied by a fermionic mode), are in one-to-one correspondence with l -dimensional holes in the corresponding independence complex:

$$\begin{array}{ccc} l\text{-dimensional holes} & \iff & H |\Omega_{l+1}\rangle = 0. \\ \text{in } I(G) & & \end{array} \quad (1.2)$$

We will return to this model in Section 6. (See [46] for a nice review of this model, its relation to the independence complex, and references therein for related work in the math literature.)

Finding l -dimensional holes in the independence complex has applications in topological data analysis (TDA) [83] and a quantum algorithm for this was presented in [61].⁵ An alternative approach to TDA is thus provided by implementing the fermionic mode Hamiltonian (1.1) – gates simulating hopping terms in fermionic models have already been implemented on near-term devices [27]. We comment on this further in Section 6, where we also propose an alternative, variational quantum algorithm for TDA.

A number of supersymmetric models have been studied in the many-body literature—see [66, 76, 25, 91, 24, 19, 33] for a non-exhaustive list. Supersymmetric systems can arise at special points in parameter space of systems that are well studied in the computational complexity literature, such as an open XXZ spin chain and adding a boundary magnetic field with a precise value or the t - J model at the point $J = -2t$ [91]. Finding ground states of close cousins of these systems remains QMA-hard (see, e.g., [18, 67]). It is thus conceivable that such simple supersymmetric systems remain QMA-hard and it would be very interesting to establish this.

The goal of this paper, however, is not to focus on particular supersymmetric Hamiltonians but rather to study the computational complexity associated to a large class of supersymmetric Hamiltonians and thus of a large class of cohomology problems. We view this as a first step towards a program of

⁴The operators $a_i^\dagger$ and a_i create and destroy a fermionic mode (or simply a fermion) at site i , respectively. These satisfy the standard anticommutation relations—see Section 2.3 for details or, e.g., [11].

⁵More precisely, the problem considered there relates to the clique complex, but this is equivalent to the independence complex of the complement graph $\bar{G}$.

determining the computational complexity of a large variety of specific homological problems using the techniques of Hamiltonian complexity, thus providing an overarching framework for studying the complexity of a variety of disparate computational problems.

As we discuss in more detail below, the perspective just described leads to *promise* versions of the (co)homology problem. The notion of promise problems – decision problems in which the input comes with a certain promise [36] – is central to quantum computation. Indeed, in the black box setting, without such promises quantum algorithms can at most provide polynomial speedups over classical algorithms [7], and in fact no complete problems are known for the class BQP, only for the class “PromiseBQP” (see [47] for a detailed discussion). It is thus essential when considering problems in homological algebra (or any other field) from a quantum computational perspective to identify appropriate promises. This is precisely what a (supersymmetric) Hamiltonian complexity perspective accomplishes for homological algebra, translating the promises natural for Hamiltonians, such as properties of the ground state or spectral gap, into promises on the homological problems, allowing us to precisely pinpoint their computational complexity. One may hope that this perspective can lead to the identification of new problems with provable quantum speed-ups (assuming widely believed conjectures in complexity theory). Indeed, as we discuss in detail below, the problem of estimating the rank of general cohomology groups (or Betti numbers) is one such problem.⁶ It would be interesting to study whether this remains the case for particular cohomology problems, such as finding the number of l -dimensional holes in a simplicial complex. Although we do not establish the hardness of these particular problems here, the quantum Hamiltonian perspective we advocate provides a concrete setting in which to study this and other similar problems.

As illustrated above, a subclass of supersymmetric Hamiltonians connect to problems in algebraic topology, thus establishing a close link between certain problems in topology and quantum computation. A well known relation between topology and quantum computation also arises in the context of topological quantum computation, a model of quantum computation based on the manipulation of anyons [28, 29] (see [72, 56] for reviews). This has also led to a number of complexity results in computational topology, in particular in relation to the computation of the Jones polynomial [55, 30, 31, 5]. The connection here is via the study of topological quantum field theory (TQFT), specifically Chern-Simons theory, and its relation to the Jones polynomial as uncovered by Witten [90]. The appearance of topology in supersymmetric quantum mechanics is less explicit, being related to supersymmetric ground states or to a particular class of “supersymmetric” operators. See [17] for a discussion of these operators in the context of quantum computation.

The high degree of symmetry of supersymmetric systems often leads to a strong computational control, sometimes even rendering them exactly solvable (see, e.g., [25, 41]). What makes supersymmetric systems most interesting from our perspective, however, is not that in some cases they may become

⁶More precisely, we show this is the case for the closely related problem of computing “quasi” Betti numbers, which we define in Section 5.

“simple” from the perspective of a mathematical physicist but the fact that they remain computationally hard, implying the hardness the corresponding problems in homological algebra.

Although the main focus of this paper is on the problem of computing (co)homology groups, and estimating their ranks, we expect this interplay to extend beyond these problems. Similarly, we note that the field of Hamiltonian complexity has led to a number of interesting complexity results (see [34] for a review) and we conjecture that many of these could be “ported” over to interesting statements about homological algebra.

1.1 Summary of results

We first show that the local Hamiltonian problem remains QMA -complete for supersymmetric Hamiltonians. This is shown in Section 3, by reduction from the standard (non-supersymmetric) local Hamiltonian problem. This result becomes relevant to the study of (co)homology when considering the perfect completeness, or QMA_1 , version of the problem, which we study in Section 4. This is the problem of deciding whether the supersymmetric Hamiltonian has any supersymmetric ground states, i.e., states with *exactly* zero energy. Since supersymmetric ground states are in one-to-one correspondence with cohomology groups it thus follows that the problem of deciding whether a cohomology group is trivial or nontrivial is QMA_1 -hard. As we discuss, the problem is generally contained in QMA and thus its complexity lies somewhere between QMA_1 and QMA .

We then move on to the problem of determining the rank of the cohomology group, or Betti number. This is equivalent to counting supersymmetric ground states which is $\#\text{BQP}$ -complete, as we discuss in Section 5.⁷ We then consider the problem of approximating Betti numbers up to some additive accuracy. As we discuss in detail, it is less clear how to precisely pinpoint the complexity of this problem, and following [40] we thus consider a more general problem, which we call “quasi Betti number estimation” and which coincides with the problem of estimating Betti numbers when the input satisfies a certain promise.⁸ Using the tools from the previous section we show that the general quasi Betti number estimation problem is DQC1 -hard, and that there is a BQP algorithm for this problem and thus its hardness lies somewhere between DQC1 and BQP . This shows that there exists an efficient quantum algorithm for this problem that cannot be dequantized, unless $\text{DQC1} \subseteq \text{BPP}$. We stress that the quasi Betti number estimation problem is identical to the Betti number estimation problem only when an additional promise is given on the system, in the form of an inverse polynomial gap in the spectrum. Unfortunately, it is not clear at the moment whether the quasi Betti number problem remains DQC1 -hard in this case and thus the complexity of the Betti number problem remains open.⁹ Studying this in

⁷ $\#\text{BQP}$ (the counting version of QMA) is equal to $\#\text{P}$ under weakly parsimonious reductions [12].

⁸Although we define the problem in terms of estimating Betti numbers, this can easily be framed as a decision problem from which the estimate can be recovered via binary search, and so it makes sense to talk about this problem in the context of complexity classes of decision problems such as BQP , DQC1 , etc.

⁹It is in fact possible to show that the true Betti number problem (counting only exactly-0 eigenvalues) is hard for a “perfect-completeness version” of DQC1 using techniques from [9] and an appropriate definition of that class. However, it

detail may be of great interest, in particular for the independence (or clique) complex mentioned above in light of possible implementations of a quantum algorithm for this problem in near-term devices. We comment further about this in Section 6.

Relation to previous work. The current paper expands on the ideas of [17], which considered various aspects at the intersection of supersymmetry and quantum computation. The problem of counting (with signs) the number of supersymmetric ground states was shown to be $\#P$ -complete for a 6-local supercharge. The problem of determining the ground state energy of supersymmetric systems was posed, suggesting that the problem may remain QMA-hard as we show here. In an independent set of developments, the question of the complexity of estimating normalized Betti numbers of the independence complex of a graph was considered by Gyurik et al. in [40]. A relaxed version of the problem, “approximate Betti number estimation” (ABNE), was defined but its hardness was left as an open problem. Instead, the authors showed that the more general problem of determining the “low-lying spectral density” (LLSD) of a Hamiltonian is DQC1-hard. The relevance of this result to the problem of estimating normalized “quasi” Betti numbers, however, was left open. This question is addressed here, in the context of our other main results, where we show that estimating normalized “quasi” Betti numbers for general cohomology groups is DQC1-hard. The question on whether this remains the case for the clique complex, however, remains open.

Historical comment. Supersymmetry was originally proposed in the late 1960s in the context of string theory and quantum field theory as a possible fundamental symmetry of spacetime, with the consequence that for every fundamental bosonic particle there should be an accompanying fundamental fermionic particle. Since then, research in supersymmetry has grown into playing a prominent role in modern theoretical physics: it is a powerful tool in charting the space of possible quantum field theories and deriving exact, non-perturbative, results; it is a cornerstone of string/M-theory and the study of quantum gravity via the holographic principle; and has led to various developments in pure mathematics including complex geometry, topology, and representation theory. As mentioned above, most important to us is the intimate relation between supersymmetry and cohomology. Returning to its origins, the question of whether supersymmetry is realized as a fundamental symmetry at the level of particle physics has been an active area of theoretical and experimental research for several decades but still remains open. The class of supersymmetric systems that we consider here, however, are not affected by the answer to this question as supersymmetry appears in these systems not as a fundamental symmetry but as a symmetry of certain many-body systems which can, in principle, be engineered in the lab or as an emergent symmetry of certain condensed-matter systems at low energies.

Experimental realizations. The systems we consider can, at least in principle, be realized in the lab as spin or qubit systems with specially tuned interactions to make them supersymmetric. An

is not clear that this result is very meaningful – a perfect completeness version of DQC1 is not a particularly well-defined model, and it is unclear whether the class would be hard to simulate classically, as we suspect is the case for DQC1.

example is given by the XXZ model with a specially tuned boundary magnetic field [91]. See [64] and references therein for proposed realizations of the fermion hard-core model in 1D using trapped atoms. Other experimentally realizable systems include relativistic supersymmetric quantum systems emerging at low energies at the boundary of a topological phase [38], or from strongly interacting Majorana zero modes [73]. Thus, the complexity results described here and in [17] can be seen as a first step towards characterizing the computational capabilities of these experimentally realizable systems. As pointed out in [17], the topological nature of certain subsectors of supersymmetric quantum mechanics suggests a natural robustness of these systems against a certain type of noise. It would be interesting to study this further but leave this and others interesting questions for future work.

Organization of the paper. The paper aims at being accessible to both physicists and computer scientists and is as self-contained as possible. To make it accessible to computer scientists we begin in Section 2 with a review of supersymmetric quantum mechanics as well as a basic introduction to cohomology, focusing on the relation to supersymmetric quantum mechanics and the study of supersymmetric ground states. To make the paper accessible to physicists, we define the basic concepts of quantum complexity theory, in particular the definition of various quantum complexity classes, as they are needed. We refer to [84] for a survey of quantum complexity theory, to [34] for a survey of Hamiltonian complexity, and to [49, 43] for a broader review of supersymmetry and its myriad applications in statistical physics and mathematics. In Section 3 we overview basic definitions in quantum complexity theory, define the local Hamiltonian problem for supersymmetric systems, and show that it remains QMA-hard. The tools developed there are the basis for the problems of most interest to us, namely that of determining cohomology groups, which we study in Section 4, and that of estimating Betti numbers which we study in Section 5. We note that although we draw much of our intuition from the SUSY local Hamiltonian problem, the k -local cohomology problem can be formulated without any reference to supersymmetry and thus for readers who are mostly interested in computational (co)homology it is possible to skip directly to Sections 4 and 5, which we have made as self-contained as possible, with little reference to previous sections. In Section 6 we focus our attention on the fermion hard-core model and its relation to the topic of topological data analysis (TDA), suggesting some approaches towards TDA on near-term devices. Finally, we conclude with a discussion of open problems in Section 7.

2 Supersymmetric Quantum Mechanics

We begin by reviewing basic aspects of supersymmetric (SUSY) quantum mechanics [66, 87–89], including the supersymmetric algebra, the structure of ground states and the role of cohomology. We mostly follow the original exposition in [88, 89] and also Chapter 10 of [43].¹⁰

¹⁰See also [49] for a textbook focusing on applications of supersymmetry to statistical mechanics.

2.1 Definition, algebra, and spectrum

The Hilbert space of any quantum mechanical theory can be decomposed as

$$\mathcal{H} = \mathcal{H}^B \oplus \mathcal{H}^F, \quad (2.1)$$

where $\mathcal{H}^B$ and $\mathcal{H}^F$ are the spaces of bosonic and fermionic states, respectively. These are distinguished by an operator, denoted $(-1)^F$, which acts as $+1$ on bosonic states and as -1 on fermionic states. By definition, a supersymmetric quantum mechanics is a quantum mechanical system in which there is a set of Hermitian operators, Q_I , $I = 1, \dots, \mathcal{N}$, sending states in $\mathcal{H}^B$ into states in $\mathcal{H}^F$ and vice versa. The operators Q_I are the generators of supersymmetry, or “supercharges,” and one says that the quantum mechanics has $\mathcal{N}$ real supercharges. A basic consistency requirement is that

$$\{(-1)^F, Q_I\} = 0, \quad (2.2)$$

where $\{\cdot, \cdot\}$ is the anticommutator. In addition, one imposes that the supercharges satisfy the algebra

$$\{Q_I, Q_J\} = \delta_{IJ} H, \quad I, J \in \{1, \dots, \mathcal{N}\}, \quad (2.3)$$

where H is, by definition, the Hamiltonian of the system. In words, an $\mathcal{N} = 1$ supersymmetric Hamiltonian is thus a Hamiltonian that can be written as the *square* of a parity-odd, Hermitian, operator Q . If there are $\mathcal{N}$ such independent operators, all anticommuting with each other, then one says that the systems “has $\mathcal{N}$ real supercharges.”

It follows directly from the algebra above that

$$[Q_I, H] = 0, \quad [(-1)^F, H] = 0, \quad (2.4)$$

and thus the operators Q_I and $(-1)^F$ are symmetries of such systems. Another crucial property of supersymmetric systems is that the spectrum is positive semi-definite:

$$\langle \psi | H | \psi \rangle = \langle \psi | \{Q_I, Q_I\} | \psi \rangle = 2 \|Q_I | \psi \rangle\|^2 \geq 0, \quad \text{for each } I. \quad (2.5)$$

In particular, a state $|\Omega\rangle$ has $E = 0$ if and only if it preserves all the supersymmetries, i.e.,

$$Q_I |\Omega\rangle = 0, \quad I = 1, \dots, \mathcal{N}. \quad (2.6)$$

Such states, if they exist, are called *supersymmetric ground states* and will play a fundamental role in what follows.

A crucial property of states with $E > 0$ is that they come in boson-fermion pairs with the same energy. To see why this is the case consider for example a bosonic state $|b\rangle$ with $E > 0$ and thus there exists at least one supercharge, say $Q := Q_1$, which does not annihilate the state. Then, the fermionic state $|f\rangle := \frac{1}{\sqrt{E}} Q |b\rangle$ is normalized and has the same energy as $|b\rangle$, since Q commutes with the Hamiltonian. One then says that the states $(|b\rangle, |f\rangle)$ form a “supersymmetric doublet” and furnish a two-dimensional representation of supersymmetry. When there are several supercharges one may apply




Figure 2: The spectrum of supersymmetric quantum mechanics. Black dots represent bosonic states and white ones fermionic states. All states with $E > 0$ come in boson/fermion pairs. Supersymmetric ground states ($E = 0$), if there are any, are not necessarily paired.

multiple Q_I 's, obtaining a longer multiplet of bosonic and fermionic states but these can be constructed by combining doublets for each supercharge. The conclusion is that, regardless of the total number of supercharges, for each bosonic state with $E > 0$ there is a corresponding fermionic state with the same energy (see Figure 2).¹¹

The situation for supersymmetric ground states, $E = 0$, on the other hand, is very different. As seen above, in this case all supercharges annihilate the state, $Q_I |\Omega\rangle = 0$, and thus supersymmetric ground states cannot be paired into doublets. Rather, they stand by themselves, in the trivial representation of supersymmetry. Thus, in general there can be an arbitrary number, $n_{E=0}^B$ and $n_{E=0}^F$, of bosonic and fermionic supersymmetric ground states, respectively.

Given a supersymmetric system, an interesting question is whether the vacuum state $|\Omega\rangle$ preserves the full supersymmetry of the system, i.e., $Q_I |\Omega\rangle = 0$ for all I . As shown above, this is the case only if the ground state has $E = 0$. When this is not the case, one says that the system breaks supersymmetry “spontaneously.” This is in general a rather difficult question to answer (in fact, it is QMA_1 -hard as shown below) as this typically requires exact, non-perturbative, methods to determine if the energy of a state is *exactly* zero. This led Witten to define the “index” [88],

$$\mathcal{I} := n_{E=0}^B - n_{E=0}^F, \quad (2.7)$$

given by the difference in the number of bosonic and fermionic supersymmetric ground states. The quantum mechanical system of Figure 2, for instance, has Witten index $\mathcal{I} = 4 - 2 = 2$. Under certain conditions, the “Witten index” (2.7) can be computed exactly by perturbative methods (see [88] for a detailed discussion). This simplicity, however, comes at the expense of only providing partial information on supersymmetry breaking; if it is nonzero the cohomology is necessarily nontrivial but if it vanishes it is inconclusive (it only implies that $n_{E=0}^B = n_{E=0}^F$ but these could vanish or not). Despite its relative simplicity, computing the Witten index is $\#\text{P}$ -hard [17].

¹¹In this paper we shall be mostly interested in finite-dimensional Hilbert spaces and thus the spectrum is discrete.

It is useful to distinguish operators acting on the Hilbert space according to the action of $(-1)^F$ on the operator. An operator is said to be bosonic or fermionic if it commutes or anticommutes with the parity operator, i.e.,

$$[(-1)^F, \mathcal{O}_B] = 0, \quad \{(-1)^F, \mathcal{O}_F\} = 0, \quad (2.8)$$

respectively. It follows that when acting on a state with definite parity, bosonic operators preserve its parity and fermionic operators flip it. The supercharges Q_I , for example, are fermionic operators and the Hamiltonian is a bosonic operator.

$\mathcal{N} = 2$ systems. From now on we will focus on systems with $\mathcal{N} = 2$ supersymmetry. In this case it is useful to combine the Hermitian supercharges into the complex combinations

$$\mathcal{Q} := \frac{1}{\sqrt{2}}(Q_1 + iQ_2), \quad \mathcal{Q}^\dagger := \frac{1}{\sqrt{2}}(Q_1 - iQ_2). \quad (2.9)$$

The supersymmetry algebra (2.3) implies these complex supercharges are nilpotent of degree two, i.e.,

$$\mathcal{Q}^2 = (\mathcal{Q}^\dagger)^2 = 0. \quad (2.10)$$

The Hamiltonian can be written as

$$H = \{\mathcal{Q}, \mathcal{Q}^\dagger\}. \quad (2.11)$$

As discussed above, supersymmetric ground states (states with exactly zero energy) must preserve all supersymmetries and thus

$$\mathcal{Q}|\Omega\rangle = \mathcal{Q}^\dagger|\Omega\rangle = 0. \quad (2.12)$$

One of the most important properties of $\mathcal{N} = 2$ supersymmetric quantum mechanics is that supersymmetric ground states are in one-to-one correspondence with the cohomology group of $\mathcal{Q}$. We review this next.

2.2 Supersymmetric ground states and cohomology

A special role is played by states that are “closed” or “exact” with respect to the supercharge. A state is said to be “ $\mathcal{Q}$ -closed” if it is in the kernel of $\mathcal{Q}$, i.e.,

$$\mathcal{Q}|\psi\rangle = 0. \quad (2.13)$$

A state is said to be “ $\mathcal{Q}$ -exact” if it is in the image of $\mathcal{Q}$, i.e., if there exists some state $|\psi'\rangle$ such that

$$|\psi\rangle = \mathcal{Q}|\psi'\rangle. \quad (2.14)$$

Consider the set of states in the kernel and image of $\mathcal{Q}$, i.e.,

$$\text{Ker}(\mathcal{Q}) = \{|\psi\rangle \mid \mathcal{Q}|\psi\rangle = 0\}, \quad \text{Im}(\mathcal{Q}) = \{|\psi\rangle \mid |\psi\rangle = \mathcal{Q}|\psi'\rangle\}. \quad (2.15)$$

The same definitions hold for $\mathcal{Q}^\dagger$ and supersymmetric ground states (2.12) are closed with respect to both $\mathcal{Q}$ and $\mathcal{Q}^\dagger$. Now, since $\mathcal{Q}^2 = 0$, all states that are $\mathcal{Q}$ -exact are also $\mathcal{Q}$ -closed and thus $\text{Im}(\mathcal{Q}) \subseteq \text{Ker}(\mathcal{Q})$. Then, the vector space

$$H(\mathcal{Q}) := \frac{\text{Im}(\mathcal{Q})}{\text{Ker}(\mathcal{Q})} \quad (2.16)$$

is well defined and known as the *cohomology group* of $\mathcal{Q}$.¹² It consists of all states in the Hilbert space which are $\mathcal{Q}$ -closed but not $\mathcal{Q}$ -exact. In fact, due to the $\mathbb{Z}_2$ -grading of the Hilbert space, $\mathcal{H} = \mathcal{H}_B \oplus \mathcal{H}_F$, there is a more refined structure and one can define a $\mathbb{Z}_2$ -graded “complex” of vector spaces,

$$C : \quad \mathcal{H}^F \xrightarrow{\mathcal{Q}} \mathcal{H}^B \xrightarrow{\mathcal{Q}} \mathcal{H}^F \xrightarrow{\mathcal{Q}} \mathcal{H}^B. \quad (2.17)$$

Then, one can define two cohomology groups:

$$H^B(\mathcal{Q}) := \frac{\text{Ker}(\mathcal{Q}) : \mathcal{H}^B \rightarrow \mathcal{H}^F}{\text{Im}(\mathcal{Q}) : \mathcal{H}^F \rightarrow \mathcal{H}^B}, \quad (2.18)$$

$$H^F(\mathcal{Q}) := \frac{\text{Ker}(\mathcal{Q}) : \mathcal{H}^F \rightarrow \mathcal{H}^B}{\text{Im}(\mathcal{Q}) : \mathcal{H}^B \rightarrow \mathcal{H}^F}, \quad (2.19)$$

consisting of all bosonic and fermionic states which are $\mathcal{Q}$ -closed but not $\mathcal{Q}$ -exact, respectively. The full cohomology group is

$$H(\mathcal{Q}) = H^B(\mathcal{Q}) \oplus H^F(\mathcal{Q}). \quad (2.20)$$

A complex for which $\text{Ker}(\mathcal{Q}) = \text{Im}(\mathcal{Q})$, and thus $H(\mathcal{Q}) = 0$, is called an *exact* sequence. Thus, cohomology groups are a measure of the failure of a sequence to be exact.¹³

Physically, an exact sequence corresponds to a system with no supersymmetric ground states. To see this, note that states with $E > 0$ that are $\mathcal{Q}$ -closed are necessarily $\mathcal{Q}$ -exact, as follows directly from the supersymmetry algebra (2.11). Indeed, applying the algebra on a state with energy E one has $\{\mathcal{Q}, \mathcal{Q}^\dagger\}|\psi\rangle = E|\psi\rangle$. Using that the state is $\mathcal{Q}$ -exact and assuming $E > 0$, implies $|\psi\rangle = \mathcal{Q}|\psi'\rangle$ with $|\psi'\rangle := \frac{1}{E}\mathcal{Q}^\dagger|\psi\rangle$ and thus the state $|\psi\rangle$ is $\mathcal{Q}$ -exact. On the other hand, all states with $E = 0$ are necessarily closed, as we have already seen, but *cannot* be exact. To show this suppose for the sake of contradiction that is the case, i.e., that there exists a state $|\psi'\rangle$ such that $|\Omega\rangle = \mathcal{Q}|\psi'\rangle$. Since $\mathcal{Q}$ commutes with H , the state $|\psi'\rangle$ must also have $E = 0$ and hence must be $\mathcal{Q}$ -closed, $\mathcal{Q}|\psi'\rangle = 0$. The upshot of this analysis is that $\text{Ker}(\mathcal{Q})|_{E>0} = \text{Im}(\mathcal{Q})|_{E>0}$ while $\text{Im}(\mathcal{Q})|_{E=0} = 0$. These two statements together imply

$$H^B(\mathcal{Q}) = \mathcal{H}_{E=0}^B, \quad H^F(\mathcal{Q}) = \mathcal{H}_{E=0}^F, \quad (2.21)$$

as claimed. That is, the cohomology of $\mathcal{Q}$ is given by the set of supersymmetric ground states. In particular, the total number of SUSY ground states is given by the dimension of the full cohomology group,

$$n_{E=0}^B + n_{E=0}^F = \dim H^B(\mathcal{Q}) + \dim H^F(\mathcal{Q}). \quad (2.22)$$

¹²In this case the cohomology group has the additional structure of a vector space but this is not necessary. See Section 4 for the general definition.

¹³For a textbook on algebraic topology see, e.g., <http://pi.math.cornell.edu/~hatcher/AT/ATplain.pdf> and for computational aspects see [21, 92].

The Witten index of such systems is then given by

$$\mathcal{I} = n_{E=0}^B - n_{E=0}^F = \dim H^B(\mathcal{Q}) - \dim H^F(\mathcal{Q}), \quad (2.23)$$

which coincides with the “Euler charactersitic” of the complex (2.17).

To summarize, finding the space of supersymmetric ground states is equivalent to finding the cohomology group $H(\mathcal{Q}) = H^B(\mathcal{Q}) \otimes H^F(\mathcal{Q})$ and computing the Witten index of the system is equivalent to computing the Euler characteristic of the cochain complex. This will play a crucial role in studying the complexity of the cohomology problem in Section 4.

Systems with $F \in \mathbb{N}_0$. In certain models, a more refined structure may appear. A case of particular interest to us is when there exists an operator F (usually called the fermion number operator) with integer-valued eigenvalues, $\mathbb{N}_0 = \{0, 1, \dots\}$, and thus one can identify the parity operator $(-1)^F$, literally as “minus one to the power F , ” or

$$(-1)^F = e^{i\pi F}. \quad (2.24)$$

The Hilbert space can then be graded more finely, with respect to eigenvalues of F . We will focus on systems in which the fermion number is bounded above and thus

$$\mathcal{H} = \mathcal{H}^0 \oplus \mathcal{H}^1 \oplus \dots \oplus \mathcal{H}^m, \quad (2.25)$$

where $\mathcal{H}^p$ is spanned by all states with $F = p$ and m is the maximum fermion number of the system. Then $\mathcal{H}_B$ and $\mathcal{H}_F$ correspond to states with even and odd fermion number, respectively.

By definition, an operator $\mathcal{O}$ has definite fermion number f , if

$$[F, \mathcal{O}] = f\mathcal{O}. \quad (2.26)$$

We will informally say that such an operator has fermion number $F = f$. Note that the action by the operator $\mathcal{O}$ with fermion number f on a state $|s\rangle$ with fermion number k gives a state $\mathcal{O}|s\rangle$ with fermion number $f+k$; $F\mathcal{O}|s\rangle = (f+k)\mathcal{O}|s\rangle$. An operator with $F = 0$ thus preserves the fermion number of the states it acts on and the operator is block diagonal in a the basis of states organized by fermion number.

Note that all that is required for defining supersymmetry is that $\mathcal{Q}$ maps states with even fermion number to states with odd fermion number and vice versa. Thus, the supercharge $\mathcal{Q}$ must have definite parity, $(-1)^F = -1$, but not necessarily definite fermion number F (the supercharge may be given, for instance, by a sum of terms each with different (odd) fermion numbers). In the case that $\mathcal{Q}$ has a definite fermion number and equal to 1, i.e.,

$$[F, \mathcal{Q}] = \mathcal{Q}, \quad (2.27)$$

the supercharge acts by sending states in $\mathcal{H}^p$ to states in $\mathcal{H}^{p+1}$ and the complex (2.17) splits into

$$C: \quad 0 \xrightarrow{\mathcal{Q}} \mathcal{H}^0 \xrightarrow{\mathcal{Q}} \mathcal{H}^1 \xrightarrow{\mathcal{Q}} \dots \xrightarrow{\mathcal{Q}} \mathcal{H}^m \xrightarrow{\mathcal{Q}} 0, \quad (2.28)$$

where we have included the empty spaces at the ends to emphasize there are no states with $F < 0$ or $F > m$. Then one may define the cohomology groups at each level,

$$H^p(\mathcal{Q}) := \frac{\text{Ker}(\mathcal{Q}) : \mathcal{H}^p \rightarrow \mathcal{H}^{p+1}}{\text{Im}(\mathcal{Q}) : \mathcal{H}^{p-1} \rightarrow \mathcal{H}^p} . \quad (2.29)$$

Then, each cohomology group H^p cohomology is in one-to-one correspondence with the space of supersymmetric ground states with fermion number p :

$$H^p(\mathcal{Q}) = \mathcal{H}_{E=0}^p . \quad (2.30)$$

The p th Betti number β_p is defined as the dimension of the p th cohomology group,

$$\beta_p := \dim H^p(\mathcal{Q}) . \quad (2.31)$$

The Witten index of the system is given by the alternating sum

$$\mathcal{I} = \sum_p (-1)^p \beta_p , \quad (2.32)$$

which coincides with the Euler characteristic of the complex (4.1).

2.3 Supersymmetry in lattice systems

A number of interesting models of supersymmetric quantum mechanics have been studied in the literature, with applications ranging from statistical mechanics to string theory and with deep connections to algebraic topology.¹⁴ In this paper we are mostly interested in the realization of $\mathcal{N} = 2$ quantum mechanics in a *discrete* system of spins or qubits, as first discussed in [66]. It is most convenient to use the fermionic representation of states. In this model one considers m vertices of a graph G , each of which can be occupied by 0 or 1 spinless fermions. A fermion at vertex i is created by an operator $a_i^\dagger$ and annihilated by a_i , satisfying the standard anticommutation relations

$$\{a_i, a_j^\dagger\} = \delta_{ij} , \quad i, j \in \{1, \dots, m\} . \quad (2.33)$$

$|0 \dots 0\rangle$ is the state with no fermionic excitations, which is annihilated by all a_i , i.e., $a_i |0 \dots 0\rangle = 0$ for all i . The fermion number operator in Fock space is defined as

$$F = \sum_{i=1}^m a_i^\dagger a_i . \quad (2.34)$$

The 2^m -dimensional Fock space is then constructed by acting with creation operators on the vacuum state, as

$$|n_1 \dots n_m\rangle = (a_1^\dagger)^{n_1} \dots (a_m^\dagger)^{n_m} |0 \dots 0\rangle . \quad (2.35)$$

¹⁴For instance, a supersymmetric version of a free particle moving in a curved manifold leads to a description of de Rham cohomology and turning on a potential leads to Morse theory [89].

Note that the fundamental modes $a_i^\dagger$ are fermionic. However, the multiparticle states $|n_1 \dots n_m\rangle$ can be either bosonic or fermionic depending on whether the total number of fermionic particles in the state is even or odd, respectively, i.e.,

$$(-1)^F |n_1 \dots n_m\rangle = \begin{cases} +|n_1 \dots n_m\rangle & \text{for } \sum_i n_i \in 2\mathbb{Z} \\ -|n_1 \dots n_m\rangle & \text{for } \sum_i n_i \in 2\mathbb{Z} + 1 \end{cases} \quad (2.36)$$

Any operator in Fock space can be constructed as a polynomial in the fermionic operators a_i and $a_i^\dagger$. The supercharge $\mathcal{Q}$ must be a polynomial with only odd degree monomials in order to be consistent with its fermionic nature. In the model of [66] the supercharges are cubic functions of the creation and annihilation operators. More generally one can write

$$\begin{aligned} \mathcal{Q} = & \sum_i (\alpha_i a_i^\dagger + \bar{\alpha}_i a_i) + \sum_{ijk} (\alpha_{ijk} a_i^\dagger a_j^\dagger a_k^\dagger + \beta_{ijk} a_i^\dagger a_j^\dagger a_k + \gamma_{ijk} a_i^\dagger a_j a_k) \\ & + \sum_{ijk} (\bar{\alpha}_{ijk} a_i a_j a_k + \bar{\beta}_{ijk} a_i a_j a_k^\dagger + \bar{\gamma}_{ijk} a_i a_j^\dagger a_k^\dagger) + \dots \end{aligned} \quad (2.37)$$

where the $\alpha_i, \bar{\alpha}_i, \alpha_{ijk}, \bar{\alpha}_{ijk}$, etc. are complex coefficients which are constrained by the nilpotency condition (2.10) and the ellipses denote possible higher power terms. Any such supercharge defines a supersymmetric Hamiltonian via (2.11) which, by construction, is parity preserving. If each of the terms above involve at most k fermionic modes we say the supercharge is k -local and the resulting Hamiltonian is (at most) $2k$ -local.

Note that the supercharge above has definite parity $(-1)^F$ but does not necessarily have a definite fermion number F , as it is given by the sum of terms with different odd values of F . We will be particularly interested in a large family of supersymmetric systems which arises from imposing two separate conditions on the supercharge. The first is that the supercharge has definite fermion number $F = 1$. We require this for later applications to the cohomology and Betti number problems. Any supercharge with $F = 1$ can be written as

$$\mathcal{Q} = \sum_i a_i^\dagger B_i(a, a^\dagger), \quad (2.38)$$

where the B_i are bosonic functions with $F = 0$ and further constrained by the condition that the supercharge is nilpotent, $\mathcal{Q}^2 = 0$. The second condition we will impose is that the supercharge is k -local and thus each B^i above can be written as

$$B_i(a, a^\dagger) = \sum_{S_i} \mathcal{B}_{S_i}(a, a^\dagger) \quad (2.39)$$

where for each i , $S_i \subseteq \{1, \dots, m\}$ is a subset of m vertices of size at most $k - 1$, and the $\mathcal{B}_{S_i}(a, a^\dagger)$ are bosonic operators all with $F = 0$ acting on the sites S_i . We thus require that each $\mathcal{B}_{S_i}$ is a $(k - 1)$ -local operator in fermionic space with fermion number $F = 0$.

Examples. The simplest model is the 1-local supercharge

$$\mathcal{Q} = \sum_i a_i^\dagger, \quad (2.40)$$

which one can easily check squares to zero and has $F = 1$. The resulting Hamiltonian, however, is proportional to the identity operator and is thus rather trivial (in particular, it has no supersymmetric ground states). Examples with a nontrivial Hamiltonian are the original models considered in [66] or the $\mathcal{N} = 2$ version of the SYK model [32], given by

$$Q = \sum_{ijk} \alpha_{ijk} a_i^\dagger a_j^\dagger a_k^\dagger, \quad (2.41)$$

where the α_{ijk} is completely antisymmetric and the supercharge has $F = 3$. This model has an exponential number of supersymmetric ground states.

Finally, a very interesting class of models is the fermion hard-core models of [26]. These are specified by a choice of a graph $G = (V, E)$ and

$$\mathcal{Q} = \sum_{i \in V} a_i^\dagger P_i, \quad P_i = \prod_{j \mid (i,j) \in E} (1 - \hat{n}_j). \quad (2.42)$$

The supercharge has $F = 1$ and is $(\delta + 1)$ -local, with δ the maximum degree of the graph. The resulting Hamiltonian is the one in (1.1). We will return to this model in Section 6 in relation to topological data analysis.

In the next few sections we will not focus on particular models but rather on studying the computational complexity of generic supersymmetric systems and its implications for the (co)homology problem.

3 The SUSY Local Hamiltonian Problem

In this section, we determine the complexity of determining the ground state of supersymmetric systems, i.e., a supersymmetric version of k -LOCAL HAMILTONIAN studied by Kitaev. It was suggested in [17] that the problem remains QMA-complete for quantum mechanical Hamiltonians with $\mathcal{N} = 2$ supersymmetry and we show here that this is indeed the case. As we discuss in detail in Section 4, a consequence of this result is that a certain promise version of the cohomology problem is QMA₁-hard and in certain cases QMA₁-complete (we define QMA and related complexity classes in Appendix A).

3.1 Definitions

We begin by reviewing k -LOCAL HAMILTONIAN, defined as follows [54]:

Definition 1 (k -LOCAL HAMILTONIAN)

Input: An n -qubit local Hamiltonian, $H = \sum_a H_a$, where each H_a acts on at most k qubits, and two numbers, $b > a \geq 0$, with $b - a > \frac{1}{\text{poly}(n)}$.

Promise: The lowest energy level E_0 of the Hamiltonian is either $E_0 \leq a$ or $E_0 \geq b$.

Problem: Output YES if $E_0 \leq a$ and NO if $E_0 \geq b$.

A well known series of results established that k -LOCAL HAMILTONIAN is QMA-complete for $k \geq 2$ [54, 52, 53]. This is also the case for 2-local Hamiltonians in fermionic Fock space [59]. As reviewed in Section 2, supersymmetric Hamiltonians are given in terms of the supercharge $\mathcal{Q}$ which, due to its anticommuting nature, is most naturally described in fermionic Fock space. We thus formulate the supersymmetric Hamiltonian problem in fermionic Fock space. Precisely, we define the $\mathcal{N} = 2$ supersymmetric version of the local Hamiltonian problem as follows:

Definition 2 (SUSY k -LOCAL HAMILTONIAN)

Input: A local fermionic supercharge, $\mathcal{Q} = \sum_a \mathcal{Q}_a$, acting on the Fock space of n fermionic modes, and two numbers $b' > a' \geq 0$, with $b' - a' > \frac{1}{\text{poly}(n)}$.

Promise: The lowest energy level E_0 of the supersymmetric k -local Hamiltonian, $H = \{\mathcal{Q}, \mathcal{Q}^\dagger\}$, is such that either $E_0 \leq a'$ or $E_0 \geq b'$.

Problem: Output YES if $E_0 \leq a'$ and NO if $E_0 \geq b'$.

Note that although a local supercharge is given as input, the degree of locality k refers to that of the Hamiltonian, which is (at most) twice the degree of locality of $\mathcal{Q}$.

To provide a description of a k -local supercharge $\mathcal{Q} = \sum_a \mathcal{Q}_a$ as input it is sufficient to explicitly provide the individual operators $\mathcal{Q}_a$. We may at times drop the restriction that an operator is k -local, and instead require that it simply be *sparse*, i.e., it has at most $\text{poly}(n)$ non-zero entries per row and column, and that the input is provided via unitaries that allow us to access the non-zero entries. We call such access *sparse access*.¹⁵ In either case we always assume that the entries of $\mathcal{Q}$ are upper bounded by a polynomial, and hence that $\|\mathcal{Q}\| = O(\text{poly}(n))$.

3.2 The SUSY local Hamiltonian problem is QMA-complete

The goal of this section is to prove the following Theorem:

Theorem 1. SUSY k -LOCAL HAMILTONIAN is QMA-complete for $k \geq 4$.

We begin by first establishing the following proposition:

Proposition 1. SUSY k -LOCAL HAMILTONIAN is in QMA for any constant k , and remains in QMA when the k -local restriction is dropped and we require only that $\mathcal{Q}$ is sparse.

¹⁵We define this as having access to classical (or quantum) subroutines that, given an index i , can output a list running over all non-zero entries in row or column i , and a unitary that allows us to query the value of the i, j th entry of d . More precisely, we assume we have access to three unitaries O_{row} , O_{col} , and O_d , where: $O_{\text{row}} : |i, l\rangle |0\rangle \mapsto |i, l\rangle |I(i, l)\rangle$ and $O_{\text{col}} : |j, l\rangle |0\rangle \mapsto |j, l\rangle |J(j, l)\rangle$, for $0 \leq i, j \leq 2^n$, $0 \leq l \leq \text{poly}(n)$, and where $I(i, l)$ denotes the l th non-zero entry in row i of d and $J(j, l)$ the l th non-zero entry in column j of d . And finally $O_d : |i, j\rangle |0\rangle \mapsto |i, j\rangle |d_{ij}\rangle$.

Proof. It suffices to prove the containment for sparse $\mathcal{Q}$ since any k -local operator is automatically sparse. Since a supersymmetric Hamiltonian is a special case of a general Hamiltonian, one may think that SUSY k -LOCAL HAMILTONIAN is automatically inside QMA. This is not so straightforward since in the supersymmetric version we are given efficient (sparse) access to $\mathcal{Q}$, which does not in general ensure efficient access to $\{\mathcal{Q}, \mathcal{Q}^\dagger\}$, which might not be sparse, and hence using Hamiltonian simulation to obtain an approximation of $e^{i\{\mathcal{Q}, \mathcal{Q}^\dagger\}}$ will generally not be efficient. This difficulty is easily overcome, however, by noting that the $\mathcal{N} = 2$ SUSY algebra (2.10) and (2.11) imply

$$(\mathcal{Q} + \mathcal{Q}^\dagger)^2 = \{\mathcal{Q}, \mathcal{Q}^\dagger\} = H, \quad (3.1)$$

and thus the eigenvalues of H are the squares of the eigenvalues of the Hermitian operator $\mathcal{Q} + \mathcal{Q}^\dagger$, to which we do have efficient access. Thus, a witness for a YES instance is an eigenstate of the supercharge $\mathcal{Q} + \mathcal{Q}^\dagger$ with eigenvalue $-\sqrt{a'} \leq \lambda \leq \sqrt{a'}$, which implies that the smallest eigenvalue of H satisfies $E_0 \leq a'$. Given such a witness we simply run phase estimation on the unitary operator $e^{i(\mathcal{Q} + \mathcal{Q}^\dagger)}$, or an approximation of it obtained via Hamiltonian simulation using, e.g., the algorithm of [8], with the witness as input. Since we are promised that either $E_0 \leq a'$ or $E_0 \geq b'$ with $b' - a' > 1/\text{poly}(n)$, we only need to obtain an estimate of the eigenvalue of $\mathcal{Q} + \mathcal{Q}^\dagger$ on the witness to additive accuracy $\epsilon = 1/O(\text{poly}(n))$ to verify that $E_0 \leq a'$. Using standard techniques, the run-time of Hamiltonian simulation plus phase estimation will be $\text{poly}(n, 1/\epsilon, \|\mathcal{Q}\|)$, which is $\text{poly}(n)$ for our choice of ϵ , and thus the problem is in QMA for sparse $\mathcal{Q}$. $\square$

As mentioned, we will prove QMA-hardness of the SUSY Hamiltonian problem by reduction from standard k -LOCAL HAMILTONIAN. Since the latter is formulated in qubit space, we will need to provide an encoding of qubit space into fermionic Fock space. There are various ways of doing this, but the most useful to our purposes is to represent a qubit as a single fermion which can be in two different modes, a_i, b_i , as in [59]. Namely, we consider a set of $m = 2n$ fermionic modes satisfying the algebra

$$\{a_i^\dagger, a_j\} = \delta_{ij}, \quad \{b_i^\dagger, b_j\} = \delta_{ij}, \quad (3.2)$$

and all other anticommutators vanishing. An n -qubit basis state $|n_1 \cdots n_n\rangle$ is mapped into a state $|\psi_{n_1 \cdots n_n}\rangle$ in fermionic Fock space as

$$|n_1 \cdots n_n\rangle \rightarrow |\psi_{n_1 \cdots n_n}\rangle := (a_1^\dagger)^{1-n_1} (b_1^\dagger)^{n_1} \cdots (a_n^\dagger)^{1-n_n} (b_n^\dagger)^{n_n} |vac\rangle \quad (3.3)$$

and $|vac\rangle$ is the vacuum state with no fermions. Importantly, note that the states $|\psi_{n_1 \cdots n_n}\rangle$ have all fermion number $F = n$.

Single qubit operators are mapped into bilinears in these fermion operators, as

$$\sigma_i^+ \rightarrow a_i^\dagger b_i, \quad \sigma_i^- \rightarrow b_i^\dagger a_i, \quad \sigma_i^z \rightarrow a_i^\dagger a_i - b_i^\dagger b_i; \quad (3.4)$$

one can easily check this is consistent with the algebra of Pauli matrices.¹⁶ Note that the map is local, sending qubit operators at site i to bilinears in fermionic operators at site i . The total fermion number

¹⁶This is simply the well known fermionic representation of spin- $\frac{1}{2}$ operators, with $a^\dagger = a_\uparrow^\dagger$ creating a spin up and $b^\dagger = a_\downarrow^\dagger$ creating a spin down.

operator is

$$F = \sum_i (a_i^\dagger a_i + b_i^\dagger b_i), \quad (3.5)$$

and all bilinears in (3.4) have $F = 0$.

Then, given a k -local qubit operator A

$$A = \sum_S A_S(\sigma_i^+, \sigma_i^-, \sigma_i^z), \quad (3.6)$$

applying the map (3.4) to each term gives

$$\hat{A}(a_i, b_i, a_i^\dagger, b_i^\dagger) := \sum_S A_S(a_i^\dagger b_i, b_i^\dagger a_i, a_i^\dagger a_i - b_i^\dagger b_i), \quad (3.7)$$

which is a k -local operator in fermionic Fock space, with fermion number $F = 0$.

Note that $\hat{A}$ acts on the entire 2^{2n} -dimensional Fock space. When restricted to states in the space spanned by (3.3), however, it coincides with the qubit operator A . Indeed, consider two arbitrary n -qubit states,

$$|\psi\rangle := \sum c_{n_1 \dots n_n} |n_1, \dots, n_n\rangle, \quad |\psi'\rangle := \sum c'_{n_1 \dots n_n} |n_1, \dots, n_n\rangle. \quad (3.8)$$

Applying the map (3.3) to each basis element gives the corresponding states in fermionic Fock space,

$$|\psi_n\rangle = \sum c_{n_1 \dots n_n} |\psi_{n_1 \dots n_n}\rangle, \quad |\psi'_n\rangle = \sum c'_{n_1 \dots n_n} |\psi_{n_1 \dots n_n}\rangle. \quad (3.9)$$

Then, one can see that

$$\langle \psi' | A | \psi \rangle = \langle \psi'_n | \hat{A} | \psi_n \rangle. \quad (3.10)$$

This will play an important role in the reduction from hard problems in qubit space to hard problems in fermionic Fock space.

Theorem 2. SUSY 4-LOCAL HAMILTONIAN is QMA-hard.

Proof. We will prove hardness by reduction from the standard k -LOCAL HAMILTONIAN. Consider an instance of k -LOCAL HAMILTONIAN specified by a Hermitian matrix $A = A^\dagger = \sum_{S=1}^L A_S$ acting on n qubits, and parameters (a, b) .¹⁷ Then, applying the map (3.4) we construct the operator

$$B(a_i, b_i, a_i^\dagger, b_i^\dagger) := \sum_{S=1}^L A_S (a_i^\dagger b_i, b_i^\dagger a_i, a_i^\dagger a_i - b_i^\dagger b_i) + B_{pen}, \quad (3.11)$$

where

$$B_{pen} = J(n) \sum_{i=1}^n [n_{a_i} n_{b_i} + (n_{a_i} - 1)(n_{b_i} - 1)], \quad (3.12)$$

with $J(n)$ a coupling constant. Note that B has fermion number $F = 0$. The term B_{pen} penalizes all states which do not have exactly one fermionic mode at each site $i = (1, \dots, n)$. Note that the

¹⁷We reserve the symbol H for the actual Hamiltonian of our system which will be related, but not identical, to A .

construction of the operator B requires only a polynomial number of operations since $L \leq \binom{n}{k} = O(n^k)$. Let us introduce an auxiliary pair of fermionic modes (a_0, b_0) at a site $i = 0$. The total set of fermionic operators is (a_0, b_0, a_i, b_i) with $i = 1, \dots, n$. Let us consider the supercharge

$$\mathcal{Q} = \frac{1}{\sqrt{2}}(a_0^\dagger + b_0^\dagger) B(a_i, b_i, a_i^\dagger, b_i^\dagger). \quad (3.13)$$

Note that the fermionic operators at site $i = 0$ do not appear inside B ; their role is simply to ensure that $\mathcal{Q}^2 = (\mathcal{Q}^\dagger)^2 = 0$ for arbitrary B and that $\mathcal{Q}$ has fermion number $F = 1$. The hardness of the problem is encoded entirely in B . Indeed, the resulting Hamiltonian is

$$H = \{\mathcal{Q}, \mathcal{Q}^\dagger\} = B^2, \quad (3.14)$$

and the eigenvalues of H are simply the square of those of B . Now, by choosing $J(n)$ large enough (in particular larger than $\|A\| = O(\text{poly}(n))$) the ground space of B is necessarily in the subspace in which $B_{\text{pen}} = 0$. Note that

$$B_{\text{pen}} = 0 \quad \Leftrightarrow \quad (n_{a_i}, n_{b_i}) = (1, 0) \text{ or } (0, 1), \quad (3.15)$$

and each site $i = (1, \dots, n)$ is occupied by exactly one fermion in this subspace. This is precisely the Fermionic Fock subspace spanned by the states on the RHS of (3.3). The operator B on this subspace coincides with the qubit operator A by (3.10). Thus, if A is a k -local Hamiltonian satisfying the promise with parameters (a, b) then the SUSY Hamiltonian is $2k$ -local and satisfies the promise with parameters $(a', b') = (a^2, b^2)$. Furthermore, note that

$$b' - a' = (b - a)(b + a) \geq (b - a)^2 > \frac{1}{\text{poly}(n)}, \quad (3.16)$$

and the separation remains inverse polynomial, as required. Thus, a YES/NO instance of k -LOCAL HAMILTONIAN implies a YES/NO instance of SUSY $2k$ -LOCAL HAMILTONIAN. Since 2-LOCAL HAMILTONIAN is QMA-hard [53], it follows that SUSY 4-LOCAL HAMILTONIAN is QMA-hard. $\square$

It would be interesting to determine the hardness of SUSY k -LOCAL HAMILTONIAN for $2 \leq k \leq 3$. It would also be interesting to study whether the problem remains hard for systems with a larger degree of supersymmetry, i.e., $\mathcal{N} > 2$. Finally, note that although in this construction the supersymmetric Hamiltonian is local, in the sense that it is given by a sum of terms involving a constant number of sites, it is not *geometrically* local, even if B is geometrically local since taking a square of a geometrically local function is in general not geometrically local. It would be interesting to study whether SUSY k -LOCAL HAMILTONIAN remains hard for geometrically local Hamiltonians.

4 The k -local Cohomology Problem

In this section, we introduce and study the computational complexity of the cohomology problem. As mentioned in the Introduction, this is a fundamental problem in homological algebra, encompassing a large number of problems from a myriad of mathematical fields including topology, group theory, and

number theory. See [15, 85] for standard textbooks on homological algebra and [50] for a more accessible introduction discussing various applications and computational aspects.¹⁸ We will restrict ourselves to a subclass of cohomology problems to which the tools of quantum Hamiltonian complexity can be readily applied but which at the same time is rich enough to encompass a wide range of interesting problems. Precisely, we introduce the problem k -LOCAL COHOMOLOGY and show that it is QMA_1 -hard and contained in QMA . The setup is heavily motivated by the supersymmetric systems studied in Section 3.2 and captures the computational problems that arise in these systems.

Although we have briefly discussed cohomology in Section 2.2 in the context of supersymmetric quantum mechanics, we now review the general definition of a cochain complex and cohomology groups in a self contained way.

Definition 1 (Cochain complex). *A cochain complex is a sequence of Abelian groups, $\{C^p\}$, together with a “coboundary” operator d :*

$$C : \quad \dots \rightarrow C^{p-1} \xrightarrow{d} C^p \xrightarrow{d} C^{p+1} \rightarrow \dots, \quad (4.1)$$

with

$$d^2 = 0. \quad (4.2)$$

The elements of C^p are called p -cochains or simply cochains.

A special role is played by the notion of “cocycles” and “coboundaries.”

Definition 2 (Cocycle, coboundary). *A cochain $c^p \in C^p$ is said to be a p -cocycle if it is in the kernel of the coboundary operator, i.e.,*

$$dc^p = 0. \quad (4.3)$$

A cochain is said to be a p -coboundary if it is in the image of d , i.e., if there exists some $c^{p-1} \in C^{p-1}$ such that

$$c^p = dc^{p-1}. \quad (4.4)$$

Note that due to the nilpotency property (4.2) if c^p is a cocycle, so is $\tilde{c}^p = c^p + dc^{p-1}$. One then says that the cocycles c^p and $\tilde{c}^p$ are in the same “cohomology” class,

$$c^p \sim \tilde{c}^p \quad \text{if} \quad c^p = \tilde{c}^p + dc^{p-1}. \quad (4.5)$$

The element c^p (or equivalently $\tilde{c}^p$) is, by definition, a representative of the cohomology class $H^p(d)$. More precisely,

Definition 3 (Cohomology group). *The p th cohomology group is*

$$H^p(d) := \frac{\text{Ker}(d) : C^p \rightarrow C^{p+1}}{\text{Im}(d) : C^{p-1} \rightarrow C^p}, \quad (4.6)$$

consisting of p -cocycles that are not p -coboundaries.

¹⁸For applications in algebraic topology see, e.g., <http://pi.math.cornell.edu/~hatcher/AT/ATplain.pdf> and [21, 92] for computational aspects.




Figure 3: Relationship between the cochains C^{p-1}, C^p, C^{p+1} , the cocycles Z^{p-1}, Z^p, Z^{p+1} , and the coboundaries B^{p-1}, B^p, B^{p+1} . The p th cohomology group consists of those elements which are cocycles but not coboundaries, i.e. $H^p(d) = Z^p/B^p$.

The collection of p -cocycles form the subgroup Z^p , and the collection of p -coboundaries the subgroup B^p , so that $H^p(d) = \frac{Z^p}{B^p}$. The relationships between these groups is illustrated in Figure 3.

Definition 4 (Betti number). *The p th Betti number is defined as the rank of the p th cohomology group,*

$$\beta_p := \text{rank } H^p(d). \quad (4.7)$$

Note that the full cohomology group is

$$H(d) := \bigotimes_p H^p(d), \quad (4.8)$$

with total rank the sum of Betti numbers,

$$\text{rank } H(d) = \sum_p \beta_p. \quad (4.9)$$

The “Euler characteristic” of the complex is defined as the alternating sum,

$$\chi := \sum_p (-1)^p \beta_p. \quad (4.10)$$

Finally, one can similarly define a *chain* complex as a sequence of Abelian groups $\{C_p\}$ connected by a boundary operator ∂ , satisfying $\partial^2 = 0$, in which the arrows are reversed compared to cohomology:

$$C : \quad \cdots \leftarrow C_{p-1} \xleftarrow{\partial} C_p \xleftarrow{\partial} C_{p+1} \leftarrow \cdots, \quad (4.11)$$

and one defines the *homology* groups

$$H_p(d) := \frac{\text{Ker}(\partial) : C_p \rightarrow C_{p-1}}{\text{Im}(\partial) : C_{p+1} \rightarrow C_p}, \quad (4.12)$$

consisting of p -cycles that are not p -boundaries.

Cohomology is often presented as the “dual” version of homology. In our setting cohomology groups will also define homology groups by conjugation of the operator d . Since the notion of cohomology

connects more naturally to supersymmetry, we thus focus on cohomology from now on. The hardness results we derive for cohomology hold also for homology.

The computational problem we wish to consider is: given a cochain complex C , to decide whether a certain cohomology group $H^l(d)$ is nontrivial. We refer to this as the “cohomology problem:”

Input: A cochain complex C and an integer l .

Problem: Output YES if $H^l(d) \neq 0$ and NO otherwise.

Note that a general complex C can consist of an infinite sequence. Since we consider C to be part of the input, we will focus on finite sequences. Furthermore, in order to make contact with quantum Hamiltonian complexity, in which the operator d is taken to act on a Hilbert space, we consider the special case in which the Abelian groups C^p have the additional structure of a vector space, which we denote by V^p . All in all, we consider finite sequences of *vector spaces*,

$$C : \quad 0 \xrightarrow{d} V^0 \xrightarrow{d} \dots \xrightarrow{d} V^m \xrightarrow{d} 0, \quad (4.13)$$

and take m to be the length of the input. We will further endow the full vector space $V := \bigoplus_{p=0}^m V^p$ with an inner product and a corresponding notion of conjugation, which lets us define the operator $d^\dagger$. As we discuss below, this is crucial in our setting for defining the appropriate promise version of the cohomology problem. From now on we assume that V is a finite dimensional vector space and thus d and $d^\dagger$ are finite dimensional operators acting on this space.

Given a notion of conjugation, it is convenient to introduce the “Laplacian” operator

Definition 5 (Laplacian operator). *For coboundary operator d , the Laplacian operator is defined as*

$$\Delta := dd^\dagger + d^\dagger d. \quad (4.14)$$

Since this is a positive semidefinite operator, it follows that

$$\Delta v = 0 \quad \Leftrightarrow \quad dv = d^\dagger v = 0. \quad (4.15)$$

It is also easy to see that an element $v \in V$ is an element of cohomology $H(d)$ if and only if it is a zero-eigenvalue of the Laplacian:

$$v \in H(d) \quad \Leftrightarrow \quad \Delta v = 0. \quad (4.16)$$

In the context of de Rham cohomology this is known as Hodge’s theorem. In the language of supersymmetry, this is the statement that elements of cohomology of the coboundary operator are in one-to-one correspondence with supersymmetric ground states and we have given a proof of this below Eq. (2.6).

Let $d^l := d|_{V^l}$ and $(d^\dagger)^l := d^\dagger|_{V^l}$ denote the restriction of the operators d and $d^\dagger$ to the space V^l , respectively.¹⁹ Similarly, let $\Delta^l := \Delta|_{V^l}$ denote the restriction of the Laplacian to V^l . Then, one can

¹⁹We will often omit the subscript on d when this is redundant to simplify notation.

write²⁰

$$\Delta^l = d^{l-1}(d^\dagger)^{l-1} + (d^\dagger)^l d^l. \quad (4.17)$$

It follows trivially from (4.16) that $v^l \in V^l$ is an element of the cohomology group $H^l(d)$ if it is a zero-eigenvalue of the corresponding Laplacian:

$$v^l \in H^l(d) \quad \Leftrightarrow \quad \Delta^l v^l = 0. \quad (4.18)$$

Although the cohomology problem is rather fundamental, and it is believed to be hard, little is known about its precise complexity. In the case of simplicial complexes, the question was posed in [51]. It was recently shown in [4] that for independence complex the problem is **NP**-hard (given a graph as input) and generally not in **NP**.²¹ As we describe below, however, a certain promise version of the cohomology problem can be placed inside QMA.

4.1 The k -local cohomology problem

We now introduce a version of the cohomology problem, inspired by supersymmetric systems discussed in previous sections. Indeed, note that the supersymmetric systems discussed in Section 2.2 precisely lead to a cochain structure of the form (4.13), with the coboundary operator identified with the supercharge and the Laplacian with the supersymmetric Hamiltonian:

$$\mathcal{Q} = d, \quad H = \Delta. \quad (4.19)$$

In fact, this is how the notion of cohomology was introduced there. The supersymmetric models of Section 2.3 are of particular interest since they describe systems with discrete degrees of freedom and are thus more amenable to a computational treatment.²² We thus take inspiration from these to define the cohomology problem we will study. First, we note that in these models the vector space V is the Hilbert space of the system, which can generally be a *subset* of the fermionic Fock space of m fermions,

$$V \subseteq \text{span}\{|n_1, \dots, n_m\rangle\}, \quad n_i = \{0, 1\}. \quad (4.20)$$

This happens whenever certain states are excluded from the Hilbert space for dynamical reasons, such as a strong repulsion forbidding certain states corresponding to nearby particles. The grading of $V = \bigoplus_{p=0}^m V^p$ is provided by the fermion number F , i.e.,

$$V^p = \text{span}\{v \in V \mid Fv = pv\}. \quad (4.21)$$

²⁰Note that $(d^\dagger)^{l+1} = (d^l)^\dagger$ and thus one can also write $\Delta^l = d^{l-1}(d^\dagger)^l + (d^\dagger)^{l+1}d^l$.

²¹The problem is described in [4] in terms of homology rather than cohomology, but as already mentioned these are equivalent for these complexes.

²²It would also be interesting to study appropriate discretizations of continuum quantum mechanical systems.

We take the coboundary operator to have fermion number 1 and thus it acts as $d : V^p \rightarrow V^{p+1}$. Given such a cochain complex $C = (V, d)$ and an integer l we would like to determine whether the cohomology group $H^l(d)$ is nontrivial. We will input the description of a cochain complex as a tuple $(V[m], d[k])$, where for clarity we use the notation $[\cdot]$ to explicitly denote a parameter of V or d . As always, we need to provide a succinct description of the input. For our purposes, we will take (unless stated otherwise) the input to satisfy the following:

- The graded vector space $V[m] = V = \bigoplus_{p=0}^m V^p$ is a vector space spanned by a *subset* of basis states in the fermionic Fock space of m fermionic modes, $|n_1, \dots, n_m\rangle$. The subset will be determined by a polynomial-sized (in m) list of constraints among the n_i . This list might be specified in a number of ways: for instance, we could be given a list of constraints via a (polynomial) set of Boolean functions, a Boolean formula in conjunctive normal form, or an arithmetic circuit over the integers. We will mostly remain agnostic about how these constraints are supplied, only that they contain enough information to define the full vector space V . Often we will consider when this space can be efficiently “sampled” from (see Definition 6).
- The coboundary operator $d[k] = d$ is a k -local operator acting on V and with fermion number $F = 1$. Any such operator can be written as

$$d = \sum_{i=1}^m a_i^\dagger B_i(a, a^\dagger), \quad (4.22)$$

where the B_i are $(k-1)$ -local functions such that $d^2 = 0$. Furthermore, each B_i has $F = 0$ which ensures that $d : V^p \rightarrow V^{p+1}$ and defines a cochain of the form (4.13).²³ Note that since V is a subspace of fermionic Fock space, there is a natural notion of conjugation given by the standard Hermitian conjugate. Indeed, we define

$$d^\dagger = \sum_{i=1}^m B_i^\dagger(a, a^\dagger) a_i. \quad (4.23)$$

The size of the input in this case is m , the number of fermionic modes. This is natural, since the full Fock space on m modes will have dimension 2^m , and the topological objects we consider (e.g. Laplacians) will often be $\approx 2^m$ -dimensional.

Let us note that if, instead, the *full* vector space V is provided as a set of all its basis elements, and the coboundary operator as the full matrix acting on these basis elements, the cohomology problem is in P (taking the dimension of V to be the size of the input). Indeed, given such a description one simply needs to diagonalize the Laplacian operator Δ and check, for a given l , whether there are any zero-eigenvalues in the subspace V^l . If so, the cohomology group $H^l(d)$ is nontrivial and trivial otherwise.

²³One could relax this condition and simply require that the B_i have even fermion number. This defines a $\mathbb{Z}_2$ -graded complex. See Section 2.2 for a discussion in the context of supersymmetric systems. Since a $\mathbb{Z}_2$ -graded complex is obtained from the complexes we study, identifying $F \bmod 2$, all the hardness results we derive here hold for $\mathbb{Z}_2$ -graded complexes as well.

Of course, the representation of V as a list of all its basis elements is not efficient, since it can contain an exponential (in m) number of elements.

We comment that the description of the complex $(V[m], d[k])$ given above is a natural generalization of some instances considered in the literature. In particular, the problem of determining whether the l th *homology* group $H_l(\partial)$ of the independence complex was recently studied in [4], where the input to the problem is taken to be a graph G . In our setting, the input $(V[n], d[k])$ is specified as follows.²⁴ The vector space is defined by the $\mathcal{O}(n^2)$ number of constraints

$$A_{ij}n_i n_j = 0, \quad (4.24)$$

where A_{ij} is the incidence matrix of G . All solutions to these constraints correspond to independent sets of G . The coboundary operator is given by

$$d[k] = \sum_i a_i^\dagger P_i, \quad P_i = \prod_j (1 - A_{ij}n_j), \quad (4.25)$$

where $k = \delta + 1$ with δ the maximum degree of the graph. Thus both $V[n]$ and $d[k]$ can be computed efficiently from G . Our description of chain complexes is more general and thus encompasses a larger number of problems.

Having discussed the form of the input, we define the following problem:

Definition 3 (k -LOCAL COHOMOLOGY)

Input: A cochain complex $(V[m], d[k])$ and an integer l .

Promise: Either there exists a state $|\Omega^l\rangle \in V^l$ such that $d|\Omega^l\rangle = d^\dagger|\Omega^l\rangle = 0$, and hence $H^l(d) \neq 0$, or otherwise $\|(d \pm d^\dagger)|\Psi^l\rangle\| \geq \epsilon$ for all $|\Psi^l\rangle \in V^l$, with $\epsilon = \frac{1}{\text{poly}(m)}$.

Problem: Output YES if the former and NO if the latter.

Note that the promise requires the notion of conjugation and an inner product, which is why the groups C^p are taken to be vector spaces endowed with an inner product and notion of conjugation.

One can similarly define the “dual” version of the problem, specified instead by a vector space $V = \bigoplus_{p=0}^m V_p$ and a k -local *boundary* operator $\partial : V_p \rightarrow V_{p-1}$, $\partial^2 = 0$, in the same way:

Definition 4 (k -LOCAL HOMOLOGY)

Input: A chain complex $(V[m], \partial[k])$ and an integer l .

Promise: Either there exists a state $|\Omega_l\rangle \in V_l$ such that $\partial|\Omega_l\rangle = \partial^\dagger|\Omega_l\rangle = 0$, and hence $H_l(\partial) \neq 0$, or

²⁴The homology groups $H_l(\partial)$ with ∂ the standard boundary operator of simplicial complexes are isomorphic to the cohomology groups of the coboundary operator in (4.25), $H_l(\partial) \simeq H^l(d)$ so these are equivalent computational problems—we expand on this in Section 6.

otherwise $\|(\partial \pm \partial^\dagger)|\Psi_l\rangle\| \geq \epsilon$ for all $|\Psi_l\rangle \in V_l$, with $\epsilon = \frac{1}{\text{poly}(m)}$.

Problem: Output YES if the former and NO if the latter.

Note that since the notion of conjugation is defined for all these complexes, the operator $d^\dagger$ defines a boundary operator ∂ and it follows that all the complexity results we derive for k -LOCAL COHOMOLOGY also hold for k -LOCAL HOMOLOGY. We will not make this explicit below.

We note that k -LOCAL COHOMOLOGY (and k -LOCAL HOMOLOGY) admits an efficient algorithm for constant l , even if V itself is exponential in size. Indeed, for constant l there are at most $O(m^l)$ elements in V^l . Thus, the operators d^l and Δ^l are polynomial-size matrices and determining if they have any zero-eigenvalues can be done in polynomial time (see [21] for an overview of various algorithms).

For non-constant l , however, the matrices can be exponentially large and we expect no algorithm (classical or quantum) to be efficient. Indeed, strong evidence for this is provided by the fact that these problems are QMA₁-hard as we show below. We will prove hardness by reduction from QUANTUM k -SAT, introduced in [10]. For reference, recall this is defined as:

Definition 5 (QUANTUM k -SAT [10])

Input: A set of k -local projectors Π_S where S are possible subsets of $\{1, \dots, n\}$ of cardinality k .

Promise: Either there exists a state $|\psi\rangle$ such that $\Pi_S |\psi\rangle = 0$ for all S or otherwise $\sum_S \langle \psi | \Pi_S | \psi \rangle \geq \epsilon$ for all $|\psi\rangle$, with $\epsilon > \frac{1}{\text{poly}(n)}$.

Problem: Output YES if the former and NO if the latter.

It was shown in [10] that this problem is QMA₁ complete for $k = 4$ and that for $k = 2$ there is an efficient algorithm. For $k = 3$ the problem was shown to be QMA₁-hard in [37].

4.2 The complexity of the k -local cohomology problem

We now prove the following Theorem:

Theorem 3. 4-LOCAL COHOMOLOGY is QMA₁-hard.

Proof. We will establish hardness by reduction from QUANTUM k -SAT. That is, our goal is to construct a cochain complex (V, d) such that a certain cohomology group $H^l(d)$ is nontrivial only if there exists a satisfying assignment of QUANTUM k -SAT. The general structure of the reduction is identical to that used in the proof of Theorem 2.²⁵

²⁵In fact, one could establish hardness of the cohomology problem by choosing B there to be Kitaev's clock Hamiltonian [54] and set the perfect completeness parameter $c = 1$. However, this would lead to a 6-local supercharge and thus the hardness result is strengthened by reduction from QUANTUM k -SAT.

Consider an instance of QUANTUM k -SAT, given by a set of k -local projectors Π_S acting on an n -qubit state $|\psi\rangle$. These have the form

$$\Pi_S = \Pi_S(\sigma_i^+, \sigma_i^-, \sigma_i^z), \quad |\psi\rangle = \sum c_{n_1 \dots n_n} |n_1, \dots, n_n\rangle, \quad (4.26)$$

where the arguments in Π_S denote the qubits $i \in S \subset \{1, \dots, n\}$ that are acted upon by each term and the $c_{n_1 \dots n_n}$ are arbitrary coefficients. Based on this data and applying the map (3.3), (3.4) we construct

$$\mathcal{P}_S := \Pi_S(a_i^\dagger b_i, b_i^\dagger a_i, a_i^\dagger a_i - b_i^\dagger b_i), \quad |\psi_n\rangle := \sum c_{n_1 \dots n_n} |\psi_{n_1 \dots n_n}\rangle. \quad (4.27)$$

Note that the $\mathcal{P}_S$ are k -local operators with $F = 0$ and $|\psi_n\rangle$ is a state in fermionic Fock space with fermion number $F = n$. We introduce an additional auxiliary site $i = 0$ with corresponding fermionic modes (a_0, b_0) . Then, we construct the coboundary operator

$$d = \frac{1}{\sqrt{2}}(a_0^\dagger + b_0^\dagger) \sum_S \mathcal{P}_S. \quad (4.28)$$

Note that the fermionic modes at the site $i = 0$ appear only outside $\mathcal{P}_S$; l th is to ensure that $d^2 = 0$ for any $\mathcal{P}_S$ and that d has fermion number $F = 1$.

In principle, the operator (4.28) acts on the full fermionic Fock space $(\mathbb{C}^2)^{\otimes m}$, $m = 2n$, spanned by the basis states $|n_{a_0}, n_{b_0}, \dots, n_{a_n}, n_{b_n}\rangle$, where $n_{a_i} = \{0, 1\}$ and $n_{b_i} = \{0, 1\}$ are the occupancy numbers of the two types of fermionic modes at site i . However, we can consistently restrict the action of d to a vector space $V \subseteq (\mathbb{C}^2)^{\otimes m}$, which we take to be the subspace spanned by the basis states satisfying the n constraints:

$$n_{a_i} n_{b_i} = 0, \quad i = (1, \dots, n). \quad (4.29)$$

Note that there are no constraints on the occupancy numbers n_{a_0} and n_{b_0} at the site $i = 0$. For the remaining n sites the constraints above impose that no more than a single mode is occupied. The vector space then has the form $V = \bigoplus_{p=0}^{n+2} V^p$, where each V^p is spanned by the set of states satisfying (4.29) and fermion number p . Note that the coboundary operator (4.28) creates states at the (unconstrained) site $i = 0$, while at the sites $i = (1, \dots, n)$ it exchanges a -modes by b -modes preserving the condition (4.29). Furthermore, since d has $F = 1$ it acts as $d : V^p \rightarrow V^{p+1}$ and this defines a cochain complex of the form

$$C : \quad 0 \xrightarrow{d} V^0 \xrightarrow{d} \dots \xrightarrow{d} V^{n+2} \xrightarrow{d} 0. \quad (4.30)$$

Having specified the cochain complex (V, d) it remains only to set the integer l specifying the particular cohomology group H^l of interest. We set $l = n + 2$, the highest possible nontrivial value in the complex (4.30). That is, we need to search for cochains in V^{n+2} which are cocycles but not coboundaries. Since the sites $i = (1, \dots, n)$ can be occupied by at most one fermionic mode, any element $|\Psi_{n+2}\rangle \in V^{n+2}$ takes the form

$$|\Psi_{n+2}\rangle = |1, 1\rangle \otimes |\psi_n\rangle, \quad (4.31)$$

where $|\psi_n\rangle$ is an n -fermion state of the form in (4.27).

Note that applying the operator d in (4.28) to a state of the form (4.31), we automatically have

$$d|\Psi_{n+2}\rangle = 0, \quad (4.32)$$

consistent with the form of the complex (4.30). That is, all $(n+2)$ -cochains are automatically $(n+2)$ -cocycles. On the other hand, applying $d^\dagger$ on such states we have

$$d^\dagger|\Psi_{n+2}\rangle = \frac{1}{\sqrt{2}}(|0,1\rangle + |1,0\rangle) \otimes \sum_S \mathcal{P}_S |\psi_n\rangle. \quad (4.33)$$

Now, given a YES instance of QUANTUM k -SAT, there exists some n -qubit state $|\psi\rangle$ such that $\Pi_S |\psi\rangle = 0$ for all S . The corresponding state $|\psi_n\rangle$ in fermionic Fock space obtained by (4.27) then satisfies $\mathcal{P}_S |\psi_n\rangle = 0$. Thus,

$$\Pi_S |\psi\rangle = 0 \quad \Rightarrow \quad d^\dagger |\Psi_{n+2}\rangle = 0. \quad (4.34)$$

This, together with (4.32) implies $H^{n+2} \neq 0$. Thus, for every YES instance of QUANTUM k -SAT there is YES instance of $(k+1)$ -LOCAL COHOMOLOGY and vice versa.

To complete the reduction it remains only to check the case of a NO instance. That is, we evaluate the norm

$$\begin{aligned} \|(d \pm d^\dagger)\|_{V^{n+2}} &= \sqrt{\langle \Psi_{n+2} | (d \pm d^\dagger)^2 | \Psi_{n+2} \rangle} = \left(\langle \psi_n | \left(\sum_S \mathcal{P}_S \right)^2 | \psi_n \rangle \right)^{1/2} \\ &\geq \langle \psi_n | \sum_S \mathcal{P}_S | \psi_n \rangle = \sum_S \langle \psi | \Pi_S | \psi \rangle, \end{aligned} \quad (4.35)$$

where in the penultimate step we used the Cauchy–Schwarz inequality and in the last step we wrote the expression in qubit space, using (3.10). In a NO instance of QUANTUM k -SAT, the last expression above is bounded below by ϵ , and we have

$$\|d \pm d^\dagger\|_{V^{n+2}} \geq \epsilon. \quad (4.36)$$

Thus, a NO instance of QUANTUM k -SAT implies a NO instance of k -LOCAL COHOMOLOGY. Since QUANTUM 3-SAT is QMA₁-hard [37] then 4-LOCAL COHOMOLOGY is QMA₁-hard. $\square$

One can also see that the problem is contained in QMA, as we show next.

Theorem 4. *k -LOCAL COHOMOLOGY is contained in QMA for cochain complexes $(V[m], d[k])$.*

Proof. Recall that a YES instance of k -LOCAL COHOMOLOGY is one where there exists some state $|\Omega^l\rangle \in V^l$ such that $d|\Omega^l\rangle = d^\dagger|\Omega^l\rangle = 0$, and in a NO instance, we have that for any state $|\psi^l\rangle \in V^l$, $\|(d \pm d^\dagger)|\psi^l\rangle\| \geq \epsilon$, with $\epsilon = 1/\text{poly}(m)$. Thus for a YES instance the witness is simply the state $|w^l\rangle = |\Omega^l\rangle$. To verify the witness, it suffices to estimate $\langle w^l | \Delta^l | w^l \rangle = \langle w^l | \Delta | w^l \rangle$ up to an inverse polynomial additive accuracy smaller than ϵ . This can easily be done by running phase estimation

on an approximation of $e^{i\Delta}$ obtained from Hamiltonian simulation, and providing $|w^l\rangle$ as the input. Since d is k -local, Δ is $2k$ -local and the local terms can be computed from d efficiently. One can use, for instance, the Trotterization method of Hamiltonian simulation to implement a unitary U such that $\|U - e^{it\Delta}\| \leq 1/\text{poly}(m)$ in time $\text{poly}(m, t)$ [60]. Hence, phase estimation can be used to obtain an estimate of $\langle w|\Delta|w\rangle$ up to additive accuracy $1/\text{poly}(m)$, with probability at least $2/3$, in polynomial time. Hence, the verifier circuit will measure an energy inverse-polynomially close to zero and accept $|w\rangle$ with probability at least $2/3$, as required.

In a NO instance, we are promised that any candidate witness $|\psi^l\rangle \in V^l$ will have $\langle \psi^l|\Delta|\psi^l\rangle \geq \epsilon^2$. However, we are not provided any promises on states outside of V^l and it is possible that we could be “tricked” into accepting a witness $|w\rangle$ with support outside V^l , and so we must first project the witness onto the space V^l . This is straightforward: we compute the Hamming weight of the state $|w\rangle$ in the witness register, and store the result in an ancilla register of $\log m$ qubits. We then proceed to measure this register and if we see any outcome $\neq l$, reject. Otherwise, we know that the *collapsed* state in the witness register is supported entirely in V^l , and we proceed as above, rejecting if we obtain an energy larger than some threshold $\epsilon^2 - 1/\text{poly}(m)$. By the problem promise and the behaviour of phase estimation, we will therefore reject with probability $\geq 2/3$ as required. $\square$

Therefore we find that, in general, the complexity of k -LOCAL COHOMOLOGY is somewhere in between QMA_1 and QMA . The question of whether $\text{QMA}_1 = \text{QMA}$ is an open problem in complexity theory [1].

An interesting question is whether the cohomology problem remains QMA_1 -hard for natural instances of the problem. As mentioned in the Introduction, a system of particular interest is the fermion hardcore model of [25], which gives rise to the well-known independence complex. A lower bound on the computational complexity of the cohomology problem for the independence complex follows from the results in [4]. In this reference, the chain complex was defined by the boundary operator ∂ and it was shown that deciding if the corresponding homology groups $H_l(\partial)$ were trivial or nontrivial is NP-hard, giving a lower bound on its complexity. This translates directly into a lower bound on the complexity of the cohomology problem since in this case the homology and cohomology groups are isomorphic, $H_l(\partial) \simeq H^l(d)$. Finally, we note that for the case of the independence complex, the Laplacian is always sparse, even though the coboundary operator is not always local [40]. Therefore, following the proof above, we can run Hamiltonian simulation directly on Δ without first obtaining a local decomposition, and hence the cohomology problem for the independence complex is inside QMA . Combining the two results, we find that the complexity of the problem lies somewhere between NP and QMA . Determining the precise complexity would be interesting, especially in relation to its applications to topological data analysis [61]. We discuss this in more detail in Section 6.

5 The Betti Number Problem

In this section, we consider the counting version of the cohomology problem. Namely, rather than deciding whether a certain cohomology group $H^l(d)$ is nontrivial, we consider the problem of determining

its rank, or Betti number $\beta_l = \text{rank } H^l(d)$.²⁶ We consider both the exact problem as well as that of finding an additive approximation to normalized Betti numbers, $\frac{\beta_l}{\dim V^l}$. It is known that computing Betti numbers can be a very hard problem. For instance, in the case of algebraic varieties the problem is PSPACE-hard [78] and in the case of the independence complex it is #P-hard (see below). To the best of our knowledge, the complexity of an additive approximation to normalized Betti numbers has not been directly considered.

As we show below, for a cochain complex $(V[m], d[O(1)])$, the problem of computing the l th Betti number is complete for the class #BQP (the counting version of QMA), which was shown to be equivalent to #P under weakly parsimonious reductions in [12]. Let us first point out that a weaker hardness result follows from [17], where it was shown that computing the Euler characteristic (or Witten index),

$$\mathcal{I} = \sum_{l=0}^m (-1)^l \beta_l, \quad (5.1)$$

of a complex is #P-complete, with the vector space and a 6-local coboundary operator given as the input. It thus follows that computing the individual Betti numbers themselves is #P-hard under a Turing reduction. This result is strengthened by the following Theorem:

Theorem 5. *Given a cochain complex $(V[m], d[6])$, the problem of computing the l th Betti number is #P-hard for $l = \Omega(n)$ and #P-complete provided the l th Laplacian Δ_l has a $1/\text{poly}(m)$ gap above its zero groundspace.*

Proof. In the proof of Theorem 3, we showed how to construct a cochain complex from an arbitrary n -qubit, k -local Hamiltonian such that the Betti number β_{n+2} is equal to the number of 0-eigenvalues of that Hamiltonian.

In [12] the problem LH(σ) of computing the number of states in a low-energy groundspace of energies $0 \leq E \leq \sigma$ of a local n qubit Hamiltonian is considered, and shown to be complete for the class #BQP for $e^{-\text{poly}(n)} \leq \sigma \leq 1/\text{poly}(n)$. In our case we are interested in computing the dimension of the zero-eigenspace of a local Hamiltonian, i.e., the problem LH(0). This can be seen to be hard for the counting version of QMA₁ – which we will call #BQP₁ – using the same proof as [12] (in particular, considering the perfect completeness version of the proof of their Theorem 11). Note that by definition we have the inclusions #P $\subseteq$ #BQP₁ $\subseteq$ #BQP. Since #BQP = #P under weakly parsimonious reductions [12], it follows that LH(0) and hence the problem of computing β_{n+2} for these complexes is #BQP-hard under weakly parsimonious reductions, and in particular that computing Betti numbers is #P-hard.

The problem can be placed inside #BQP whenever the Laplacian Δ_l has at least an inverse-polynomial gap above 0. Indeed, since the Betti numbers β_l correspond to the dimension of the zero-eigenspace of the l th Laplacian, Δ_l , one can construct a BQP verifier that distinguishes zero from non-zero eigenvectors so long as Δ_l has at least an inverse-polynomial gap above 0. The problem of

²⁶Note that Betti numbers only count the rank of the freely acting part of the cohomology group. In the absence of torsion coefficients, as we assume here, the Betti number captures the rank of the full group. Thus, a vanishing Betti number implies a trivial cohomology (and dual homology) group.

computing Betti numbers then reduces to the problem of counting the number of inputs that cause the BQP verifier to accept with high probability, which is inside $\#BQP$ and by the results of [12] is inside $\#P$. In the construction used in the proof of Theorem 3 mentioned above, we do indeed have an inverse-polynomial spectral gap above 0 (this can be seen from Lemma 12 of [12] plus the perfect completeness modification mentioned above), and hence it follows that the Betti number problem is $\#P$ -complete. $\square$

We now focus on the problem of additively approximating Betti numbers. More precisely, we focus our attention on the relatively easier problem of estimating a *normalized* version of Betti numbers. Concretely, we define the following problem:

Definition 6 (BETTI NUMBER ESTIMATION (BNE))

Input: A cochain complex $(V[m], d[k])$, an integer $0 \leq l \leq m$, and an accuracy parameter $\epsilon = \frac{1}{O(\text{poly}(m))}$ and success probability $\mu > 1/2$

Problem: Output an approximation χ s.t. $\left| \chi - \frac{\beta_l}{\dim V^l} \right| \leq \epsilon$ with probability $\geq \mu$.

Intuitively, BNE should be much easier to solve than the cohomology problem considered in Section 4. Indeed, we describe a BQP algorithm for a general family of complexes in Section 5.2, showing that for a large class of complexes BNE is tractable for quantum computers. In contrast, no efficient classical algorithms are known for estimating normalized Betti numbers (even for “easier” complexes like the clique complex), and indeed it seems reasonable to expect that the problem should remain hard for a large family of complexes.

Pinpointing the precise complexity of BNE, however, is a bit subtle. As we have discussed, Betti numbers correspond to the dimension of the zero-eigenspace of the Laplacian operator Δ , and this creates difficulties from “both directions” when attempting to prove completeness results. First, distinguishing zero-eigenvalues from non-zero eigenvalues is in general a hard task for quantum computers in the absence of any promise on a spectral gap above 0 (as we had for the problems considered in Section 4), which can make it difficult to design efficient quantum algorithms. Second, any constructions used in hardness proofs will need to use information contained only in the number of zero-eigenvectors of a Laplacian, which can, at least using our techniques, make it difficult to prove hardness results for complexity classes that don’t have a notion of perfect completeness (it is for this reason that we could only prove QMA_1 -hardness in Section 4).

To allow for a more complete characterisation of the problem, we instead consider a relaxed version corresponding to counting the number of small eigenvalues (below an inverse-polynomial threshold) of Δ_l , and argue below that this is in fact a natural problem. To capture the notion that this problem is very similar in nature to estimating Betti numbers, and in fact coincides with it when we are promised that Δ_l has a spectral gap above its zero groundspace, we call this problem the *quasi-Betti number estimation* (QBNE) problem. This problem was also considered for the case of clique complexes in [40],

where it was called the *Approximate Betti number estimation* (ABNE) problem, although its hardness was not considered there. We therefore use the name QBNE to distinguish this more general version of the problem from the restriction to clique complexes, where it is called ABNE.

First we define the notion of spectral density *within a subspace*. Let $A|_V$ be the restriction of a matrix A to its action on the space V , then for a positive semi-definite Hermitian matrix $A \in \mathbb{C}^{2^n \times 2^n}$ and a threshold $b \in \mathbb{R}^+$, we define the *spectral density within a subspace* of A with respect to threshold b and space V as

$$N_A(b, V) := \frac{1}{\dim(V)} \sum_{j: \lambda_j(A|_V) \leq b} 1 \quad (5.2)$$

where $\lambda_j(A|_V)$ denotes the j th eigenvalue of $A|_V$. When the space V is taken to be the full 2^n -dimensional complex vector space, this quantity is just the ordinary low-lying spectral density defined in [9], denoted by $N_A(b)$.

The central problem that we consider in this section is the following:

Definition 7 (QUASI-BETTI NUMBER ESTIMATION (QBNE))

Input: A cochain complex $(V[m], d[k])$, an integer $0 \leq l \leq m$, precision parameters $\delta, \epsilon = \Omega(1/\text{poly}(m))$, and a success probability $\mu > 1/2$.

Problem: Output an approximation χ that satisfies, with probability at least μ ,

$$N_\Delta(b, V^l) - \epsilon \leq \chi \leq N_\Delta(b + \delta, V^l) + \epsilon, \quad (5.3)$$

where Δ is the Laplacian of the complex.

This problem is the generalization of the LLSD problem [40] from arbitrary Hermitian matrices to Laplacians of cochain complexes. For completeness, and because we make use of it later on, we define the LLSD problem here.

Definition 8 (Low-lying spectral density (LLSD) [40])

Input: A positive-semidefinite Hermitian matrix $A \in \mathbb{C}^{2^n \times 2^n}$, a real number $b = \Omega(1/\text{poly}(n))$, precision parameters $\delta, \epsilon = \Omega(1/\text{poly}(n))$, and a success probability $\mu > 1/2$

Problem: Output an approximation χ that satisfies, with probability at least μ ,

$$N_A(b) - \epsilon \leq \chi \leq N_A(b + \delta) + \epsilon. \quad (5.4)$$

We argue that the fraction of small eigenvalues of the l th Laplacian Δ_l is in fact a natural quantity to consider. For the ordinary graph Laplacian (Δ_0 of the graph's simplicial complex) it is well known that the number of 0-eigenvalues (i.e. β_0) gives the number of connected components of the graph. Cheeger's inequality and its variants provide an approximate version of this fact for the case of two components: they state that a graph has a sparse cut if and only if there are at least two eigenvalues that are close to zero [6]. Intuitively, this says that if the graph Laplacian has an “almost-zero” eigenvalue, then only a few

edges with small total weight need to be removed from the graph to cut it into two connected components. It was long conjectured that an analogous characterization should hold for higher multiplicities: that there are k eigenvalues close to zero if and only if the vertex set can be partitioned into k subsets, each defining a sparse cut. Surprisingly, this conjecture was only confirmed quite recently, where a series of works culminated in the so-called higher-order Cheeger inequality [58]. These results have been used to give theoretical backing to the use of spectral clustering algorithms for optimally partitioning the vertices of a graph into well-connected clusters [70].

Currently it is not known whether such strong results also hold in general for higher-order Laplacians, but there are some recent encouraging results in this direction [69, 42, 81, 39], which have motivated work on spectral clustering for simplicial complexes [68]. With these results in mind, we suggest that the problem of estimating the number of small eigenvalues of higher-order Laplacians is a well-motivated one.

The quantum algorithms that we consider can in some instances be implemented on a non-universal model of quantum computer known as the one clean qubit model, whose associated complexity class is known as **DQC1** (See Appendix A for a description of the model and associated complexity class). This essentially rules out the possibility that the problem is **BQP**-hard, unless $\text{DQC1} = \text{BQP}$, but raises the question of whether the problem of estimating Betti numbers is a **DQC1**-hard problem. This makes intuitive sense in light of the fact that the cohomology problem is QMA_1 -hard. Indeed, it is often useful to view the class **DQC1** as an “averaged” version of the class **QMA**: in **QMA** we care about the existence of a single input to a quantum circuit that causes it to accept with high probability; in **DQC1** we care about the average acceptance probability of a quantum circuit over all possible inputs from some Hilbert space. This connection is made explicit by some of the complete problems for each class: for **QMA**, the canonical complete problem is that of finding the value of the smallest eigenvalue of a Hamiltonian; for **DQC1**, a complete problem is estimating the average of an (exponential) number of low-lying eigenvalues of a Hamiltonian. More generally, many **QMA**-complete problems are concerned with finding a specific state or energy, whilst many **DQC1**-complete problems focus on finding more “averaged” quantities such as (normalized) partition functions, or estimating the Jones polynomial of the trace closure of braids (the exact variant of which is $\#P$ -hard).

In the next section we prove that the **QBNE** problem is indeed **DQC1**-hard. One might also consider the hardness of the true **BNE** problem. It is in fact possible to show that it is hard for a “perfect completeness” version of **DQC1**; however, such a class is not a particularly natural one, and a hardness result for this class may not be of interest (indeed, it is not immediately clear that this class would not be classically simulable) and so we do not focus on that here.

5.1 QUASI-BETTI NUMBER ESTIMATION is DQC1-hard

Recall the definition of cochain complexes given in Section 4, and their input as a tuple $(V[m], d[k])$. As discussed there, for the Betti number estimation problem to be well defined one needs to be able to be given an efficient description of the input vector space, $V[m] = V = \bigoplus_{p=0}^m V^p$. This efficient description automatically allows one to efficiently *check* whether a given (basis) vector lies in the space V^p or not, since it involves checking only a polynomial (in m) number of constraints, each of which can be checked in polynomial time. For the purpose of describing quantum algorithms for Betti number estimation, however, we will also need to be able to efficiently *sample* elements from a space V^p . This won't in general be possible, and so this imposes an additional restriction on the instances of the problem that we consider. We define what we mean by an *efficiently sample-able vector space* below.

Definition 6 (Efficiently sample-able space). *Let $V = \bigoplus_{p=0}^m V^p \subseteq \mathbb{C}^{2^m}$ be a vector space described by a polynomially long (in m) list of constraints. We say that V is efficiently sample-able if there exists a polynomial-sized quantum circuit $\mathcal{C}$ that can produce an m -qubit state*

$$\rho_{V^p} := \frac{1}{\dim(V^p)} \sum_{v^p \in V^p} |v^p\rangle\langle v^p|, \quad (5.5)$$

starting from the all-zeros state on $O(\text{poly}(m))$ qubits, and where we take the v^p to be some orthonormal basis over V^p , labelled by m -bit strings. Additionally, we require that $\mathcal{C}$ itself must be constructed in polynomial time by a classical Turing Machine given only the list of constraints defining V . For instance, if V is a subspace of the fermionic Fock space of m fermions, then ρ_{V^p} can be taken to be the maximally mixed state over all fermionic states from V with fermion number p , encoded in qubits (e.g. via some fermion-to-qubit encoding).

The requirement that the quantum circuit be constructed by a polynomial-time classical Turing machine ensures that no hardness is hidden in the sampling of the vector space.²⁷ Note that it might also be sufficient to *approximate* ρ_{V^p} by, say, sampling only approximately uniformly from a basis for V^p . For some spaces this might be computationally easier than sampling exactly uniformly, however we won't consider this further.

We show that QBNE is DQC1-hard for log-local²⁸ coboundary operators acting on efficiently sample-able spaces $V[m]$ and $l = \Theta(m)$, by constructing a particular instance of a complex for which this problem is hard. Note that the problem can only get harder if we drop the restriction that $V[m]$ is efficiently sample-able, and so that harder version of the problem remains DQC1-hard. The constructions presented in Section 4 make the proof of this fact quite straightforward.

²⁷E.g. if the list of constraints was given as some arbitrary 3SAT formula, then sampling from the space would be equivalent to sampling from its set of solutions – something we don't expect to be able to do efficiently.

²⁸A log-local operator B acting on n fermionic modes / qubits is one of the form $B = \sum_{j=1}^m B_j$, where each B_j acts on at most $\log(n)$ fermionic modes / qubits, and $m = O(\text{poly}(n))$.

Theorem 6. QBNE is DQC1-hard for cochain complexes $(V[m], d[\log m])$, $l = \Theta(m)$, and remains hard when $V[m]$ is efficiently sample-able.

Proof. LLSD was already shown to be DQC1-hard for arbitrary log-local, positive semi-definite Hermitian operators in [40].²⁹ Our approach will be to define an instance of the QBNE problem, such that solving that instance allows one to solve a DQC1-hard instance of the LLSD problem, which then implies the theorem. To this end, let A_{DQC1} be any (n -qubit, log-local, positive semi-definite) Hermitian operator for which the LLSD problem is DQC1-hard with threshold b , precision parameters δ and ϵ , and success probability μ .

We use the same construction as in the proof of Theorem 3, using A_{DQC1} as our starting point. Recall that there the complex $(V[m], d[\log n])$ is of the form

$$C : \quad 0 \xrightarrow{d} V^0 \xrightarrow{d} \dots \xrightarrow{d} V^{n+2} \xrightarrow{d} 0,$$

where $V = \bigoplus_{p=0}^m V^p$ is a subspace of fermionic Fock space on $m = 2(n+1)$ fermionic modes. Our coboundary operator $d : V^p \rightarrow V^{p+1}$ will be a log-local fermionic operator of the form $d = \frac{1}{\sqrt{2}}(a_0^\dagger + b_0^\dagger)B$, where B is the fermionic operator obtained by applying the map (3.4) and (3.3) to A_{DQC1} . Recall that there the Laplacian takes the form

$$\Delta = dd^\dagger + d^\dagger d = B^2. \quad (5.6)$$

As noted in our proof of Theorem 3, the only states supported in V^{n+2} are states of the form $|1, 1\rangle \otimes |\psi_n\rangle$, for $|\psi_n\rangle$ an n -fermion state obtained via (3.3). Moreover, by (3.10), the operator B restricted to the space $\tilde{V}^n$ spanned by the states $|\psi_n\rangle$ coincides precisely with A_{DQC1} . Since B has fermion number 0 by construction, it necessarily preserves the subspace $\tilde{V}^n$ and so, combined with (3.10), we have $\langle \psi_{n+2} | \Delta | \psi_{n+2} \rangle = \langle \psi | A_{\text{DQC1}}^2 | \psi \rangle$ for all $|\psi\rangle \in V^n$, where $|\psi_{n+2}\rangle = |1, 1\rangle \otimes |\psi_n\rangle$. Putting everything together, this implies that the eigenvalues of $\Delta|_{V^{n+2}}$ are the squares of the eigenvalues of A_{DQC1} .

It follows that the spectral density of $\Delta|_{V^{n+2}}$ below a threshold t is the same as the spectral density of A_{DQC1} below threshold $\sqrt{t}$. Hence, to solve the LLSD problem on A_{DQC1} with parameters b, δ, ϵ, μ , it suffices to solve the QBNE problem for this complex $(V[m], d[\log])$ and $l = n+2 = m/2 + 1$, with parameters $b', \delta', \epsilon, \mu$ for $b' = b^2$ and $\delta' = 2b\delta + \delta^2$. Since the LLSD problem with input A_{DQC1} is DQC1-hard, this construction gives a DQC1-hard instance of the QBNE problem.

The only thing that remains is to verify that the space V is efficiently sample-able as per Definition 6. This is straightforward to see: the space V is a subspace of fermionic Fock space on $m = 2(n+1)$ modes subject to the constraints (4.29). To sample from the space V^p , we can classically sample from the set of Hamming weight p , m -bit strings of the form $n_{a_0}, n_{b_0}, \dots, n_{a_n}, n_{b_n}$ that satisfy those same constraints. This set is easy to sample from uniformly: choose p indices from $0, \dots, n$ uniformly at random (without replacement). If we chose index 0 then we set either $n_{a_0} = 1, n_{b_0} = 1$, or both, each with equal

²⁹Note that there is a claimed proof of this result in [9]. Unfortunately, the proof contains a (major) mistake. See [40] for further details on this point.

probability. If we set both occupancies to 1, we remove an index from our choice at random (to ensure we create a state with fermion number p). From each index i of our remaining indices, we set either $n_{a_i} = 1$ or $n_{b_i} = 1$ with equal probability.

Each m -bit string will correspond to a unique basis state from V^p , and so by sampling these uniformly at random and creating the corresponding state $|n_{a_0}, n_{b_0}, \dots, n_{a_n}, n_{b_n}\rangle$ (and “forgetting” which state we produced), the quantum state will be precisely

$$\rho_{V^p} = \frac{1}{\dim V^p} \sum_{v^p \in V^p} |v^p\rangle\langle v^p|, \quad (5.7)$$

which is the state we need to produce to efficiently sample from V^p . □

5.2 QUASI-BETTI NUMBER ESTIMATION is in BQP

In this section, we present a general algorithm for estimating Betti numbers of complexes defined by a coboundary operator d and a vector space V . To keep things simple, we assume that the vector spaces V are subspaces of an n -qubit Hilbert space, and that we are given access to boundary operators that are defined in qubit space. This is contrary to the previous sections, however if we are instead given access to boundary operators in *fermionic* space, we note that there exist fermion-to-qubit transformations that can be computed efficiently and give at most a logarithmic overhead in terms of locality of the underlying operators [11, 79], and so this assumption is without loss of generality.

Our algorithm is essentially a simplification of the Lloyd, Garnerone, and Zanardi (LGZ) algorithm of [61] that applies to log-local boundary operators, but generalised to apply to a wider class of complexes (i.e. not just the clique complex). If instead the boundary operators are sparse (but not local), then one can solve the QBNE problem only when either the Laplacian is guaranteed to have an inverse-polynomial spectral gap above zero, or if it is guaranteed to be sparse itself (and we can efficiently construct sparse access to it). We comment more on this at the end of the section.

Theorem 7. *QBNE is in BQP for cochain complexes $(V[n], d[O(\log n)])$, any l , and when $V[n]$ is efficiently sample-able.*

Proof. By its definition, solving the QBNE problem is the same as solving the LLSD problem for the operator $\Delta|_{V^l}$ for $\Delta = (d + d^\dagger)^2$ the Laplacian of the cochain complex. If d is log-local, then clearly Δ is also log-local, and hence sparse. Quantum algorithms for solving LLSD on sparse Hermitian matrices were presented in [40]; these algorithms involve sampling uniformly random eigenvalues of the target matrix by running phase estimation with an input of the maximally mixed state over the full n -qubit Hilbert space. Here, we are only interested in the eigenvalues of the matrix $\Delta|_{V^l}$, and so we can run the same algorithm but instead use the maximally mixed, n -qubit state

$$\rho_{V^l} = \frac{1}{\dim(V^l)} \sum_{v^l \in V^l} |v^l\rangle\langle v^l| \quad (5.8)$$

over the space V^l as input to the phase estimation routine. By definition, all eigenvectors of $\Delta|_{V^l}$ are fully supported in this space, and so the algorithm run in this way will sample uniformly random eigenvalues of $\Delta|_{V^l}$. Finally, we remark that this state can be constructed in polynomial time by a quantum circuit if the space V is efficiently sample-able in the sense of Definition 6.

For completeness (and because we will refer to parts of this algorithm later on), we briefly describe the steps of the quantum algorithm for LLSD of the matrix $\Delta|_{V^l}$. Informally, the steps are:

1. Use Hamiltonian simulation to produce a unitary $U \approx e^{i\Delta}$
2. Run phase estimation on U with input ρ_{V^l} . This produces random samples of (approximations of) the eigenvalues of $\Delta|_{V^l}$, since the maximally mixed state ρ_{V^l} can always be expressed as a uniform mixture over eigenvectors of $\Delta|_{V^l}$.
3. Repeat a polynomial number of times, keeping a count of how many eigenvalues are output that fall below the input threshold b .

From this we can produce an approximation satisfying Equation 5.3 in polynomial time via a Hoeffding bound. Of course, one must check that the errors arising from the Hamiltonian simulation and phase estimation routines can be bounded appropriately. For a complete description of the algorithm, as well as a full error analysis, we defer to [40]. $\square$

Combined with Theorem 6, it follows that the complexity of estimating Betti numbers of the family of complexes considered in this paper lies somewhere between DQC1 and BQP. Finally, we make the observation that for some instances of the problem, the complexity actually lies inside DQC1, which can be seen by making use of a DQC1 version of the quantum algorithm for LLSD from [40] as a subroutine. We prove the following in Appendix B.

Corollary 1. *QBNE is in DQC1 for cochain complexes $(V[n], d[O(\log n)])$, any l , and when the vector space $V = \bigoplus_{p=0}^n V^p$ satisfies $\frac{\dim(V^l)}{2^n} = \Omega\left(\frac{1}{\text{poly}(n)}\right)$ and membership in V^l can be tested using a workspace of at most $O(\log n)$ bits.*

The requirement that membership in the input Hilbert space can be tested using only a logarithmically many workspace bits/qubits is somewhat unsatisfactory, but seems to be necessary to fit the rejection sampling step of the algorithm into DQC1 (see the proof in Appendix B). In [3] Aaronson et al. define the class TQC (for “trace computing quantum polynomial time”) and its sampling variant, sampTQC. They find that $\text{TQC} = \text{DQC1}$, but that sampTQC lies somewhere between DQC1 and BQP. Informally, the class is defined as any problem that can be solved using a one-clean-qubit computer, with the additional power that one is allowed to learn, after the one-clean-qubit computation, what state from the maximally mixed part of the input was actually given to the circuit. Provided with this information, one can carry out some extra post-processing on the result of the one-clean-qubit computation. With this model, we could perform the rejection-sampling step inside the post-processing part of a sampTQC computation, meaning that the requirement that membership in the Hilbert space can be checked using only a small workspace can be dropped. This implies the following corollary.

Corollary 2. QBNE is in **sampTQC** for cochain complexes $(V[n], d[O(\log n)])$, any l , and when the vector space $V = \bigoplus_{p=0}^n$ satisfies $\frac{\dim(V^p)}{2^n} = \Omega\left(\frac{1}{\text{poly}(n)}\right)$.

This raises the intriguing question of whether Betti number estimation is also *complete* for the class **sampTQC**.

Finally, we comment on our earlier remark that one can only solve the QBNE problem for sparse (but not local) boundary operators when either the Laplacian Δ itself is guaranteed to be sparse, or is guaranteed to have an inverse-polynomial spectral gap above zero (in which case QBNE is equivalent to BNE).

If the Laplacian is itself sparse, then we can efficiently use Hamiltonian simulation to obtain an approximation of $e^{i\Delta}$, and the algorithm can proceed as usual.³⁰ If it is not sparse, then one has to take a different approach. In particular, it is possible to instead simulate the operator $B = (d + d^\dagger)$, which will be sparse as long as d is. Since the eigenvalues of Δ are the squares of the eigenvalues of B , one would hope that by inputting the maximally mixed state over V^p to phase estimation on the unitary e^{iB} , and estimating the fraction of eigenvalues falling below the threshold b , it would be possible to obtain an estimate of the low-lying spectral density of $\Delta|_{V^p}$, and hence solve the QBNE problem.

However this is not necessarily the case: in QBNE we are interested in the number of (orthogonal) eigenvectors $|\psi^p\rangle \in V^p$ of Δ such that $\langle \psi^p | \Delta | \psi^p \rangle \leq b$ for some threshold b . However, any such state $|\psi^p\rangle$ cannot be an eigenvector of B unless $\langle \psi^p | \Delta | \psi^p \rangle = 0$: acting on $|\psi^p\rangle$ with $B = (d + d^\dagger)$ will yield a state with support in both V^{p+1} and V^{p-1} , unless $d|\psi^p\rangle = 0$ and $d^\dagger|\psi^p\rangle = 0$, in which case $|\psi^p\rangle$ must be a zero-eigenvector of $\Delta|_{V^p}$. Hence, estimating the number of small non-zero eigenvalues of B within the space V^p tells us nothing about the number of small non-zero eigenvalues of Δ within that space, and hence this approach will not, in general, allow us to solve the QBNE problem.

In some cases however, it *will* work: when $\Delta|_{V^p}$ has a spectral gap of at least b above zero, then any eigenvalues below b must be zero-eigenvalues, and in this case the low-lying spectral density of $\Delta|_{V^p}$ is equivalent to the low-lying spectral density of $B|_{V^p}$, and this quantity just counts the fraction of zero-eigenvalues of Δ , i.e. solves the true BNE problem. However, we emphasize that without this promise the method of inputting the maximally mixed state over V^p to phase estimation on the unitary e^{iB} will not allow one to estimate the low-lying spectral density of $\Delta|_{V^p}$, because the non-zero eigenvectors of B and Δ do not coincide.

6 Betti Numbers in Topological Data Analysis

As discussed in the Introduction, the cohomology problem encompasses a large number of computational problems, including ones of practical interest. In this section, we focus on applications to topological data analysis (TDA), specifically to the interesting problem of extracting topological features of large

³⁰For the clique/independence complex, one can show that the Laplacian actually *is* sparse [40].




Figure 4: An example of the TDA pipeline: one starts with a point cloud (left), adds edges between vertices within a certain distance of each other, resulting in a graph G (middle right), from which the clique complex is constructed (right). In this example the complex is connected and has a single one-dimensional hole, thus $\beta_0 = \beta_1 = 1$ and all other Betti numbers vanish.

data sets. As we discuss, this amounts to determining the ground states of the fermion hard-core model reviewed in the Introduction.

Understanding the ground state properties of physical systems (very much including condensed matter systems) is conjectured to be one of, if not the, most promising use-cases of quantum computers [14, 63]. Various tools for doing so have been developed over the past few years and we propose that these could be re-purposed for TDA (more generally, for many (co)homology problems). As we argue, it is entirely possible that current state-of-the-art computational techniques for studying condensed matter systems can be applied to this problem. We first provide an overview of the problem and describe the precise Hamiltonian to be considered, comment on the conditions to obtain a local Hamiltonian, and then discuss implications for the LGZ algorithm in Section 6.2. Finally, in Section 6.3 we propose a new variational approach to TDA.

6.1 TDA and the fermionic hard-core model

Topological data analysis is an approach for analyzing data using tools from topology, whose initial motivation was to study the “shape of data,” for some mathematically rigorous notion of shape. In this setting one is given a large data set (collection of points in space) and the task is to extract topological features such as the number of l -dimensional holes (corresponding to Betti numbers β_l) that can then be used to classify and analyse it.³¹ The main tool is so-called “persistent homology” – studying the homology of simplicial complexes derived from point-cloud data in order to look for *persistent* holes in the data, or rather in the topological space that the data is sampled from. A first step towards

³¹See e.g. [35] for an introductory overview.

persistent homology is simply to compute homology, i.e., to compute Betti numbers of a particular simplicial complex, which in this context is usually the *Vietoris-Rips complex* [35].

More formally, consider a collection of points $\{x_i\}_{i=1}^n$ embedded in $\mathbb{R}^m$. From this collection of points, one constructs a graph G by associating data points to vertices, and adding an edge between two vertices i and j if they are close, $\|x_i - x_j\|_2 \leq \epsilon$, where ϵ is termed the *grouping scale*. One then constructs the clique complex of G , $C(G)$, a simplicial complex whose simplices are the cliques of the graph G – see Figure 4. A complex built in this way is called the Vietoris-Rips complex. Let $C_k(G) \subset \{0,1\}^n$ denote the set of k -cliques in G , where an element $|s_k\rangle \in C_k(G)$ is represented by a Hamming weight k , n -bit string whose non-zero entries specify the vertices that constitute the clique. As with any simplicial complex, the boundary operator ∂ acts as

$$\partial |s_k\rangle = \sum_{l=1}^k (-1)^l |\widehat{s_{k-1}}(l)\rangle, \quad (6.1)$$

where $|\widehat{s_{k-1}}\rangle \in C_{k-1}(G)$ is obtained by setting the l th 1 of $|s_k\rangle$ to 0. That is, the boundary operator maps k -cliques to an alternating sum of the $(k-1)$ -cliques that it contains (the signs in the sum are determined by choosing a fixed ordering over the simplices of the complex). The homology groups $H_l(\partial)$ are then spanned by l -dimensional holes in the complex at a given grouping scale ϵ . Holes that persist over various grouping scales are deemed to be persistent holes, and form the basis for persistent homology. Hence, in order to perform topological data analysis in this way, one needs to determine the homology groups $H_l(\partial)$ or Betti numbers $\beta_l = \text{rank } H_l(\partial)$.

As mentioned, there is a direct connection between TDA in the above sense and the properties of a particular (supersymmetric) many-body system. Indeed, as discussed in the Introduction, the supersymmetric ground states of the fermionic hard-core model (1.1) for a graph $G = (V, E)$ precisely captures the cohomology groups of the independence complex $I(G)$. Since we are interested in the clique complex, which is the dual of the independence complex, we consider the model defined on the complement graph $\bar{G}$ and define the coboundary operator

$$d = \sum_{i \in \bar{V}} a_i^\dagger P_i, \quad P_i = \prod_{(i,j) \in \bar{E}} (1 - \hat{n}_j), \quad (6.2)$$

where $\bar{V}$ and $\bar{E}$ denote the vertex and edge set of $\bar{G}$. We now note that the operator $d^\dagger$ acts on cliques $|s_k\rangle$ in the same way as the coboundary operator ∂ does in (6.1). We leave as an exercise to verify that the fermionic anti-commutation relations automatically impose the correct signs in the sum of (6.2), something that results from fixing an ordering over fermionic modes, which in turn defines an orientation of the simplices. On the other hand, in the absence of torsion coefficients (as is the case here) one has $H_l(d^\dagger) \cong H^l(d)$ and thus

$$H_l(\partial) \cong H^l(d). \quad (6.3)$$

Elements of cohomology are then zero-eigenstates of the Laplacian of the (co)chain complex, given by

$$\Delta = \sum_{(i,j) \in \bar{E}} P_i a_i^\dagger a_j P_j + \sum_{i \in \bar{V}} P_i, \quad (6.4)$$

which is simply the supersymmetric Hamiltonian for the fermion hard-core model on the graph $\bar{G}$ (c.f., Eq. (1.1)). Thus, the problem of finding l -dimensional holes in the clique complex amounts to counting supersymmetric ground states of this system with fermion number l .

In the case of highly structured graphs, the supersymmetric ground states of the fermion hard-core model can be studied analytically or numerically. For instance, on a square lattice with periodic boundary conditions, the supersymmetric (zero-energy) ground states of the system are found at fillings in the range $1/5 \leq f \leq 1/4$ [46] and are in one-to-one correspondence with configurations that arise from tiling the plane using four rhombuses. The ground states are also known for the so-called martini and octagon-square lattices [23]. For more general graphs, however, the cohomology of the independence complex is much harder to compute. For instance, for the triangular and hexagonal lattices, the ground state structure of the fermion hard-core model is not known, although it is understood that the ground states occur in a finite window of filling [48], and that there is extensive ground state entropy [82, 45]. For a general graph G , however, there is no systematic approach to determining supersymmetric ground states and indeed we expect the problem to be intractable.³²

6.2 Existing quantum algorithms for TDA

A quantum algorithm for TDA that estimates normalized Betti numbers of the clique complex of a graph (in fact it solves the QBNE problem from Definition 5 for clique complexes) was described by Lloyd, Garnerone, and Zanardi (LGZ) [61]. The algorithm is essentially the one outlined in Section 5.2, in which a maximally mixed state over cliques is used as input to phase estimation on a unitary obtained via Hamiltonian simulation of either the “Dirac operator” or Laplacian of the complex. Recent work by Gyurik et al. [40] considered the possibility of implementing the LGZ algorithm on near-term quantum devices, and suggested that only a modest number of error-corrected qubits (≈ 80) might be required to see quantum advantage (i.e. to run the algorithm on input sizes beyond the reach of the best known classical methods) if various optimistic conditions are met.

The computational bottlenecks of the algorithm are the complexities of the subroutines used to implement Hamiltonian simulation of either the combinatorial Laplacian or the Dirac operator $B = (\partial + \partial^\dagger)$ constructed from the boundary operator ∂ , as well as the cost of generating the maximally mixed state over all k -simplices of the complex. The latter is dealt with only by assuming that the graph is sufficiently “clique-dense” (which will likely mean that the graph itself is very dense), allowing either rejection sampling or Grover search to be used to create the maximally mixed state. The Hamiltonian simulation subroutines can be expensive because although the boundary operator and the Laplacian are known to be sparse, they do not have an obvious decomposition into local terms, meaning that one must make use of the more resource-heavy algorithms for sparse Hamiltonian simulation instead of “lightweight” techniques for simulating local Hamiltonians. For a detailed discussion of this point,

³²Some partial information about supersymmetric ground states such as the Witten index could be extracted more easily. The Witten index for these systems is $\#P$ -hard, but finding an additive approximation to the normalized Witten index is in BPP [17].

see [40].

Let us consider the Hamiltonian (6.4). Note that the projectors P_i are nontrivial only when there is no edge between vertex i and vertex j in the input graph G , and hence the Laplacian contains terms which act on up to a total of $(n - \delta)$ fermionic modes, with δ the *minimum* degree of the original graph G . The Laplacian is thus k -local, provided

$$\delta = n - k, \quad (6.5)$$

with k a constant, i.e., for very edge-dense graphs G . Thus, by switching our attention to the coboundary operator of the independence complex, the Laplacian of interest will have a decomposition into local terms provided (6.5) holds. Moreover, this local decomposition is naturally expressed in terms of fermionic operators, suggesting that it could be possible to make use of modern techniques from the quantum chemistry literature for simulating fermionic Hamiltonians to obtain hardware efficient implementations for the Hamiltonian simulation part of the algorithm.

Graphs that are very edge-dense are likely to be natural candidates for obtaining large quantum speedups using the LGZ algorithm. Since the algorithm is efficient only when the graph G is sufficiently l -clique dense, it is likely that the graph is also edge dense – i.e., all vertices have very high degree. Indeed, a sufficient condition for l -clique denseness is when all vertices have degree at least γn , for $\gamma > \frac{l-2}{2(l-1)}$ [74]. Hence, a class of graphs for which the LGZ algorithm could exhibit quantum advantage will be when the degree of every vertex is at least $n - k$ for constant k , in which case the complement of the graph will have maximum degree k , and so the Laplacian (6.4) will be k -local. Hamiltonian simulation of local Fermionic operators is likely to be a less resource intensive task on near-to-medium-term devices than simulation of general sparse operators. For example, one can take an efficient Fermion-to-qubit mapping (e.g. the Bravyi-Kitaev transform [11]) to obtain a log-local qubit operator $\hat{d}$ corresponding to the coboundary operator d . Hamiltonian simulation of the log-local qubit operator $\hat{B} = (\hat{d} + \hat{d}^\dagger)$, or of the Laplacian itself, can then be implemented via Trotterization using no extra ancilla qubits. If we have information about the underlying structure of the graph G , it is possible that this simulation could be made even more efficient [16]. Moreover, since we are interested only in the low-energy subspace of these operators, it is possible that some recent results (e.g., [75]) regarding Hamiltonian simulation within low-energy subspaces could be exploited to make the simulation even more efficient. We leave these questions open for future work.

6.3 Alternative quantum and quantum-inspired approaches to TDA

The quantum algorithm described in the previous section suffers from a couple of drawbacks. First, it is only efficient when the graph is clique dense, which will likely restrict the family of graphs to which it can be applied. Second, it computes only *normalized* Betti numbers, and it is not entirely clear that these quantities have any immediate practical application. In particular, being able to estimate normalized l th Betti numbers up to inverse-polynomial additive error only allows one to distinguish complexes with very many l -dimensional holes from those with very few l -dimensional holes, which again might limit the family of graphs to which this algorithm can usefully be applied.

The description of elements of clique/independence homology as ground states of the (supersymmetric) fermion hard-core model reviewed above, a nontrivial quantum mechanical system of interacting fermions, suggests that the problem of clique/independence homology is quantum mechanical in nature. This raises the prospect of repurposing standard techniques for studying ground states of quantum mechanical systems for TDA, such as VQE or [71, 62] or QAOA [22]. Similarly, there are a plethora of *non-quantum* computational tools for approximating ground states of quantum many-body systems, including quantum Monte Carlo and tensor networks [57], which might yield efficient *classical* algorithms for topological data analysis. One immediate observation however, in both cases, is that often these techniques are designed for systems on very structured graphs (e.g. 2D lattices), and hence they might struggle to perform well on more arbitrary input graphs. Understanding exactly when these classical and quantum techniques perform well is an interesting question, and could shed light on when (if at all) we might expect to obtain quantum advantage for TDA, which we leave for future work.

7 Discussion and Outlook

We have argued that the intersection of many-body physics, complexity theory, and supersymmetry is a natural arena to study the complexity of problems in homological algebra, in particular the (co)homology problem. Indeed, this perspective reveals that the problems of deciding whether a (co)homology group is nontrivial and of estimating its rank are intrinsically quantum mechanical and serve as natural hard problems for various quantum complexity classes. Precisely, we defined k -LOCAL COHOMOLOGY and showed that its complexity lies between QMA_1 and QMA . Similarly, the complexity of an additive approximation to (quasi) Betti numbers lies between DQC1 and BQP .

These complexity results may come as a surprise at first, as the computational problems we discuss are defined purely in the context of (co)homology and are not manifestly quantum mechanical. As we have observed, however, these *can* be seen as quantum mechanical when viewed as problems in *supersymmetric* quantum mechanics, and relate to properties of the ground states of supersymmetric Hamiltonians. Indeed, the complexity of k -LOCAL COHOMOLOGY ultimately follows from the fact that the k -local Hamiltonian problem for supersymmetric systems remains QMA -complete, as shown in Section 3.

A set of interesting questions arises when considering simpler instances of supersymmetric Hamiltonians. An obvious question, for instance, is whether QMA -hardness of the local Hamiltonian problem for systems with $\mathcal{N} = 2$ SUSY is maintained for $2 \leq k \leq 3$. One may also consider systems with higher degrees of supersymmetry, $\mathcal{N} > 2$, and study whether there is a complexity transition at a certain value of $\mathcal{N}$ for each k . It is often the case, at least in the context of supersymmetric QFTs, that beyond a value of $\mathcal{N}$ the systems become so constrained that the entire class of Hamiltonians is specified by a small set of parameters.³³ We note, however, that although systems with extended supersymmetry can be easily

³³An example is the case of 4d QFTs with sixteen (Poincaré) supercharges, which are specified by a choice of gauge group and one complex coupling constant. These and other maximally supersymmetric systems are also often integrable. Although integrability does not imply tractability in a computational sense (see, e.g., [2] for a counterexample) it would

constructed in continuum quantum mechanics and quantum field theory, much less is known about spin systems with extended supersymmetry (see however the $\mathcal{N} = 4$ systems of [77]). We also stress that although the Hamiltonian for which we have proven hardness is local, in the sense that it is given by a sum of terms acting at most on a constant number of sites, it is in general *geometrically* non-local. It would be interesting to study whether the SUSY local Hamiltonian problem remains hard for systems constrained by geometric locality or other physical considerations. All this applies also to the problem of Betti number estimation. Finally, it would be interesting to study the complexity of these problems in the case of the fermion hard-core model of [26], especially in relation to topological data analysis as discussed in Section 6. In particular, if the local Hamiltonian problem remains QMA-hard for this model, as it is known to be for similar models like Fermi-Hubbard [67], then this may imply QMA₁-hardness of the cohomology problem for the independence complex, and likely also DQC1-hardness for QBNE of clique complexes. The latter would imply that the LGZ algorithm cannot be dequantized unless DQC1 $\subseteq$ BPP.

More generally, we note that our hardness constructions give a general approach for mapping hard problems in Hamiltonian complexity into hard problems in algebraic homology. This suggests that the approach could be exploited to elucidate the complexity of problems beyond those considered in this paper. In fact, this suggests defining new problems in computational topology altogether, which are natural from a quantum Hamiltonian complexity perspective. For instance, one may consider the standard QCMA-complete problem of deciding whether a local Hamiltonian has an “efficiently describable” ground state, the QMA(2)-complete problem of deciding if there is a ground state with a particular tensor-product structure, or finding the ground energy of stoquastic Hamiltonians. It would be interesting to determine whether these additional constraints correspond to any natural constraints or promises on the corresponding problems in homological algebra.

On the more practical side, as we discussed in Section 6, viewing the problem of topological data analysis as the problem of finding the ground states of a (supersymmetric) Hamiltonian suggests new approaches to this problem. In particular, this opens the door to using modern techniques from quantum chemistry for topological data analysis and cohomology problems more generally, as long as these can be formulated in terms of a fermionic coboundary operator acting on Fock space. Before the arrival of fault-tolerant quantum computers, it may also be interesting to consider the application of *classical* techniques (e.g. traditional tools for studying condensed-matter systems, or even machine learning techniques [44]) for studying ground states of many-body systems to topological data analysis or other (co)homological problems.

Acknowledgements

We are grateful to Ronald de Wolf and Kareljan Schoutens for helpful discussions and feedback on an earlier version of this paper, and Vedran Dunjko, Casper Gyurik, Ido Niesen, and Ian Marshall for useful

be interesting to consider the computational hardness of these systems.

discussions. CC was supported by QuantERA project QuantAlgo 680-91-034. PMC is supported by the European Union's Horizon 2020 Research Council grant 724659 Massive-Cosmo ERC-2016-COG and the STFC grant ST/T000791/1.

A Complexity Classes

In this section, we define the complexity classes QMA , QMA_1 , and DQC1 , which play a central role in our work.

We begin by defining the quantum complexity class QMA . Loosely speaking, this is the class of problems whose solution can be verified in polynomial time by a quantum computer, a quantum analog of the class NP . More formally, the quantum complexity class $\text{QMA}[c, s]$ is defined as follows:

Definition 7 ($\text{QMA}[c, s]$). *A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in $\text{QMA}[c, s]$ if there exists a polynomial-time quantum circuit (the “verifier”) V_x for any n -bit input string x such that*

- *If $x \in A_{\text{yes}}$, there exists a $\text{poly}(n)$ -qubit witness state $|w\rangle$ such that $\Pr[V_x(|w\rangle) = 1] \geq c$,*
- *If $x \in A_{\text{no}}$, then for any $\text{poly}(n)$ -qubit witness state $|w\rangle$, $\Pr[V_x(|w\rangle) = 1] \leq s$.*

That is, for a yes-instance, there must exist a quantum witness that causes the verifier to accept with probability $\geq c$. Otherwise, the verifier must reject all witnesses with probability $\geq 1 - s$.

The parameters c and s are known as the completeness and soundness parameters, respectively. The class $\text{QMA}[2/3, 1/3]$ is simply denoted by QMA and is viewed as the quantum analogue of the class NP . The choice of $2/3$ and $1/3$ is arbitrary; in fact we have that $\text{QMA} = \text{QMA}[1 - 1/\exp(n), 1/\exp(n)] = \text{QMA}[c, c - 1/\text{poly}(n)]$. A class that will be relevant for us is the “perfect completeness” variant of QMA , $\text{QMA}[1, 1 - 1/\text{poly}(n)] = \text{QMA}_1$. Whether $\text{QMA}_1 = \text{QMA}$ is a longstanding open question in quantum complexity theory.

A comment on perfect completeness. The notion of a perfect completeness complexity class introduces some nuances. In particular, the notion of a “universal gate set” becomes tricky, and one must fix a gate set in advance and require that any quantum circuits (i.e. the verifier circuit in the case of QMA_1) be constructed from this gate set. It is not known whether the definition of QMA_1 can be made independent of the choice of gate set. To illustrate this point, consider a problem contained in QMA_1 with respect to some choice of universal gate set $\mathcal{G}$, meaning that the verifier circuit, constructed out of gates from $\mathcal{G}$, accepts any valid witness with probability 1. Now consider moving to another universal gate set $\mathcal{G}'$: of course the verifier circuit can still be constructed out of gates from $\mathcal{G}'$, but only *approximately*, even though the approximation can be exponentially close to the target circuit. For perfect completeness, however, this is not enough: we require that the circuit built using $\mathcal{G}'$ also accepts valid witnesses with certainty, and not with “merely” a probability exponentially close to 1. Hence, the problem would be considered to be in QMA_1 with respect to $\mathcal{G}$, but *not* in QMA_1 with respect to $\mathcal{G}'$. In [37], Gosset and Nagaj use a standard universal gate set $\mathcal{G} = \{\hat{H}, T, \text{CNOT}\}$, and for our purposes we choose this gate set also for our definition of QMA_1 . For a detailed discussion of this point, we refer the reader to [37].

Next we define the one-clean-qubit model, and its associated complexity class **DQC1**. Here, we are given access to a *fixed* input

$$\rho = |0\rangle\langle 0| \otimes \frac{I_n}{2^n} \quad (\text{A.1})$$

consisting of a single qubit in a pure state $|0\rangle$, and n qubits in a maximally mixed state. We are allowed to apply a polynomial (in n) sized circuit, without intermediate measurements, to this state, and then we can measure the clean qubit. We are allowed to perform this procedure a polynomial number of times, and the probability of obtaining a 0 on the output is considered the output of the model. The associated complexity class is called **DQC1** and is defined as follows.

Definition 8 (**DQC1**). *A language $L = (L_{yes}L_{no})$ is in **DQC1** if, for every r -bit input $x \in L$, there exists a quantum circuit U_x acting on $n = \text{poly}(r)$ qubits and consisting of $O(\text{poly}(n))$ quantum gates, such that*

- If $x \in L_{yes}$,

$$p_{yes} = \text{Tr} \left[\left(U_x \left(|0\rangle\langle 0| \otimes \frac{I}{2^{n-1}} \right) U_x^\dagger \right) |1\rangle\langle 1| \otimes I \right] \geq a$$

- If $x \in L_{no}$,

$$p_{yes} = \text{Tr} \left[\left(U_x \left(|0\rangle\langle 0| \otimes \frac{I}{2^{n-1}} \right) U_x^\dagger \right) |1\rangle\langle 1| \otimes I \right] \leq b$$

for $a - b \geq 1/O(\text{poly}(n))$.

The promise gap between a and b means that, to solve any problem in **DQC1**, it is enough to estimate the “average acceptance probability” of U_x , p_{yes} , up to inverse-polynomial additive error.

B DQC1 Algorithm for Estimating Quasi Betti Numbers

Here we provide a proof of Corollary 1.

Proof. Observe that the condition $\frac{\dim(V^l)}{2^n} = \Omega\left(\frac{1}{\text{poly}(n)}\right)$ immediately implies that V^l is efficiently sampleable, and so the set of instances satisfying this constraint is a subset of those in **BQP**: if the space makes up a polynomially large fraction of the entire n -qubit Hilbert space, one can use rejection sampling to sample uniformly from V^l by repeatedly generating random n -bit strings until finding one that is contained in the set S^l of strings that give the computational basis states that span V^l . Each string has at least a $1/\text{poly}(n)$ probability of being in S^l , and hence this rejection sampling takes polynomial time (and succeeds with probability arbitrarily close to 1). We will show that in the case that membership in V^l can be tested using only logarithmically many ancilla (qu)bits, this rejection sampling step can effectively be performed inside **DQC1**.

First, we note that if the co-boundary operator d is log-local, then the (Hermitian) operator $A = (d + d^\dagger)^2$ that we use for Hamiltonian simulation will also be log-local. Hamiltonian simulation of log-local Hamiltonians can be performed in polynomial time in **DQC1** using the vanilla simulation technique of Trotterization ([13], Section 3.4). Hence, as pointed out in [40], each step of the algorithm

from Theorem 7 can be performed in DQC1, since the Hamiltonian simulation step requires no clean ancilla qubits, and the phase estimation routine requires at most $O(\log n)$ of them. The same error analysis applies in this case as in the proof of Theorem 7, and so for the rest of the proof we will assume that each part of the quantum algorithm works perfectly, and instead focus on the problem of doing rejection sampling within the one-clean-qubit model. We will refer to the circuit that performs steps 1 and 2 of the algorithm sketched in the proof of Theorem 7 (i.e. phase estimation on (an approximation of) e^{iA}) as C , and assume that it works perfectly – i.e. for an eigenstate $|\psi\rangle$ of A , it outputs the corresponding eigenvalue λ perfectly and with certainty.

In DQC1, we have at our disposal $O(\log n)$ clean qubits (i.e. in the state $|0\rangle$), and a maximally mixed register of n qubits in the state

$$\rho = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} |x\rangle\langle x|.$$

The algorithm of Theorem 7 requires a maximally mixed state over the space V^l , and so we need a way to emulate that input within DQC1 given access to only the full maximally mixed state ρ . Provided there exists a (classical) circuit to check whether a given bit-string is contained in S^l that uses no more than $O(\log n)$ workspace bits, we can perform this check within the one clean qubit model itself. This allows us to build the rejection sampling step into the DQC1 algorithm, in one of two ways.

If we are allowed to measure 2 qubits at the end of the computation, rather than the usual single qubit, then we can perform the rejection sampling as a post-selection step at the end of the computation. Such a model was considered in, e.g., [65], although this model does not technically give rise to the same complexity class DQC1. Alternatively, we could estimate the quantity $\frac{\dim(V^l)}{2^n}$ classically (i.e. in BPP), and estimate the quantity $\frac{\tilde{\beta}_l}{2^n}$ using a DQC1 algorithm, where $\tilde{\beta}_l$ is used to denote the number of eigenvalues of A_{V^l} below the given threshold. By dividing the latter by the former, and ensuring that the estimation accuracies are chosen appropriately, we can obtain an estimate of $\frac{\tilde{\beta}_l}{\dim(V^l)}$, which will give us a solution to the NORMALIZED QUASI-BETTI NUMBER ESTIMATION problem after errors from the circuit are taken into account. The class DQC1 is defined³⁴ to be the set of problems solvable in polynomial time, given access to a one-clean-qubit quantum computer plus polynomial time classical post-processing, and hence this would be a true DQC1 algorithm.

We begin by considering the first approach. As described in the proof of Theorem 6, if we were to input the state $|0\rangle\langle 0|_A \otimes \rho_{V^l}$ to C , where the first register consists of the $O(\log n)$ ancilla qubits used by the circuit, we would obtain an approximation of a uniformly random eigenvalue of A_{V^l} in the first register. We can append an extra ancilla qubit and use this as a flag that we set to $|1\rangle$ when the

³⁴The accepted definition of DQC1 is quite hard to pin down. We take it to be the set of problems solvable in polynomial time given access to a one-clean-qubit quantum computer (which we can run non-adaptively) plus polynomial-time probabilistic classical computation, which feels quite natural – after all, one must have access to *some* form of classical computation in order to, say, post-process the results of a polynomial number of measurement outcomes in order to output a final answer to the decision problem. This is the view taken also by Brandão in [9]

eigenvalue in the first register is $\leq b$, and $|0\rangle$ otherwise. Then the probability of measuring a 1 on this qubit at the end of the computation will be equal to the fraction of eigenvalues of $A_{|V^l}$ that are $\leq b$. However, since DQC1 requires us to input the maximally mixed state ρ over the full n -qubit Hilbert space we will also obtain estimates of eigenvalues of the other parts of A . To remedy this, we can attach another ancilla qubit and use it as a flag that we set to $|1\rangle$ at the beginning of the computation only when the state on the maximally mixed register is a computational basis state from V^l (here is where it is important that this checking only requires $O(\log n)$ many “work” qubits, which we must provide as yet more clean ancilla qubits). Finally, if we can measure the two flag qubits at the end of the computation, then by keeping a count of the number of times we see a state from V^l , and the number of times we see an eigenvalue $\leq b$, we can estimate the fraction of eigenvalues smaller than b out of all eigenvalues of $A_{|V^l}$ using a Hoeffding bound.

Now we consider the second approach. We proceed in exactly the same way as above, except we add a single extra ancilla qubit in which we store the logical AND of the two flag qubits. We then measure this qubit at the end of the computation as permitted by the canonical definition of DQC1. The circuit for this is shown in Figure 5, where we call the unitary for the entire circuit V .




Figure 5: DQC1 circuit for quasi Betti number estimation. The gate “ $\in V^l$?” is a gate that checks (using at most $O(\log n)$ ancillas) whether a (computational basis) state is contained in V^l or not, QPE is quantum phase estimation (for some choice of unitary), and the “ $\leq b$?” gate checks whether the value stored in the phase register is smaller than b or not.

Collecting the clean qubits into a single register A , the probability of measuring a 1 on the first qubit is

$$\Pr[\text{measure } 1] = \text{Tr} \left[V \left(|0\rangle\langle 0|_A \otimes \frac{I}{2^n} \right) V^\dagger (|1\rangle\langle 1| \otimes I) \right]. \quad (\text{B.1})$$

Since $(d + d^\dagger)^2$ is block-diagonal with each block acting on a separate V^l , every one of its eigenvectors,

and therefore those of A , are either entirely supported within some V^l , or else not supported at all. By expanding the equation above and taking the maximally mixed state to be a uniform mixture over eigenvectors of A (plus any extra orthogonal vectors required to complete a basis for $\mathbb{C}^{2^n}$), we can see that the probability of measuring a 1 at the end of the circuit is just the probability that some input state $|\psi\rangle$ chosen from the ensemble is contained in V^l and that this state is an eigenstate of A with eigenvalue $\leq b$. Hence, the probability of measuring 1 at the end of the computation, taken over all such states $|\psi\rangle$, is

$$\Pr_{|\psi\rangle}[\text{measure 1}] = \Pr_{|\psi\rangle}[|\psi\rangle \in V^l] \cdot \Pr_{|\psi\rangle}[\langle\psi|A|\psi\rangle \leq b | |\psi\rangle \in V^l] \quad (\text{B.2})$$

$$= \frac{\dim(V^l)}{2^n} \cdot \frac{\tilde{\beta}_l}{\dim(V^l)} = \frac{\tilde{\beta}_l}{2^n}. \quad (\text{B.3})$$

By repeating the computation a polynomial number of times, we can obtain an estimate of this probability up to additive accuracy $1/\text{poly}(n)$. Hence, we can obtain a quantity

$$\frac{\tilde{\beta}_l}{2^n} \pm \frac{1}{p(n)}, \quad (\text{B.4})$$

for $p(n)$ some polynomial in n .

Next, we can classically estimate the fraction of all n -qubit (computational basis) states that lie inside V^l – i.e. estimate the relative dimension of the space V^l . To do this we can sample n -bit strings uniformly at random and check whether they satisfy the description of V^l provided as input. Once again, in polynomial time and by a Hoeffding bound we can obtain an estimate

$$\frac{\dim(V^l)}{2^n} \pm \frac{1}{r(n)} \quad (\text{B.5})$$

for some polynomial $r(n)$. Finally, dividing the quantity in B.4 by the one in B.5, we obtain an estimate

$$\begin{aligned} \left(\frac{\tilde{\beta}_l}{2^n} \pm \frac{1}{p(n)} \right) \left(\frac{\dim(V^l)}{2^n} \pm \frac{1}{r(n)} \right)^{-1} &= \left(\frac{\tilde{\beta}_l}{2^n} \pm \frac{1}{p(n)} \right) \frac{2^n}{\dim(V^l)} \left(1 \pm \frac{2^n}{r(n) \dim(V^l)} \right)^{-1} \\ &= \left(\frac{\tilde{\beta}_l}{\dim(V^l)} \pm \frac{2^n}{p(n) \dim(V^l)} \right) \left(1 \pm \frac{2^n}{r(n) \dim(V^l)} \right)^{-1}. \end{aligned} \quad (\text{B.6})$$

Using the fact that $\frac{\dim(V^l)}{2^n} \geq 1/q(n)$ for some polynomial $q(n)$, our estimate satisfies

$$\left(\frac{\tilde{\beta}_l}{\dim(V^l)} \pm \frac{q(n)}{p(n)} \right) \left(1 \pm \frac{q(n)}{r(n)} \right)^{-1}. \quad (\text{B.7})$$

Choosing $r(n)$ and $p(n)$ sufficiently large so that $\frac{q(n)}{p(n)} = \epsilon = 1/\text{poly}(n)$ and $\frac{q(n)}{r(n)} = 1/\text{poly}(n) = \epsilon'$, we have an estimate

$$\left(\frac{\tilde{\beta}_l}{\dim(V^l)} \pm \epsilon \right) (1 \pm \epsilon')^{-1}. \quad (\text{B.8})$$

We can then see that

$$(1 \pm \epsilon')^{-1} = 1 \mp \frac{\epsilon'}{1 \pm \epsilon'} = 1 \pm \frac{1}{O(\text{poly}(n))},$$

and hence the estimate in Equation B.8 satisfies

$$\frac{\tilde{\beta}_l}{\dim(V^l)} \pm \frac{1}{O(\text{poly}(n))} \quad (\text{B.9})$$

for some polynomial depending on ϵ and ϵ' .

In the end we obtain an estimate of $\frac{\tilde{\beta}_l}{\dim(V^l)}$ up to inverse polynomial additive error in polynomial time using a DQC1 algorithm, proving the statement. $\square$

References

- [1] Scott Aaronson. On perfect completeness for QMA. *Quantum Information & Computation*, 9(1): 81–89, 2009. DOI: 10.48550/arXiv.0806.0450. [arXiv:0806.0450](#).
- [2] Scott Aaronson, Adam Bouland, Greg Kuperberg, and Saeed Mehraban. The computational complexity of ball permutations. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 317–327, 2017. DOI: 10.1145/3055399.3055453. [arXiv:1610.06646](#).
- [3] Scott Aaronson, Adam Bouland, Greg Kuperberg, and Saeed Mehraban. The computational complexity of ball permutations. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 317–327, 2017. DOI: 10.1145/3055399.3055453. [arXiv:1610.06646](#).
- [4] Michał Adamaszek and Juraj Stacho. Complexity of simplicial homology and independence complexes of chordal graphs. *Computational Geometry*, 57:8 – 18, 2016. ISSN 0925-7721. DOI: <https://doi.org/10.1016/j.comgeo.2016.05.003>.
- [5] Dorit Aharonov, Vaughan Jones, and Zeph Landau. A polynomial quantum algorithm for approximating the jones polynomial. *Algorithmica*, 55(3):395–421, 2009. DOI: 10.48550/arXiv.quant-ph/0511096. [arXiv:quant-ph/0511096](#).
- [6] Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986. DOI: 10.1007/BF02579166.
- [7] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM (JACM)*, 48(4):778–797, 2001. DOI: 10.1145/502090.502097. [arXiv:quant-ph/9802049](#).
- [8] Dominic W Berry, Andrew M Childs, and Robin Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 792–809. IEEE, 2015. DOI: 10.1109/FOCS.2015.54. [arXiv:1501.01715](#).
- [9] Fernando GSL Brandão. Entanglement theory and the quantum simulation of many-body physics. [arXiv:0810.0026](#), 2008. DOI: 10.48550/arXiv.0810.0026.
- [10] Sergey Bravyi. Efficient algorithm for a quantum analogue of 2-SAT. *Contemporary Mathematics*, 536:33–48, 2011. DOI: 10.1090/conm/536/10552.

- [11] Sergey B Bravyi and Alexei Yu Kitaev. Fermionic quantum computation. *Annals of Physics*, 298 (1):210–226, 2002. DOI: 10.1006/aphy.2002.6254. [arXiv:quant-ph/0003137](#).
- [12] Brielin Brown, Steven T. Flammia, and Norbert Schuch. Computational difficulty of computing the density of states. *Physical Review Letters*, 107(4), Jul 2011. ISSN 1079-7114. DOI: 10.1103/physrevlett.107.040501. [arXiv:1010.3060](#).
- [13] Chris Cade and Ashley Montanaro. The quantum complexity of computing Schatten p -norms. In *13th Conference on the Theory of Quantum Computation, Communication and Cryptography*, 2018. DOI: 10.4230/LIPIcs.TQC.2018.4. [arXiv:1706.09279](#).
- [14] Yudong Cao, Jonathan Romero, Jonathan P Olson, Matthias Degroote, Peter D Johnson, Mária Kieferová, Ian D Kivlichan, Tim Menke, Borja Peropadre, Nicolas PD Sawaya, et al. Quantum chemistry in the age of quantum computing. *Chemical reviews*, 119(19):10856–10915, 2019. DOI: 10.1021/acs.chemrev.8b00803. [arXiv:1812.09976](#).
- [15] Henry Cartan and Samuel Eilenberg. *Homological Algebra (PMS-19)*. Princeton University Press, 2016. DOI: doi:10.1515/9781400883844. URL <https://doi.org/10.1515/9781400883844>.
- [16] Andrew M Childs, Yuan Su, Minh C Tran, Nathan Wiebe, and Shuchen Zhu. A theory of Trotter error. [arXiv:1912.08854](#), 2019. DOI: 10.1103/PhysRevX.11.011020.
- [17] P. Marcos Crichigno. Supersymmetry and Quantum Computation. 11 2020.
- [18] Toby Cubitt and Ashley Montanaro. Complexity classification of local Hamiltonian problems. *SIAM Journal on Computing*, 45(2):268–316, 2016. DOI: 10.1137/140998287. [arXiv:1311.3161](#).
- [19] Jan de Gier, Gyorgy Z. Feher, Bernard Nienhuis, and Magdalena Rusaczek. Integrable supersymmetric chain without particle conservation. *J. Stat. Mech.*, 1602(2):023104, 2016. DOI: 10.1088/1742-5468/2016/02/023104.
- [20] A Drucker and R de Wolf. Quantum proofs for classical theorems. *Theory of Computing*, 2011. DOI: 10.48550/arXiv.0910.3376. [arXiv:0910.3376](#).
- [21] Herbert Edelsbrunner and John L. Harer. *Computational Topology: An Introduction*, volume 47. 2009. DOI: 10.1007/978-3-540-33259-6_7.
- [22] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. 2014. DOI: 10.48550/arXiv.1412.6062.
- [23] Paul Fendley and Kareljan Schoutens. Exact results for strongly correlated fermions in $2+1$ dimensions. *Physical review letters*, 95(4):046403, 2005. DOI: 10.1103/PhysRevLett.95.046403. [arXiv:cond-mat/0504595](#).
- [24] Paul Fendley and Kareljan Schoutens. Cooper pairs and exclusion statistics from coupled free-fermion chains. *J. Stat. Mech.*, 0702:P02017, 2007. DOI: 10.1088/1742-5468/2007/02/P02017.
- [25] Paul Fendley, Bernard Nienhuis, and Kareljan Schoutens. Lattice fermion models with supersymmetry. *Journal of Physics A: Mathematical and General*, 36(50):12399–12424, Dec 2003. ISSN 1361-6447. DOI: 10.1088/0305-4470/36/50/004. [arXiv:cond-mat/0307338](#).
- [26] Paul Fendley, Kareljan Schoutens, and Jan de Boer. Lattice models with $\mathcal{N} = 2$ supersymmetry. *Physical review letters*, 90(12):120402, 2003. DOI: 10.1103/PhysRevLett.90.120402. [arXiv:hep-th/0210161](#).

- [27] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Zijun Chen, K. Satzinger, and R. Barends. Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms. *Physical Review Letters*, 125(12), Sep 2020. ISSN 1079-7114. DOI: 10.1103/physrevlett.125.120504. [arXiv:2001.08343](#).
- [28] M. Freedman. P/NP, and the quantum field computer. *Proceedings of the National Academy of Sciences of the United States of America*, 95 1:98–101, 1998. DOI: 10.1073/pnas.95.1.98.
- [29] Michael Freedman, Alexei Kitaev, Michael Larsen, and Zhenghan Wang. Topological quantum computation. *Bulletin of the American Mathematical Society*, 40(1):31–38, 2003. [arXiv:quant-ph/0101025](#).
- [30] Michael H. Freedman, Alexei Kitaev, and Zhenghan Wang. Simulation of topological field theories by quantum computers. *Communications in Mathematical Physics*, 227(3):587–603, Jun 2002. ISSN 1432-0916. DOI: 10.1007/s002200200635. [arXiv:quant-ph/0001071](#).
- [31] Michael H Freedman, Michael Larsen, and Zhenghan Wang. A modular functor which is universal for quantum computation. *Communications in Mathematical Physics*, 227(3):605–622, 2002. DOI: 10.1007/s002200200645. [arXiv:quant-ph/0001108](#).
- [32] Wenbo Fu, Davide Gaiotto, Juan Maldacena, and Subir Sachdev. Supersymmetric Sachdev-Ye-Kitaev models. *Phys. Rev. D*, 95(2):026009, 2017. DOI: 10.1103/PhysRevD.95.026009. [Addendum: *Phys.Rev.D* 95, 069904 (2017)].
- [33] Wenbo Fu, Davide Gaiotto, Juan Maldacena, and Subir Sachdev. Supersymmetric sachdev-ye-kitaev models. *Physical Review D*, 95(6), Mar 2017. ISSN 2470-0029. DOI: 10.1103/physrevd.95.069904.
- [34] Sevag Gharibian, Yichen Huang, Zeph Landau, and Seung Woo Shin. Quantum Hamiltonian complexity. *Theoretical Computer Science*, 10(3):159–282, 2014. DOI: 10.1561/04000000066. [arXiv:1401.3916](#).
- [35] Robert Ghrist. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008. DOI: 10.1090/S0273-0979-07-01191-3.
- [36] Oded Goldreich. On promise problems: A survey. In *Theoretical computer science*, pages 254–290. Springer, 2006. DOI: 10.1007/11685654_12.
- [37] David Gosset and Daniel Nagaj. Quantum 3-sat is QMA_1 -complete. *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, Oct 2013. DOI: 10.1109/focs.2013.86. [arXiv:1302.0290](#).
- [38] Tarun Grover, D.N. Sheng, and Ashvin Vishwanath. Emergent Space-Time Supersymmetry at the Boundary of a Topological Phase. *Science*, 344(6181):280–283, 2014. DOI: 10.1126/science.1248253.
- [39] Anna Gundert and May Szedlák. Higher dimensional cheeger inequalities. In *Proceedings of the thirtieth annual symposium on Computational geometry*, pages 181–188, 2014. DOI: 10.1145/2582112.2582118. [arXiv:1401.2290](#).
- [40] Casper Gyurik, Chris Cade, and Vedran Dunjko. Towards quantum advantage for topological data analysis. [arXiv:2005.02607](#), 2020. DOI: 10.22331/q-2022-11-10-855.
- [41] Christian Hagendorf, Thessa B Fokkema, and Liza Huijse. Bethe ansatz solvability and supersymmetry of the M2 model of single fermions and pairs. *Journal of Physics A: Mathematical and*

- Theoretical*, 47(48):485201, Nov 2014. ISSN 1751-8121. DOI: 10.1088/1751-8113/47/48/485201. [arXiv:1408.4403](#).
- [42] Danijela Horak and Jürgen Jost. Spectra of combinatorial laplace operators on simplicial complexes. *Advances in Mathematics*, 244:303–336, 2013. DOI: 10.1016/j.aim.2013.05.007. [arXiv:1105.2712](#).
 - [43] K. Hori, S. Katz, A. Klemm, R. Pandharipande, R. Thomas, C. Vafa, R. Vakil, and E. Zaslow. *Mirror symmetry*, volume 1 of *Clay mathematics monographs*. AMS, Providence, USA, 2003. DOI: 10.1088/1126-6708/2001/07/022.
 - [44] Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V. Albert, and John Preskill. Provably efficient machine learning for quantum many-body problems. 2021. DOI: DOI: 10.1126/science.abk3333.
 - [45] L. Huijse, Dhagash Bharatbhai Mehta, N. Moran, K. Schoutens, and J. Vala. Supersymmetric lattice fermions on the triangular lattice: Superfrustration and criticality. *New Journal of Physics*., 14:073002, 2012. DOI: 10.1088/1367-2630/14/7/073002.
 - [46] Liza Huijse, Kareljan Schoutens, et al. Supersymmetry, lattice fermions, independence complexes and cohomology theory. *Advances in Theoretical and Mathematical Physics*, 14(2):643–694, 2010. DOI: 10.48550/arXiv.0903.0784. [arXiv:0903.0784](#).
 - [47] Dominik Janzing and Pawel Wocjan. A simple PromiseBQP-complete matrix problem. *Theory of computing*, 3(1):61–79, 2007. 10.4086/toc.2007.v003a004.
 - [48] Jakob Jonsson. Certain homology cycles of the independence complex of grids. *Discrete & Computational Geometry*, 43(4):927–950, 2010. DOI: 10.1007/s00454-009-9224-9.
 - [49] Georg Junker. *Supersymmetric Methods in Quantum and Statistical Physics*. Theoretical and Mathematical Physics. Springer-Verlag Berlin Heidelberg, 1996. DOI: 10.1007/978-3-642-61194-0.
 - [50] Tomasz Kaczynski, Konstantin Michael Mischaikow, Marian Mrozek, and Konstantin Mischaikow. *Computational homology*. Applied mathematical sciences (Springer-Verlag New York Inc.); v. 157. Springer, New York, 2004. ISBN 0387408533.
 - [51] Volker Kaibel and Marc E. Pfetsch. Some Algorithmic Problems in Polytope Theory. *arXiv Mathematics e-prints*, art. math/0202204, February 2002.
 - [52] Julia Kempe and Oded Regev. 3-local hamiltonian is QMA-complete. *Quantum Information & Computation*, 3(3):258–264, 2003. [arXiv:quant-ph/0302079](#).
 - [53] Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local Hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006. DOI: 10.1137/S0097539704445226. [arXiv:quant-ph/0406180](#).
 - [54] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. Graduate Studies in Mathematics, 2002.
 - [55] A.Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, Jan 2003. ISSN 0003-4916. DOI: 10.1016/s0003-4916(02)00018-0. [arXiv:quant-ph/9707021](#).
 - [56] Ville Lahtinen and Jiannis Pachos. A short introduction to topological quantum computation. *SciPost Physics*, 3(3), Sep 2017. ISSN 2542-4653. DOI: 10.21468/scipostphys.3.3.021. [arXiv:1705.04103](#).

- [57] James PF LeBlanc, Andrey E Antipov, Federico Becca, Ireneusz W Bulik, Garnet Kin-Lic Chan, Chia-Min Chung, Youjin Deng, Michel Ferrero, Thomas M Henderson, Carlos A Jiménez-Hoyos, et al. Solutions of the two-dimensional hubbard model: Benchmarks and results from a wide range of numerical algorithms. *Physical Review X*, 5(4):041041, 2015. DOI: 10.1103/PhysRevX.5.041041. [arXiv:1505.02290](#).
- [58] James R Lee, Shayan Oveis Gharan, and Luca Trevisan. Multiway spectral partitioning and higher-order cheeger inequalities. *Journal of the ACM (JACM)*, 61(6):1–30, 2014. DOI: 10.1145/2665063. [arXiv:1111.1055](#).
- [59] Yi-Kai Liu, Matthias Christandl, and F. Verstraete. Quantum computational complexity of the n -representability problem: QMA complete. *Physical Review Letters*, 98(11), Mar 2007. ISSN 1079-7114. DOI: 10.1103/physrevlett.98.110503. [arXiv:quant-ph/0609125](#).
- [60] Seth Lloyd. Universal quantum simulators. *Science*, pages 1073–1078, 1996. DOI: science.273.5278.1073. [arXiv:quant-ph/9703054](#).
- [61] Seth Lloyd, Silvano Garnerone, and Paolo Zanardi. Quantum algorithms for topological and geometric analysis of data. *Nature Communications*, 7(1):10138, 2016. DOI: 10.1038/ncomms10138. [arXiv:1408.3106](#).
- [62] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016. DOI: 10.1088/1367-2630/18/2/023023. [arXiv:1509.04279](#).
- [63] Jarrod R McClean, Nicholas C Rubin, Joonho Lee, Matthew P Harrigan, Thomas E O’Brien, Ryan Babbush, William J Huggins, and Hsin-Yuan Huang. What the foundations of quantum computer science teach us about chemistry. [arXiv:2106.03997](#), 2021. DOI: 10.1063/5.0060367.
- [64] Jiří Minář, Bart van Voorden, and Kareljan Schoutens. Kink dynamics and quantum simulation of supersymmetric lattice Hamiltonians. [arXiv:2005.00607](#), 2020.
- [65] Tomoyuki Morimae, Keisuke Fujii, and Joseph F Fitzsimons. Hardness of classically simulating the one-clean-qubit model. *Physical review letters*, 112(13):130502, 2014. DOI: 10.1103/PhysRevLett.112.130502. [arXiv:1312.2496](#).
- [66] H. Nicolai. Supersymmetry and Spin Systems. *Journal of Physics A*, 9:1497–1506, 1976. DOI: 10.1088/0305-4470/9/9/010.
- [67] Bryan O’Gorman, Sandy Irani, James Whitfield, and Bill Fefferman. Electronic structure in a fixed basis is qma-complete. [arXiv:2103.08215](#), 2021. DOI: 10.48550/arXiv.2103.08215.
- [68] Sourabh Palande and Bei Wang. Spectral sparsification of simplicial complexes for clustering and label propagation. *Journal of Computational Geometry*, 11(1):176–211, 2020. DOI: 10.20382/jocg.v11i1a8. [arXiv:1708.08436](#).
- [69] Ori Parzanchevski, Ron Rosenthal, and Ran J Tessler. Isoperimetric inequalities in simplicial complexes. *Combinatorica*, 36(2):195–227, 2016. DOI: 10.1007/s00493-014-3002-x. [arXiv:1207.0638](#).
- [70] Richard Peng, He Sun, and Luca Zanetti. Partitioning well-clustered graphs: Spectral clustering works! In *Conference on Learning Theory*, pages 1423–1455. PMLR, 2015. DOI: 10.1137/15M1047209. [arXiv:1411.2021](#).

- [71] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5(1):1–7, 2014. DOI: 10.1038/ncomms5213. [arXiv:1304.3061](#).
- [72] John Preskill. Lecture notes for Physics 219: Quantum computation. *Caltech Lecture Notes*, 1999. [Link](#).
- [73] Armin Rahmani, Xiaoyu Zhu, Marcel Franz, and Ian Affleck. Emergent Supersymmetry from Strongly Interacting Majorana Zero Modes. *prl*, 115(16):166401, October 2015. DOI: 10.1103/PhysRevLett.115.166401.
- [74] Christian Reiher. The clique density theorem. *Annals of Mathematics*, pages 683–707, 2016. DOI: 10.48550/arXiv.1212.2454. [arXiv:1212.2454](#).
- [75] Burak Sahinoglu and Rolando Somma. Hamiltonian simulation in the low energy subspace. *Bulletin of the American Physical Society*. DOI: 10.1038/s41534-021-00451-w. [arXiv:2006.02660](#).
- [76] Hubert Saleur and Nicholas P Warner. Lattice models and $n = 2$ supersymmetry. In *Quantum Field Theory and String Theory*, pages 335–377. Springer, 1995. [arXiv:hep-th/9311138](#).
- [77] Raoul Santachiara and Kareljan Schoutens. Supersymmetric model of spin-1/2 fermions on a chain. *Journal of Physics A Mathematical General*, 38(24):5425–5439, June 2005. DOI: 10.1088/0305-4470/38/24/003.
- [78] Peter Scheiblechner. On the complexity of deciding connectedness and computing betti numbers of a complex algebraic variety. *Journal of Complexity*, 23(3):359–379, 2007. DOI: 10.1016/j.jco.2007.03.008.
- [79] Jacob T Seeley, Martin J Richard, and Peter J Love. The Bravyi-Kitaev transformation for quantum computation of electronic structure. *The Journal of chemical physics*, 137(22):224109, 2012. DOI: 10.1063/1.4768229. [arXiv:1208.5986](#).
- [80] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, Oct 1997. ISSN 1095-7111. DOI: 10.1137/s0097539795293172. [arXiv:quant-ph/9508027](#).
- [81] John Steenbergen, Caroline Klivans, and Sayan Mukherjee. A cheeger-type inequality on simplicial complexes. *Advances in Applied Mathematics*, 56:56–77, 2014. DOI: 10.1016/j.aam.2014.01.002. [arXiv:1209.5091](#).
- [82] Hendrik van Eerten. Extensive ground state entropy in supersymmetric lattice models. *Journal of mathematical physics*, 46(12):123302, 2005. DOI: 10.1063/1.2142836. [arXiv:cond-mat/0509581](#).
- [83] Larry Wasserman. Topological data analysis. *Annual Review of Statistics and Its Application*, 5: 501–532, 2018. DOI: 10.1146/annurev-statistics-031017-100045. [arXiv:1609.08227](#).
- [84] John Watrous. *Quantum Computational Complexity*, pages 7174–7201. Springer New York, New York, NY, 2009. ISBN 978-0-387-30440-3. DOI: 10.1007/978-0-387-30440-3_428. [quant-ph/0804.3401](#).
- [85] Charles A. Weibel. *An Introduction to Homological Algebra*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1994. DOI: 10.1017/CBO9781139644136.

- [86] Charles A. Weibel. Chapter 28 - History of Homological Algebra. In I.M. James, editor, *History of Topology*, pages 797–836. North-Holland, Amsterdam, 1999. ISBN 978-0-444-82375-5. DOI: <https://doi.org/10.1016/B978-044482375-5/50029-8>.
- [87] Edward Witten. Dynamical Breaking of Supersymmetry. *Nuclear Physics B*, 188:513, 1981. DOI: 10.1016/0550-3213(81)90006-7.
- [88] Edward Witten. Constraints on Supersymmetry Breaking. *Nuclear Physics B*, 202:253, 1982. DOI: 10.1016/0550-3213(82)90071-2.
- [89] Edward Witten. Supersymmetry and Morse theory. *J. Differential Geometry*, 17(4):661–692, 1982. DOI: 10.4310/jdg/1214437492.
- [90] Edward Witten. Quantum Field Theory and the Jones Polynomial. *Communications in Mathematical Physics*, 121:351–399, 1989. DOI: 10.1007/BF01217730.
- [91] Xiao Yang and Paul Fendley. Non-local spacetime supersymmetry on the lattice. *Journal of Physics A Mathematical General*, 37(38):8937–8948, September 2004. DOI: 10.1088/0305-4470/37/38/003.
- [92] Afra J. Zomorodian. *Topology for Computing*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2005. DOI: 10.1017/CBO9780511546945.