Quantum SAT Problems with Finite Sets of Projectors Are Complete for a Plethora of Classes

Ricardo Rivera Cardoso
RCQI, Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia

Alex Meiburg
Perimeter Institute for Theoretical Physics, Waterloo, Canada
Institute for Quantum Computing, University of Waterloo, Canada

Daniel Nagaj
RCQI, Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia

Abstract

Previously, all known variants of the Quantum Satisfiability (QSAT) problem – consisting of determining whether a k-local (k-body) Hamiltonian is frustration-free – could be classified as being either in P; or complete for NP, MA, or QMA_1. Here, we present new qubit variants of this problem that are complete for BQP_1, coRP, QCMA, $\mathsf{PI}(\mathsf{coRP}, \mathsf{NP})$, $\mathsf{PI}(\mathsf{BQP}_1, \mathsf{NP})$, $\mathsf{PI}(\mathsf{BQP}_1, \mathsf{MA})$, $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$, $\mathsf{SoPU}(\mathsf{BQP}_1, \mathsf{NP})$, and $\mathsf{SoPU}(\mathsf{BQP}_1, \mathsf{MA})$. Our result implies that a complete classification of quantum constraint satisfaction problems (QCSPs), analogous to Schaefer's dichotomy theorem for classical CSPs, must either include these 13 classes, or otherwise show that some are equal. Additionally, our result showcases two new types of QSAT problems that can be decided efficiently, as well as the first nontrivial BQP_1-complete problem.

We first construct QSAT problems on qudits that are complete for BQP_1, coRP, and QCMA. These are made by restricting the finite set of Hamiltonians to consist of elements similar to H_{init}, H_{prop}, and H_{out}, seen in the circuit-to-Hamiltonian transformation. Usually, these are used to demonstrate hardness of QSAT and *Local Hamiltonian* problems, and so our proofs of hardness are simple. The difficulty lies in ensuring that all Hamiltonians generated with these three elements can be decided in their respective classes. For this, we build our Hamiltonian terms with high-dimensional data and clock qudits, ternary logic, and either monogamy of entanglement or specific clock encodings. We then show how to express these problems in terms of qubits, by proving that any QCSP can be reduced to a qubit problem while maintaining the same complexity – something not believed possible classically. The remaining six problems are obtained by considering "sums" and "products" of some of the QSAT problems mentioned here. Before this work, the QSAT problems generated in this way resulted in complete problems for PI and SoPU classes that were trivially equal to NP, MA, or QMA_1. We thus commence the study of these new and seemingly nontrivial classes.

While [Meiburg, 2021] first sought to prove completeness for coRP, BQP_1, and QCMA, we note that those constructions are flawed. Here, we rework them, provide correct proofs, and obtain improvements on the required qudit dimensionality.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Quantum complexity theory, quantum satisfiability, circuit-to-Hamiltonian, pairwise union of classes, pairwise intersection of classes

Digital Object Identifier 10.4230/LIPIcs.TQC.2025.6

Related Version *Full Version*: `https://arxiv.org/abs/2506.07244` [14]

Funding RRC and DN were supported by projects APVV-22-0570 (DeQHOST) and VEGA No. 2/0164/25 (QUAS). AM was supported by the Natural Sciences and Engineering Research Council of Canada through grant number RGPIN-2019-04198. IQC and the Perimeter Institute are supported in part by the Government of Canada through ISED and the Province of Ontario.

Acknowledgements We thank Alex Grilo and Dorian Rudolph for helpful discussions, and the authors of Ref. [12] for the inspiration of Figure 8 in the full version of this paper.

20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025).
Editor: Bill Fefferman; Article No. 6; pp. 6:1–6:24
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Many of the interesting and puzzling phenomena in many-body physics occurs at the ground state of materials. One way to study quantum systems in this state is through their ground state energy, as this quantity can be used to provide information about physical and chemical properties of the system. It is thus of great interest to calculate or even estimate this quantity. This task is embodied by the k-Local Hamiltonian (k-LH) problem. Specifically, given a *k-local* (k-body) Hamiltonian – an operator of the form $H = \sum_i h_i$ where each h_i acts on at most k qubits – and two numbers $a, b \in \mathbb{R}$ with $b - a \geq 1/poly(n)$, this problem consists of distinguishing between the cases where H has an eigenvalue less than a or greater than b. Kitaev [30] showed that k-LH with $k \geq 5$ (and later improved to $k \geq 2$ [28]) is unlikely to be decided efficiently with a classical or quantum computer. In complexity theory terms, k-LH with $k \geq 2$ is QMA-complete.[1]

The LH problem is considered a "weak" quantum constraint satisfaction problem (QCSP) as states with energy less than a do not necessarily minimize the energy of each h_i. For this reason, LH is often compared to MAX-k-SAT instead of the "strong" CSP k-SAT. Due to the immense importance of SAT in classical complexity and other hard sciences, Bravyi [6] defined the Quantum k-SAT (k-QSAT) problem. Given a set of k-local projectors (also referred as *clauses* or *constraints*) and a number $b \in \mathbb{R}$, this problem consists of distinguishing between the cases where there exists a state that simultaneously lies in the null space of all projectors, or for all states, the penalty incurred by violations of the constraints is greater than b.[2] Bravyi showed that 2-QSAT on qubits is in P while k-QSAT with $k \geq 4$ (and later improved to $k \geq 3$ [22]) is QMA_1-complete when using the Clifford+T gate set $\mathcal{G}_8 = \{H, \mathrm{CNOT}, T\}$.[3]

Interestingly, these two problems have in common that they are in P for a certain k but appear to become much harder for $k + 1$: LH is in P for $k \leq 1$ and becomes QMA-complete for $k > 1$, while QSAT is in P for $k \leq 2$ and $\mathsf{QMA}_1^{\mathcal{G}_8}$-complete for $k > 2$. This is not entirely surprising since the Hamiltonians considered in the problems have no restriction other than their locality, and perhaps the difficulty lies in deciding "unphysical" Hamiltonians. Following this line of thought, others have considered variations of these problems where the h_i are drawn from more realistic and relevant sets that satisfy some property or correspond to a physical model. To name a few, these may be stoquastic [7], commuting [9], fermionic [31], bosonic [41], or from models like the Heisenberg [39] and Bose-Hubbard [15]. In addition, one might also consider placing restrictions on the geometry of the problem [33, 23, 2, 26, 36].

In a landmark result, Cubitt and Montanaro [18] showed that any LH problem where the h_i are drawn from a finite set of at most 2-local qubit Hermitian matrices can be classified as being either in P, NP-complete, StoqMA-complete, or QMA-complete.[4] As decision problems in the latter three classes are not known to be efficiently solvable in either classical or quantum computers, they showed that the only Hamiltonians of this type for which the LH problem can be solved efficiently are those with only 1-local terms. This is significant,

[1] The class QMA can be thought of as the quantum analog of NP, or more accurately MA since the class has probabilistic acceptance and rejection.

[2] Alternatively, this problem can be defined with local Hamiltonians instead of projectors, in which case, the problem is equivalent to determining whether the Hamiltonian is frustration-free.

[3] QMA_1 is the one-sided error variation of QMA with perfect completeness, i.e. instances for which the answer is "yes" (in this case frustration-free Hamiltonians) are accepted with certainty. The notation $\mathcal{G}_8$ stems from Ref. [4] and denotes the Clifford-cyclotomic gate set of degree of 8. The reason why it is necessary to specify the gate set for classes with perfect completeness is discussed in Section A.2.

[4] StoqMA is the class of problems equivalent to estimating the ground state energy of the transverse-field Ising model [8].

as many relevant Hamiltonians in nature can be approximated by 2-local Hamiltonians of this type (e.g. all those supported on Pauli operators like Heisenberg and Ising spin glass models), and it is then likely that estimating their ground state energy efficiently lies outside of reach. Moreover, their result has led to a much larger repertoire of problems from which to construct reductions and potentially show the complexity of other computational problems.

Prior to our work, all known QSAT problems with finite or infinite sets of local interactions could be classified as being either in P, NP-complete, MA-complete, or QMA_1-complete, but this list is not known to be exhaustive in either case. The fact that QSAT has resisted classification can be attributed to two factors. First, is that since most relevant instances of QSAT can be decided classically (2-QSAT is in P), there is a lack of interest to search for a classification of QSAT problems with $k > 2$. This is unlike in the LH problem where most relevant instances were hard (2-LH is QMA-complete), motivating the study of Cubitt and Montanaro. Second, is the fact that QSAT problems are usually complete for classes that are harder to work with as they seem to depend on gate sets. In this work, it is our goal to concretize the implications that such a theorem may have, and hence motivate its study.

1.1 Summary of results

Our main result establishes that the QSAT problem SLCT-QSAT is $\mathsf{BQP}_1^{\mathcal{G}_8}$-complete. However, as the construction and analysis of this problem is contrived, we first show that the simpler and less optimized version of this problem, LCT-QSAT, is also complete for this class.

▶ **Theorem 1.** *The problem* LINEAR-CLOCK-TERNARY-QSAT (LCT-QSAT) *with* 4*-local clauses acting on* 17*-dimensional qudits is* $\mathsf{BQP}_1^{\mathcal{G}_8}$*-complete.*

An interesting feature of this problem, and one that may be of independent interest, is that this problem makes clever use of the principle of monogamy of entanglement to strongly constrain the structure of input instances, facilitating the task of deciding whether they are frustration-free.[5] Unfortunately, this trick comes at a price of high qudit dimensionality. Our main result shows that by relaxing the constraint on the instance's structure and instead study the instances more closely, we can obtain a similar problem with the same complexity but with reduced qudit dimensionality.

▶ **Theorem 2.** *The problem* SEMILINEAR-CLOCK-TERNARY-QSAT (SLCT-QSAT) *with* 4*-local clauses acting on* 6*-dimensional qudits is* $\mathsf{BQP}_1^{\mathcal{G}_8}$*-complete.*

Recently, among many other interesting results, Rudolph [37] demonstrated that $\mathsf{BQP}_1^{\mathcal{G}_{2^i}} = \mathsf{BQP}_1^{\mathcal{G}_{2^j}}$ for any $i, j \in \mathbb{N}$. In other words, any problem in BQP_1 using a Clifford-cyclotomic gate set of degree 2^i can be perfectly simulated with one of degree 2^j for all $i, j \in \mathbb{N}$. For us, this then implies that:

▶ **Corollary 3.** *The problems* LCT-QSAT *and* SLCT-QSAT *are* BQP_1*-complete with any gate set* $\mathcal{G}_{2^l}$ *with* $l \in \mathbb{N}$*.*

Subsequently, by performing slight modifications to the clauses of SLCT-QSAT, we also obtain QCMA-complete and coRP-complete problems:

▶ **Theorem 4.** *The problem* WITNESSED SLCT-QSAT *with* 4*-local clauses acting on* 8*-dimensional qudits is* QCMA*-complete.*

[5] This construction is the most faithful to those considered by Meiburg in Ref. [32].

▶ **Theorem 5.** *The problem* Classical SLCT-QSAT *with* 5*-local clauses acting on* 8*-dimensional qudits is* coRP*-complete.*

Then, using a similar application of monogamy of entanglement as in LCT-QSAT, we demonstrate that we can reduce any QCSP on qudits to another one on qubits.

▶ **Theorem 6** (informal). *Every QCSP $\mathcal{C}$ on qudits is equivalent in difficulty to some other QCSP $\mathcal{C}'$ on qubits.*

▶ **Corollary 7.** *Together, Theorems 2 and 4–6 imply:*
1. $\mathrm{SLCT\text{-}QSAT}_2$ *is a* $\mathsf{BQP}_1^{\mathcal{G}_8}$*-complete problem on qubits with* 48*-local clauses.*
2. Witnessed $\mathrm{SLCT\text{-}QSAT}_2$ *is a* QCMA*-complete problem on qubits with* 48*-local clauses.*
3. Classical $\mathrm{SLCT\text{-}QSAT}_2$ *is a* coRP*-complete problem on qubits with* 60*-local clauses.*

We refer to these problems by the same name as before, except that we now add a subindex to represent that the problem refers to the qubit version, e.g. $\mathrm{SLCT\text{-}QSAT}_2$ is the QSAT problem that results from the reduction of SLCT-QSAT.

Finally, there is a notion of *direct product* "$\otimes$" and *direct sum* "$\oplus$" (Definitions 17 and 18) for both CSPs and QCSPs, which we use to show that there are six new QSAT problems that are complete for classes $\mathsf{PI}(A, B)$ and $\mathsf{SoPU}(A, B)$, where A and B are themselves complexity classes. $\mathsf{PI}(A, B)$ stands for the *pairwise intersection of classes* (Definition 11), and $\mathsf{SoPU}(A, B)$ for the *star of pairwise union of classes* (Definition 12). Roughly, these two classes correspond to the sets of problems that can be expressed as the intersection and union (respectively) of a problem in A and a problem in B.[6] We show:

▶ **Theorem 8.** *Let "$\otimes$" and "$\oplus$" denote the* direct product *and* direct sum *for quantum constraint satisfaction problems. Pairwise combinations of the four* QSAT *problems –* 3-SAT*,* Classical $\mathrm{SLCT\text{-}QSAT}_2$*,* $\mathrm{SLCT\text{-}QSAT}_2$*, and* Stoquastic 6-SAT *– yield the following complete problems:*
1. Classical $\mathrm{SLCT\text{-}QSAT}_2 \otimes$ 3-SAT *is* $\mathsf{PI}(\mathsf{coRP}, \mathsf{NP})$*-complete.*
2. Classical $\mathrm{SLCT\text{-}QSAT}_2 \oplus$ 3-SAT *is* $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$*-complete.*
3. $\mathrm{SLCT\text{-}QSAT}_2 \otimes$ 3-SAT *is* $\mathsf{PI}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{NP})$*-complete.*
4. $\mathrm{SLCT\text{-}QSAT}_2 \oplus$ 3-SAT *is* $\mathsf{SoPU}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{NP})$*-complete.*
5. $\mathrm{SLCT\text{-}QSAT}_2 \otimes$ Stoquastic 6-SAT *is* $\mathsf{PI}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{MA})$*-complete.*
6. $\mathrm{SLCT\text{-}QSAT}_2 \oplus$ Stoquastic 6-SAT *is* $\mathsf{SoPU}(\mathsf{BQP}_1^{\mathcal{G}_8}, \mathsf{MA})$*-complete.*

Finally, given that the QSAT problems in Corollary 7 and Theorem 8 consist of finite sets of projects with $\mathcal{O}(1)$-local qubit clauses, and similarly 2-SAT, 3-SAT, Stoquastic 6-SAT, and 3-QSAT (which are respectively in P, NP-complete, MA-complete and $\mathsf{QMA}_1^{\mathcal{G}_8}$-complete), our results imply that:

▶ **Corollary 9.** *A complete classification theorem for strong QCSPs with $\mathcal{O}(1)$-local clauses acting on qubits must either include at least* 13 *classes, or otherwise indicate that some of these are equal.*

The relationship between the 13 classes mentioned here is shown in Figure 1.

[6] These classes are not to be confused with $A \cap B$ and $A \cup B$. $A \cap B$ corresponds to the set of problems that are in both A and B, while $A \cup B$ corresponds to those that are in either A or B.

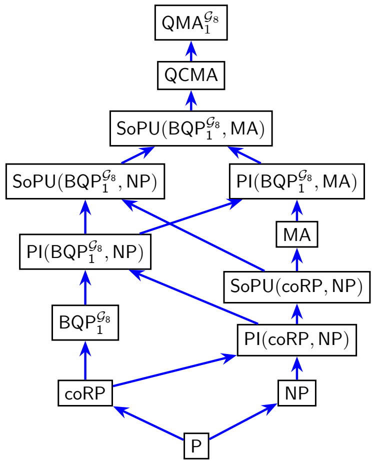

Figure 1 The classes for which we now have a complete strong QCSP, and their corresponding inclusions. In this work, we show completeness for quantum complexity classes with perfect completeness using the Clifford+T gate set $\mathcal{G}_8 = \{H, T, \text{CNOT}\}$. Rudolph's result [37] further strengthens ours by showing that $\mathsf{BQP}_1^{\mathcal{G}_8} = \mathsf{BQP}_1^{\mathcal{G}_{2^l}}$ for all $l \geq 1$. We discuss some of the inclusions in this figure in Section 2.6 and Section A.2.

2 Contributions

In this section, we summarize the main ideas and proof techniques related to the results presented in Section 1.1. In particular, we detail the main roadblocks in the construction of each QSAT problem, and how we overcome them. The full proofs of the statements here can be found in the full version of the text [14].

Section A covers the notation and background information used here. For the rest of this section, we fix the gate set $\mathcal{G}_8$ and omit the superscript when referring to BQP_1 and QMA_1, except when needed for emphasis.

2.1 BQP_1-complete problem

The goal of the construction is to design a QSAT problem that can encode the computation of any quantum circuit in BQP_1, while keeping all its instances solvable in quantum polynomial-time with perfect completeness and bounded soundness. We define the problem using projectors Π_{init}, $\Pi_{prop,U}$, and Π_{out} similar to P_{init}, $P_{prop,U}$, and P_{out} defined in Equation (6).[7] To see why our projectors must differ from the original ones, consider the QSAT problem built with $\{P_{init}, P_{prop,U}, P_{out}, P_{start}, P_{clock}, P_{end}\}$. Showing that the problem is BQP_1-hard is straightforward, as we can encode the circuit that computes the answer to a BQP_1 problem in a similar way as that shown in Section A.3. This time however, all data particles in the instance should be initialized, instead of having free particles whose role is to accommodate a witness state. The difficulty lies in demonstrating that every instance generated with a polynomial number of these projectors can also be decided in BQP_1. There is a fundamental and a practical limitation for this:

[7] The projectors P_{start}, P_{clock}, and P_{stop} associated with the clock encoding remain unchanged and are integrated into the definitions of Π_{init}, Π_{prop}, and Π_{out}.

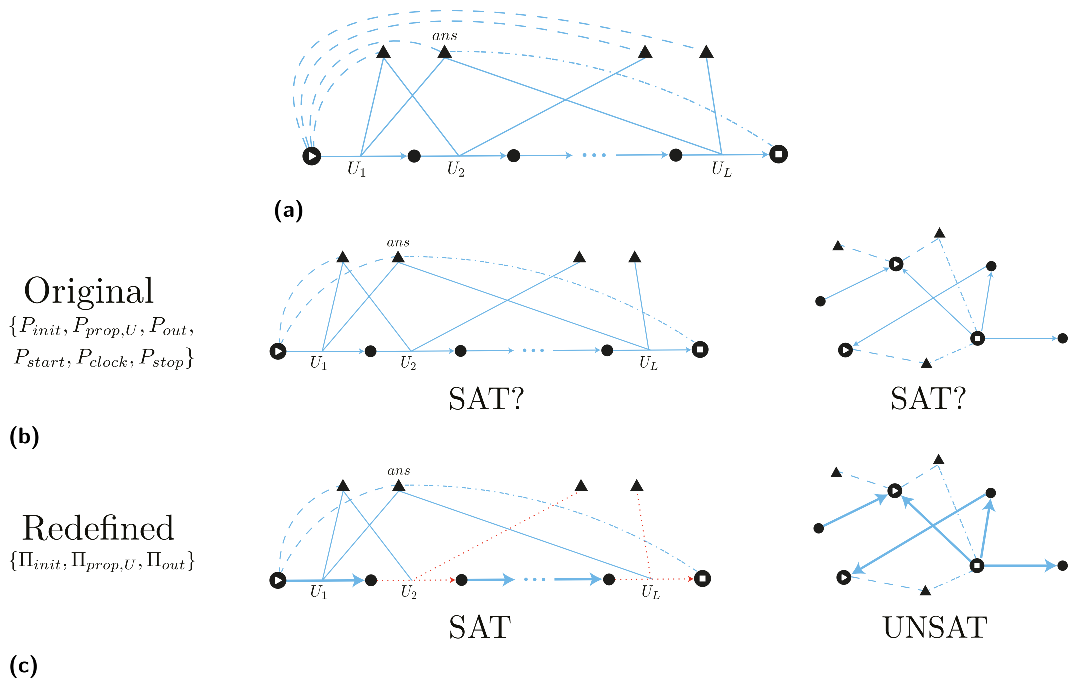

Figure 2 (a) A typical instance that encodes the computation of a BQP_1 circuit $U_L \dots U_1$. The satisfiability of the instance can also be decided in BQP_1. (b) Examples of troublesome instances whose satisfiability is not known to be decidable with a BQP_1 algorithm. (c) The above instances recast with the new set of projectors $\{\Pi_{init}, \Pi_{prop,U}, \Pi_{out}\}$. The bold blue arrows represent the $\Pi_{prop,U}$ clauses which now also indicate the particles should be maximally entangled, and the dotted red arrows those that are connected to undefined logical qudits. With these projectors, their satisfiability can be more easily decided. The left instance is satisfiable due to the undefined clauses, while the one on the right is unsatisfiable, as any potential satisfying state violates monogamy of entanglement. The instance in (a) has the same meaning/satisfiability with either set of projectors.

- Instances which encode the computation of a QMA_1 problem, e.g. the instance in Figure 4 and the left instance in Figure 2b, are valid inputs. This is problematic since it is unknown how to decide these instances in BQP_1 (and doing so would show that $\mathsf{BQP}_1 = \mathsf{QMA}_1$).
- Input instances may form intricate structures complicating the task of deciding if a satisfying state exists, e.g. the right instance in Figure 2b.

We define the projectors Π_{init}, Π_{prop}, and Π_{out} to address these two difficulties (see Figure 2c). Importantly, these projectors do not significantly alter the proof that the problem is BQP_1-hard and can proceed as mentioned. Now, let us briefly discuss how we overcome both difficulties.

Instances like those in Figure 4, which have a proper structure and uninitialized data particles, are prototypical examples of QMA instances. These "free" particles give one the freedom to guess if there exists a state they can be in such that the instance can be satisfied (or equivalently be provided with such a state which we verify). To address this issue, we remove the need to guess a satisfying state by introducing a new *undefined* basis state $|?\rangle$ (making the data particles 3-dimensional), such that setting the free data particles to this state always results in a satisfiable instance. More specifically, we achieve this by defining $\Pi_{prop,U}$ so that if any data particle in the clause is in state $|?\rangle$, the clause is satisfied without needing to apply the associated unitary.[8] Then, for these instances, the satisfying state is given by a truncated version of the history state (without a witness) since the computation is no longer required to elapse past the first $\Pi_{prop,U}$ clause acting on an undefined state. We say the instance is now "trivially satisfiable" as its structure alone suffices to determine its satisfiability.

To determine the satisfiability of intricate instances, the projectors are now also defined to leverage the principle of monogamy of entanglement. Each clock particle is equipped with two 2-dimensional auxiliary subspaces *CA* and *CB* (making them 12-dimensional) and the $\Pi_{prop,U}$ clauses are then defined to require that the *CB* subspace of the predecessor clock particle forms a $|\Phi^+\rangle$ Bell pair with the *CA* subspace of its successor. Then, if a *CA* or *CB* subspace is required to form more than one Bell pair, the principle of monogamy of entanglement states that only one of these clauses can be satisfied, and so the instance is unsatisfiable. Therefore, instances that are not deemed unsatisfiable because of this reason must form one-dimensional chains with a unique "time" direction. Finally, to guarantee that Π_{init} and Π_{out} only act on the ends of the chain, these make use of a new *endpoint* particle consisting of a single two-dimensional space *EC* and require that it also forms a Bell pair with either the *CA* (for Π_{init}) or *CB* (for Π_{out}) subspace of a clock particle. These modifications thus yield a 17-dimensional local Hilbert space: a 3-dimensional data subspace, plus a 2-dimensional endpoint subspace, plus a 12-dimensional clock subspace.

Although these modifications do not get rid off all difficulties, a comprehensive analysis of the resulting instances can be used to demonstrate that a hybrid algorithm can determine the satisfiability status of all input instances. Briefly, the classical part of the algorithm evaluates the structure of the clauses in the instance and concludes whether it is trivially unsatisfiable, trivially satisfiable, or is one requiring the assistance of a quantum subroutine. Trivially unsatisfiable instances are those whose clause arrangement imply one or several clauses cannot be simultaneously satisfied, like those that violate monogamy of entanglement. On the other hand, trivially satisfiable instances are those whose clauses do not create any conflicts but whose structure is simple enough that the satisfying state can be inferred, like those with a proper structure and uninitialized data particles. We show that the only type of instances that are not in either one of these cases, are those like Figure 2a which express the computation of a quantum circuit on initialized ancilla qubits. For these instances, the classical algorithm makes use of a quantum subroutine that executes the quantum circuit expressed by the instance, while simultaneously measuring the eigenvalues of relevant projectors. The measurement outcomes indicate whether the instance should be accepted or rejected.

2.2 Reducing the qudit dimensionality

This section argues that even by removing the projectors that demand successive clock (or endpoint) particles must be entangled with each other, the satisfiability of instances remains the same. Specifically, we argue that the propagation rules, the choice of clock encoding, and the requirement to maintain a consistent clock register state at all times suffice to show that any instance in which the clock particles are not arranged linearly and do not point in the same direction is unsatisfiable. Consequently, there is no longer a need for auxiliary

[8] Although the data particles are 3-dimensional and the unitaries are gates from a set designed to act on qubits, these cause no conflicts as the gates will never act on undefined data particles.

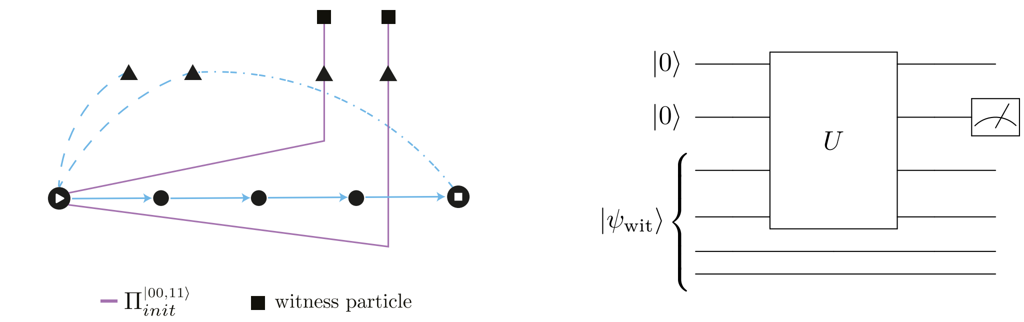

(a) **(b)**

Figure 3 (a) Toy example of an input "quantum" instance with a TACC of length $L = 4$, acting on four logical qudits and two witness qudits. Although not illustrated, the Π_{prop} clauses are assumed to have unitaries $U_1, \ldots, U_4$ which act only on the logical qudits of the instance. These unitaries define a circuit $U = U_4U_3U_2U_1$. (b) Quantum circuit representing the instance on the left.

subspaces or endpoint particles. Together, these results show that while the use of monogamy of entanglement in the construction does facilitate some proofs, it is not crucial for the construction. Removing these elements reduce the local dimension from 17 down to 6.

The main challenge in this construction stems from the weaker constraints that the Π_{init} and Π_{out} clauses set instead of the endpoint particles. In summary, instances with more than a single Π_{init}/Π_{out} pair may now be satisfiable. Part of the proof of this section requires showing that if such sub-instances are potentially satisfiable, they can be further separated into smaller linear instances, each with a single Π_{init}/Π_{out} pair. Each of these smaller pieces is then satisfied by a history state, while the clauses connecting them together (arranged in any shape) can be satisfied trivially. For this reason, we have used the term *semilinear* in the name of the resulting problem.

2.3 QCMA-complete problem

The construction from Section 2.2 can be modified to generate a $\mathsf{QCMA}_1^{\mathcal{G}_8}$-complete problem. Moreover, since $\mathsf{QCMA}_1^{\mathcal{G}_8} = \mathsf{QCMA}$ [27], this results in a QCMA-complete problem. Although there are already many problems known to be complete for this class [19, 42, 21, 24, 40], none of them are strong QCSPs.[9]

In Section 2.1, we argued that the unconstrained or "free" logical qudits of an instance allowed one to guess what state of these qudits (the witness state) might satisfy the instance. This freedom made the problem more difficult and thus not likely contained in BQP_1. For this reason, we introduced the undefined state $|?\rangle$, which simplified these instances and made them decidable in BQP_1. In this construction, we seek to construct a problem that sits in between these two classes so it is QCMA-complete. To accomplish this, we desire to have "free" logical qudits to accommodate a witness state that helps verify whether the instance is satisfiable, but have some sort of constraint to demand that the state is classical.[10]

In practice, creating these constraints is challenging since any superposition of two satisfying states will also satisfy the clause. Instead, we set the constraints such that if there exists a quantum witness state that is part of a satisfying state, there is also a classical

[9] While Ref. [40] also defines a QCMA-complete QSAT problem, it requires additional promise conditions.
[10] We continue using the undefined state for logical qudits whose initial state is not constrained so the difficulty of the problem does not become QMA.

witness state. Loosely, we accomplish this by defining new *witness qudits* and create a new constraint $\Pi_{init}^{|00,11\rangle}$ that connects a witness qudit with a logical one, and require that they are both either $|00\rangle$, $|11\rangle$, or in a superposition of the two.[11] In this way, the two qudits are partially "free" as there is some freedom to their state, yet posses some desired structure. Importantly, we ensure that the witness qudits do not form part of the computation after this initial point.

To see why this leads to the desired effect, consider the toy instance of Figure 3 and suppose there exists a state $|\psi_{\text{wit}}\rangle$ of the four "free" qudits that leads to a satisfying state. Observe that to satisfy the $\Pi_{init}^{|00,11\rangle}$ clauses, this state must be of the form $|\psi_{\text{wit}}\rangle = (\alpha_{00}|0000\rangle + \alpha_{01}|0011\rangle + \alpha_{10}|1100\rangle + \alpha_{11}|1111\rangle)_{L_1,W_1,L_2,W_2}$ with $\sum_{b\in\{0,1\}^2}|\alpha_b|^2 = 1$, which we can rewrite in a more convenient form as $|\psi_{\text{wit}}\rangle = \sum_{b\in\{0,1\}^2}\alpha_b|b\rangle_L \otimes |b\rangle_W$. Then, a state that satisfies all clauses of the instance is the history state

$$\begin{aligned}|\psi_{hist}\rangle &= \frac{1}{\sqrt{5}}\sum_{t=0}^{4}[U_t\dots U_0|00\rangle\otimes|\psi_{\text{wit}}\rangle]\otimes|\underbrace{d\dots d}_{t}a_t\underbrace{r\dots r}_{4-t}\rangle \\ &= \frac{1}{\sqrt{5}}\sum_{t=0}^{4}\sum_{b\in\{0,1\}^2}\alpha_b|\xi_b^t\rangle\otimes|b\rangle_W\otimes|\underbrace{d\dots d}_{t}a_t\underbrace{r\dots r}_{4-t}\rangle,\end{aligned} \tag{1}$$

where $|\xi_b^t\rangle := U_t\dots U_0|00\rangle\otimes|b\rangle_L$. Now, let us argue that there is also a classical witness that leads to a satisfying history state. First, observe that any basis state $|b\rangle_L\otimes|b\rangle_W$ with $|\alpha_b|>0$ from the decomposition of the witness satisfies the $\Pi_{init}^{|00,11\rangle}$ clauses. Consequently, the history state above but with initial state $|00\rangle\otimes|b\rangle_L\otimes|b\rangle_W$ satisfies the $\Pi_{init}^{|0\rangle}$, $\Pi_{init}^{|00,11\rangle}$, and Π_{prop} clauses of the instance. Finally, to show that this state also satisfies the Π_{out} clause, recall that this clause is satisfied if at time $t=4$, the probability that the second qubit yields outcome "1" when measured is 1. As shown in Equation (3), this probability can be written as

$$\Pr(\text{outcome } 1) = \sum_{b,b'\in\{0,1\}^2}\alpha_b\alpha_{b'}^*\langle\xi_{b'}^4|\otimes\langle b'|\Pi^{(1)}|\xi_b^4\rangle\otimes|b\rangle = \sum_{b\in\{0,1\}^2}|\alpha_b|^2\langle\xi_b^4|\Pi^{(1)}|\xi_b^4\rangle$$

where $\Pi^{(1)} := |1\rangle\langle 1|_2\otimes I_{rest}$, and in the last equality we observed that $\langle b'|b\rangle = \delta_{b,b'}$. Then, by the assumption that the instance is satisfiable, it must be that $\langle\xi_b^4|\Pi^{(1)}|\xi_b^4\rangle = 1$ for all basis states $|b\rangle$ with $|\alpha_b|>0$. This can also be understood as the probability that at the end of the circuit, the second qubit yields outcome "1" when the witness is the basis state $|b\rangle\otimes|b\rangle$. Therefore, the history state of Equation (1) with classical witness $|\psi_{\text{wit}}\rangle = |b\rangle_L\otimes|b\rangle_W$ also satisfies all clauses of the instance.

The remaining parts of the proof require showing that all new possible qudit connections with new $\Pi_{init}^{|00,11\rangle}$ clause can still be handled, as well as demonstrate perfect completeness and bounded soundness of the hybrid algorithm. For the latter, the majority of the arguments from the BQP_1 construction also directly apply here.

2.4 coRP-complete problem

In Section 2.1, we mentioned that the satisfiability of some instances is decided through a quantum circuit. In particular, this circuit was used to verify the satisfiability of the simultaneous Π_{prop} clauses and final Π_{out} clauses. For the latter, the circuit executed

[11] The local Hilbert space is then 8-dimensional, as it composed of a 3-dimensional data subspace, a 3-dimensional clock subspace, and a 2-dimensional witness subspace.

the quantum circuit $U_T \dots U_1$ on input $|0\rangle^{\otimes q}$, measured some of the qubits, and accepted or rejected depending on the measurement outcomes. Intuitively, to generate the coRP-complete problem, we would like to replace the universal quantum circuit by a universal classical *reversible* circuit $R = R_T \dots R_1$ (reversibility is needed since the best potentially satisfying state is still a *quantum* history state) and introduce randomness into the instance by initializing p ancilla qubits to $|+\rangle$. Then, for these new sub-instances, we could analogously verify the Π_{out} clauses by sampling a bitstring $b \in \{0,1\}^p$, evaluating the circuit R on input $(0^q, b)$ and deciding on its satisfiability based on the final state of the bits. While this idea is close to the actual construction, for reasons mentioned in the full version of the text, it is not sufficient to decide all instances.

For the construction to work, we also incorporate elements of the QCMA problem from the previous section. Namely, we modify the $\Pi_{init}^{|00,11\rangle}$ clause so it initializes both a witness (now referred to as *auxiliary* qudit as we remove the freedom) and a logical qudit to the maximally entangled state $|\Phi^+\rangle$. This new clause is denoted $\Pi_{init}^{|\Phi^+\rangle}$.

Again using the toy example of Figure 3 (replacing the $\Pi_{init}^{|00,11\rangle}$ clauses by $\Pi_{init}^{|\Phi^+\rangle}$ clauses and the unitaries U_i by reversible classical gates R_i), let us illustrate how this construction allows us to verify the satisfiability of Π_{out} clauses. If the instance is satisfiable, the satisfying state must be the history state

$$\begin{aligned} |\psi_{hist}\rangle &= \frac{1}{\sqrt{5}} \sum_{t=0}^{4} \left[R_t \dots R_0 |00\rangle \otimes |\Phi^+\rangle^{\otimes 2} \right] \otimes |\underbrace{d \dots d}_{t} a_t \underbrace{r \dots r}_{4-t}\rangle \\ &= \frac{1}{\sqrt{5}} \sum_{t=0}^{4} \sum_{b \in \{0,1\}^2} \frac{1}{2} |\xi_b^t\rangle \otimes |b\rangle_{Aux} \otimes |\underbrace{d \dots d}_{t} a_t \underbrace{r \dots r}_{4-t}\rangle , \end{aligned}$$

where in the second line we observed that $|\Phi^+\rangle^{\otimes p} = 2^{-\frac{p}{2}} \sum_{b \in \{0,1\}^p} |b\rangle_L \otimes |b\rangle_{Aux}$ for any $p \in \mathbb{N}$, and defined $|\xi_b^t\rangle := R_t \dots R_1 |00\rangle \otimes |b\rangle_L$. The Π_{out} clause is satisfied if at time $t = 4$, the probability that the second qubit yields outcome "1" when measured is 1. This probability is given by

$$\Pr(\text{outcome } 1) = \frac{1}{4} \sum_{b,b' \in \{0,1\}^2} \langle \xi_{b'}^4| \otimes \langle b'| \, \Pi^{(1)} \, |\xi_b^4\rangle \otimes |b\rangle = \frac{1}{4} \sum_{b \in \{0,1\}^2} \langle \xi_b^4| \, \Pi^{(1)} \, |\xi_b^4\rangle ,$$

from where it is evident that if the instance is satisfiable, $\langle \xi_b^4| \Pi^{(1)} |\xi_b^4\rangle = 1$ for all $b \in \{0,1\}^2$. Hence, it is possible to verify the Π_{out} clause by sampling one of the strings b, running circuit R on input $(0^2, b)$, and measuring the state of the second qubit.

Another important consideration is that the classical reversible gate set must be chosen with care. Although not covered in this version of the paper, we usually desire that $\mathcal{G}$ is a gate set such that all gates in the set change the basis states upon application, and so $V(\Pi_i)$ of Equation (2) can be implemented perfectly with gates from this set. Here, only the first property is relevant. We choose $\mathcal{G} = \{X, (X \otimes X \otimes X)\text{Toffoli}\}$, which clearly satisfies this property and is also a universal gate set for reversible classical computation. As a consequence, the QSAT problem of this section has 5-local clauses since the Π_{prop} clauses may use a Toffoli. This is the best locality we can achieve as it is also well known that any universal gate set for reversible quantum computation must include a 3-bit gate.

2.5 Universality of qubits for QCSPs

In previous sections, we showed that there are QSAT problems acting on qudits that are complete $\mathsf{BQP}_1^{\mathcal{G}_8}$, QCMA, and coRP. Here, we refine these statements and show that there are QSAT problems on qubits (albeit with higher locality) that are also complete for these classes. To achieve this, we show that any QCSP on qudits can be reduced to another QSAT problem on qubits using little computational power. We note that this section, apart from some changes in the exposition, stems directly from Ref. [32].

At first glance, this statement may seem trivial as operations on qubits are universal for quantum computation, i.e. we can emulate a d-qudit with a $\lceil \log_2(d) \rceil$ qubits and carry out unitaries on those qubits. For our QSAT problems, it is true that any instance retains its satisfiability status when expressed in terms of qubits. However, it is not clear if all input instances generated with these new qubit clauses are contained within this class. For a successful reduction, we must have both.

For an even more explicit example, let us first represent the basis clock states using qubits as: $|r\rangle := |00\rangle$, $|a\rangle := |01\rangle$, and $|d\rangle := |10\rangle$. The $\Pi_{start} = |r\rangle\langle r|$ clause (defined exactly as P_{init} in Equation (7)) can now be written as $\Pi_{start} = |00\rangle\langle 00| + |11\rangle\langle 11|$, where the last term is to prevent the fourth basis state $|11\rangle$, which did not exist before. The clause $(|00\rangle\langle 00| + |11\rangle\langle 11|)_{1,2} + (|00\rangle\langle 00| + |11\rangle\langle 11|)_{2,3}$ acting on three qubits is now valid and is satisfied by the state $|0_1 0_2 1_3\rangle$. This state, however, presents some ambiguity: either we have $|r\rangle$ on qubits 1 and 2, or $|a\rangle$ on qubits 2 and 3. In general, decomposing all clauses into qubits and considering all input instances that may occur adds a significant level of complexity to the problem, making it difficult to determine if it remains in the same class. Moreover, we remark that this is not only particular to our QSAT problems, but in fact applies to all CSPs and QCSPs defined on qudits or non-Boolean variables! In general, the issue is that we cannot ensure that the new qubit clauses are applied to qubits in a consistent fashion based on its parent qudit problem. For example, a qubit clause might treat a particular qubit as “qubit 1” of a previous d-qudit, while another clause might refer to the same qubit as “qubit 2”. Moreover, the qubit clauses could also “mix and match”, combining “qubit 1” from one previous d-qudit with “qubit 2” from another d-qudit (as in the example with Π_{start}). Overall, these lead to constraints that were unrealizable in the parent qudit problem.

Our main result of this section shows that with a more clever reduction than directly decomposing a d-qudit into $\lceil \log_2(d) \rceil$ qubits, we can guarantee that a satisfiable/unsatisfiable instance on qubits maps to one on qudits with the same satisfiability status. This is something that is not known to be possible classically! More formally, we show that

▶ **Theorem 10** (Theorem 6; formal)**.** *For any QCSP $\mathcal{C}$ on d-qudits, there is another QCSP $\mathcal{C}'$ on qubits, and AC^0 circuits f and g, such that f reduces $\mathcal{C}$ to $\mathcal{C}'$, and g reduces $\mathcal{C}'$ to $\mathcal{C}$. If $\mathcal{C}$ is k-local, then $\mathcal{C}'$ can be chosen to be $4 \cdot 2^{\lceil \log_2(\lceil \log_2(d) \rceil) \rceil} k$ local (that is, $O(\log(d))$ times larger.)*

The main idea behind the proof is that in the quantum world, we can fix the issues mentioned above by again using monogamy of entanglement to bind together our constituent qubits into ordered, entangled larger systems. Ultimately, each clause in the resulting qubit problem incorporates new projectors that force a particular ordering of qubits, and any two clauses that try to “mix and match”, or use the same set of qubits but with different ordering, are necessarily frustrated.

If in Theorem 10 we do not require the reductions to be in AC^0, and instead allow P-reductions, a locality of $4\lceil \log_2(d) \rceil k$ suffices. This is used in Corollary 7.

2.6 Direct sum and direct product problems

There is a notion of *direct sum* (denoted by “$\oplus$”) and *direct product* (denoted by “$\otimes$”) on CSPs and QCSPs that allow us to define the remaining six complete problems. To be clear, these are operations on languages themselves, not on instances. For example, we can talk about the languages 3-Colorable $\oplus$ 4-SAT and 3-Colorable $\otimes$ 4-SAT. Although this notion appears to be quite natural, we were unable to find many sources discussing such ideas – possibly because the classical theory is not as exciting, for reasons we also discuss. The relevant task here is to demonstrate that sum and product QCSPs inherit completeness properties from their constituents. In this way, we are able to construct QCSPs that are complete for PI and SoPU classes, defined as follows.

▶ **Definition 11** (Pairwise intersection of classes). *If $\mathcal{C}_1$ and $\mathcal{C}_2$ are two complexity classes (any sets of languages), then $\mathsf{PI}(\mathcal{C}_1, \mathcal{C}_2)$ is the class that denotes the* pairwise intersection *of $\mathcal{C}_1$ and $\mathcal{C}_2$. In other words, it is the class of languages that can be written as the intersection, i.e. the logical AND, of a language in $\mathcal{C}_1$ and a language in $\mathcal{C}_2$.*

▶ **Definition 12** (Star of pairwise unions of classes). *If $\mathcal{C}_1$ and $\mathcal{C}_2$ are two complexity classes, then their* star of pairwise unions*, denoted $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2)$, is a complexity class defined as follows: for each language $L_1 \in \mathcal{C}_1$ and $L_2 \in \mathcal{C}_2$, let d be a fresh symbol that is not in the alphabet of L_1 or L_2. Then, the language $(dL_1|dL_2)^*$ is in $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2)$. $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2)$ is the closure of all such languages under L (logspace reductions).*

Definition 12 merits a brief explanation. For a pair of languages L_1 and L_2, what do the strings in the language $L := (dL_1|dL_2)^*$ look like? Given an input string like $d010011d101101d101001$, it will belong to L if and only if each of the three bitstrings $\{010011, 101101, 101001\}$ belongs to either L_1 or L_2. If C is a complexity class powerful enough to break apart the individual bitstrings from the d-delimited string, as well decide both L_1 and L_2, then $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2) \in C$.

We begin discussing that there are CSPs that are complete for these classes, and then extend this to the quantum setting since the latter follows almost identically.

2.6.1 Direct product of constraint satisfaction problems

To begin, it is useful to recall the precise definition of a CSP. A *constraint satisfaction problem* is a triple (V, D, C), where $V = \{v_1, \ldots, v_n\}$ is a finite set of variables, each taking a value from the domain D. If the domain is $D = \{0, 1\}$, then we have a Boolean CSP, and can be generalized to dits if D is instead $D = \{0, \ldots, d\}$. C is a set of constraints, where each constraint $c \in C$ restricts the values that a subset of the variables may take.

Now, let L_1 and L_2 be two CSPs with domains D_1 and D_2, and allowed constraints C_1 and C_2, respectively.

▶ **Definition 13** (Direct product of CSPs). *Given the CSPs L_1 and L_2, their* direct product *$L_1 \otimes L_2$ is a CSP whose domain is the Cartesian product $D_1 \times D_2$. Each constraint $c_i \in C_1$ (resp. C_2) of locality k leads to a constraint c_i' in $L_1 \otimes L_2$, also of locality k, as follows. A tuple $(v_1, v_2, \ldots, v_k) \in (D_1 \times D_2)^k$, where each entry $v_i = (v_{i,1}, v_{i,2})$, belongs to c_i' if the tuple $(v_{i,1}, \ldots, v_{k,1})$ belongs to c_i. Each constraint in $L_1 \otimes L_2$ arises this way from a constraint in L_1 or L_2.*

The goal of this subsection is to show that when one CSP is complete for a complexity class A, and another is complete for a class B, the product problem is complete for the complexity class $\mathsf{PI}(A, B)$. Formally,

▶ **Theorem 14** (Completeness of direct products for PI). *Let M be a set of functions closed under composition with local functions, and closed under concatenations (i.e. if some f, $g : \Sigma_1^* \to \Sigma_2^*$ are each in M, then $h : x \to f(x)g(x)$ is as well). Let L_1 be a CSP complete under*

M-reductions for a class $\mathcal{C}_1$, and likewise L_2 be complete for $\mathcal{C}_2$. Assume that each of $\mathcal{C}_1$ and $\mathcal{C}_2$ are closed under reductions by local functions, closed under intersections, and contain the language ALL *of all strings, Σ^*. Then, the direct product $L_1 \otimes L_2$ is complete under M-reductions for* $\mathsf{PI}(\mathcal{C}_1, \mathcal{C}_2)$.

The assumptions in this theorem are mild and satisfied even for AC^0-reductions and most complexity classes. In other words, this theorem essentially states that if we have CSPs complete for "reasonable" classes $\mathcal{C}_1$ and $\mathcal{C}_2$, the product CSP is complete for $\mathsf{PI}(\mathcal{C}_1, \mathcal{C}_2)$.

2.6.2 Direct sum of constraint satisfaction problems

If direct products let us express (informally) a "two-input logical AND" of two CSPs, then direct sums let us express "unbounded-fanin AND of fanin-2 ORs".

▶ **Definition 15** (Direct sum of CSPs). *Given the CSPs L_1 and L_2, their* direct sum *$L_1 \oplus L_2$ is a CSP whose domain is the disjoint union $D_1 \uplus D_2$. Each constraint in $L_1 \oplus L_2$ is either of the form $c_i \cup (D_2^k)$, where $c_i \in C_1$ is a constraint of locality k; or it is $c_i \cup (D_1^k)$ for some $c_i \in C_2$.*

To better understand this definition, consider an instance of $L_1 \oplus L_2$ with a single connected component, and assume that it is satisfiable.[12] The definition of the problem and this assumption imply that any satisfying state must either have all variables set to values from D_1, or all of them must be from D_2. Then, for a general instance of $L_1 \oplus L_2$, solving the problem amounts to identifying all of the connected components, and for each one determine whether it can be satisfied entirely from values of D_1 or D_2. The instance is satisfiable iff all components are as well.

One might expect that, by analogy with the direct product, the sum of CSPs should then be complete for the pairwise union of two classes, $\mathsf{PU}(A, B)$. This would be true if we only had to worry about problems that formed a single connected component, which is not the case. This is why we must define the "star of pairwise unions" as in Definition 12. The goal of this section is to demonstrate this fact formally.

▶ **Theorem 16** (Completeness of direct sums for SoPU). *Let be M a set of functions closed under composition with logspace-computable functions (such as the set of logspace functions themselves,* FL*). Let L_1 be a CSP complete under M-reductions for a class $\mathcal{C}_1$, and likewise L_2 be M-complete for $\mathcal{C}_2$. Assume that each of $\mathcal{C}_1$ and $\mathcal{C}_2$ are closed under M-reductions, and contain the language* NONE *of no strings, $\emptyset$. Then, the direct sum $L_1 \oplus L_2$ is complete under M-reductions for* $\mathsf{SoPU}(\mathcal{C}_1, \mathcal{C}_2)$.

2.6.3 Quantum sums and products

The constructions above transfer in a very natural way to the quantum setting. Now, instead of domains that are a Cartesian product or disjoint union, the Hilbert spaces are a tensor product or direct sum. The clauses are accordingly built as tensor products and direct sums.

▶ **Definition 17** (Direct product of QCSPs). *Given the QCSPs L_1 and L_2, their* direct product *$L_1 \otimes L_2$ is a QCSP whose domain is the tensor product Hilbert space $D_1 \otimes D_2$. Each operator H_i in L_1 leads to an operator $H_i \otimes I$, a tensor product with the identity, and likewise for L_2.*

[12] A *connected component* of a CSP is a connected component in the graph for that CSP, where the vertices are variables, and there is an edge between variables if they share a constraint.

▶ **Definition 18** (Direct sum of QCSPs). *Given the QCSPs L_1 and L_2, their* direct sum *$L_1 \uplus L_2$ is a QCSP whose domain is the direct sum Hilbert space $D_1 \oplus D_2$. Each operator H_i in L_1 leads to an operator $H_i \oplus 0$, a direct sum with the 0 operator, requiring that a frustration-free state lies in the null space of H_i* or *the right half of the direct sum (or a linear combination). Likewise for operators in L_2.*

These have the same essential properties as the direct product and sum for classical CSPs, where we can produce product and sum instances that are satisfiable iff both (resp. either) of the original instances are satisfiable.

Given the discussion above on languages and strings of symbols, one might think that we must talk about quantum states and concatenations of strings of qubits. This is not the case. The strings of symbols are just the encoding of the constraints, which are classical data even for a QCSP. The only quantum-specific requirements involve checking that tensor products or embeddings of satisfying states yield another satisfying state; and the appropriate converse properties. These follow directly from the definition of tensor products and direct sums.

2.6.4 Basic class properties

Here, we state some basic properties of general $\mathsf{PI}(A, B)$ and $\mathsf{SoPU}(A, B)$ classes.

▶ **Lemma 19.** *If the class B includes the language* ALL *of all strings, then $A \subseteq \mathsf{PI}(A, B)$. Similarly, if B includes the language* NONE *of no strings, then $A \subseteq \mathsf{SoPU}(A, B)$.*

▶ **Lemma 20.** PI *and* SoPU *respect the inclusion order of complexity classes. That is, $A \subseteq C$ and $B \subseteq D$ implies $\mathsf{PI}(A, B) \subseteq \mathsf{PI}(C, D)$ and $\mathsf{SoPU}(A, B) \subseteq \mathsf{SoPU}(C, D)$.*

SoPU generally leads to a more powerful class than PI, that is:

▶ **Lemma 21.** *If classes A and B are closed under reductions by local functions, then $\mathsf{PI}(A, B) \subseteq \mathsf{SoPU}(A, B)$.*

This is apparent from the definition of these classes since SoPU is also required to compute the AND of multiple inputs. It is also true that PI and SoPU do not increase the power of classes by combining a class A with something weaker. Formally:

▶ **Lemma 22.** *If A is closed under intersection, and $B \subseteq A$, then $\mathsf{PI}(A, B) \subseteq A$. Moreover, if A is closed under logspace reductions, unions, and delimited concatenation, then $\mathsf{SoPU}(A, B) \subseteq A$.*

2.7 New complete problems

As mentioned previously, while the notion of product and sum of constraint problems seems natural, classical constraint problems do not seem to offer such a rich theory. This is due to the fact that most classes with complete CSPs are contained within each other. Indeed, Allender *et al.*'s refinement of Schaefer's dichotomy theorem states that all Boolean CSPs are either in co-NLOGTIME; or are complete for L, NL, ⊕L, P or NP under AC^0 reductions [3]. With the exception of NL and ⊕L, all possible pairs from this list have an obvious containment relation, so the only nontrivial consequence would be that there exists a CSP, on a domain of size four, that is $\mathsf{PI}(\oplus\mathsf{L}, \mathsf{NL})$-complete under AC^0 reductions. However, under the more common P-reductions, the complexity of these problems becomes either in P or NP-complete.[13] Then,

[13] The same is true for CSPs defined on qudits [43].

since $\mathsf{P} \subseteq \mathsf{NP}$ and these classes meet all properties discussed in Lemma 22, it follows that products or sums of these problems result in complexity classes that are trivially equal to NP. For example, 2-SAT $\oplus$ 3-SAT and 2-SAT $\otimes$ 3-SAT are complete problems for $\mathsf{PI}(\mathsf{P}, \mathsf{NP})$ and $\mathsf{SoPU}(\mathsf{P}, \mathsf{NP})$, but these classes are trivially equal to NP.

This is no longer the case in this work. The seven classes we have discussed so far which have complete CSPs are P, coRP, BQP_1, NP, QCMA, and QMA_1. Importantly, these all have the closure properties discussed in this section so far: union, intersection, logspace reductions, and delimited concatenation; and they all include the trivial problems ALL and NONE. Among these classes, most pairs $\{A, B\}$ have $A \subseteq B$, in which case $\mathsf{PI}(A, B) = \mathsf{SoPU}(A, B) = B$. However, there are three pairs that are not known to contain each other, these are: $\mathsf{coRP} \overset{?}{\subseteq} \mathsf{NP}$, $\mathsf{BQP}_1 \overset{?}{\subseteq} \mathsf{NP}$, and $\mathsf{BQP}_1 \overset{?}{\subseteq} \mathsf{MA}$. Each of these pairs leads to two new classes $\mathsf{PI}(A, B)$ and $\mathsf{SoPU}(A, B)$, that are not obviously equal to some other known class. Together, we obtain six more complexity classes with complete QCSPs.

2.7.1 Relations to other classes

Notably, the pair coRP and NP involves only classical classes, and accordingly there is more theory already developed around them.

From the lemmas stated earlier in this section, one can show that $\mathsf{NP} \subseteq \mathsf{PI}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{MA}$. In addition to this, we can relate $\mathsf{PI}(\mathsf{coRP}, \mathsf{NP})$ to the class $\mathsf{DP} := \mathsf{PI}(\mathsf{NP}, \mathsf{coNP})$ studied in Ref. [34]. This class forms the second layer of the *boolean hierarchy* BH, i.e. $\mathsf{DP} = \mathsf{BH}_2$ [13]. Since $\mathsf{coRP} \subseteq \mathsf{coNP}$, Lemma 20 tells us that $\mathsf{PI}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{DP}$. On the other hand, $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$ does not obviously lie in the Boolean hierarchy. If the class was a simple pairwise union (instead of the "*star* of pairwise unions"), it would lie in BH_3 – the pairwise union of DP and NP. However, it seems unlikely that $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$ falls within this class, as doing so would require showing that one could condense the long list of checks required to decide a SoPU problem down to only two queries. In this line of thought, we know that queries to an NP oracle do not need to depend on each other adaptively, so $\mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP})$ is contained in $\mathsf{P}^{||\mathsf{NP}} = \mathsf{P}^{\mathsf{NP}[\log]}$, studied in Refs. [11, 25].[14]

These classes are also related to two interesting collapse statements. First, observe that if $\mathsf{P} = \mathsf{RP}$ (derandomization), then $\mathsf{coRP} = \mathsf{P} \subseteq \mathsf{NP}$ and so $\mathsf{NP} \subseteq \mathsf{PI}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP}) = \mathsf{SoPU}(\mathsf{NP}, \mathsf{NP}) = \mathsf{NP}$. Moreover, an even weaker version of derandomization where $\mathsf{NP} = \mathsf{MA}$ would also lead to a collapse. Here, since $\mathsf{coRP} \subseteq \mathsf{MA}$, we have $\mathsf{NP} \subseteq \mathsf{PI}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{SoPU}(\mathsf{coRP}, \mathsf{NP}) \subseteq \mathsf{SoPU}(\mathsf{MA}, \mathsf{NP}) = \mathsf{SoPU}(\mathsf{NP}, \mathsf{NP}) = \mathsf{NP}$. Second, we see that if $\mathsf{NP} = \mathsf{coNP}$ (concise refutations), then $\mathsf{coRP} \subseteq \mathsf{coNP} = \mathsf{NP}$.

For the PI and SoPU classes that involve BQP_1, NP, or MA, it seems difficult to state other inclusions, besides the fact that they lie above BQP_1 and below QCMA.

3 Discussion and open questions

Perhaps the most interesting points of discussion are the implications of Corollary 9. In the latter case, if a complete classification theorem for QSAT problems shows that there are fewer than 13 classes, this would present exciting implications as equalities between some of these classes tackle many interesting and open questions (see Figure 1). This is true even for adjacent classes. For instance, $\mathsf{P} = \mathsf{coRP}$ would imply that probabilistic algorithms

[14] $\mathsf{P}^{||\mathsf{NP}}$ is the class of problems that can be solved by a P machine with polynomially many nonadaptive NP queries, or alternatively, logarithmically many adaptive queries.

with perfect completeness can be derandomized, and $\mathsf{QCMA} = \mathsf{QMA}_1$ would imply that any quantum-verifiable problem (with perfect completeness) could be verified using a classical witness state. Even for the PI and SoPU classes defined here, we have that if $\mathsf{PI}(A, B) = A$, then $B \subseteq A$. As mentioned in Section 2.7, it is expected through derandomization conjectures that some of these classes are in fact equal to each other. Even if this classification theorem proves any of these conjectures, it would be a great result since such proofs have eluded us for many decades. In the former case of Corollary 9, a classification showing that there are more than 13 classes would be a stark contrast with classical CSPs, which can be completely classified as being either in P or NP-complete [38, 43]. This would highlight the more rich and complex panorama of strong QCSPs, and establish a larger repertoire of problems from which to construct reductions and potentially describe the complexity of other problems.

This last point also raises the question whether there could be other classes with complete QSAT problems. Considering those corresponding to polynomial-time computation and verification, we think that this is unlikely. For example, we have not mentioned complete QCSPs for BPP, BQP, or QMA. Since $\mathsf{coRP} \subseteq \mathsf{BPP}$, $\mathsf{BQP}_1 \subseteq \mathsf{BQP}$, and $\mathsf{QMA}_1 \subseteq \mathsf{QMA}$, there are clearly strong QCSPs in these classes. However, the challenge lies in proving their hardness: as shown in Section A.3, these proofs usually require encoding a probabilistic circuit into an instance of this problem. As is also shown there, perfect completeness is critical for the construction, and thus does not work for a circuit with two-sided error. Adressing this would require a different technique.[15] A positive resolution could arise if these classes admit a scheme that boosts their acceptance probabilities to 1. Jordan *et al.* [27] showed that this was possible for QCMA (demonstrating that $\mathsf{QCMA} = \mathsf{QCMA}_1$), but whether this is possible for BPP, BQP, or QMA remains an open question. Another set of classes we have not considered, are those with no error. Little is known about these classes as they appear to be extremely difficult to work with since the perfect soundness requirement implies that no incorrect instance is ever accepted. Besides classes related to polynomial-time computation and verification, there could be other classes with complete QCSPs. After all, the complexity class landscape is vast.

In Theorems 2 and 5, we describe two new types of QSAT problems that can be solved efficiently with a quantum or probabilistic classical computer. Unfortunately, the projectors used in these problems are artifacts built to achieve these results and do not immediately correspond to QSAT problems of interest, even in the qubit case. Recent developments in the fields of quantum chemistry [5], high-energy physics [35] and nuclear physics [10, 16, 17] have shown that 3- or 4-local Hamiltonians are sometimes necessary to explain emergent physics. The QSAT problems for these Hamiltonians are not immediately tractable as $k \geq 3$, so it would be exciting to determine if these problems, or others, fall within these complexity classes. We hope that having demonstrated that such problems exist, our results inspire others to search for more relevant cases.

Finally, Theorem 8 adds an additional six classes to the set of classes with strong QCSP complete problems. Beyond the inclusions shown in Figure 1, little is known about them. It would thus be interesting to investigate how these classes relate to others, and which other problems fall within them.

[15] Another technique is to reduce an already known hard problem into an instance of the target problem. For the LH problem, this is done via perturbation theory gadgets [28, 39, 18]. However, these gadgets rely on approximations and hence do not preserve perfect completeness.

References

1 Scott Aaronson. On perfect completeness for QMA. *Quantum Inf. Comput.*, 9(1&2):81–89, 2009. doi:10.26421/QIC9.1-2-5.

2 Dorit Aharonov, Oded Kenneth, and Itamar Vigdorovich. On the Complexity of Two Dimensional Commuting Local Hamiltonians. In *13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2018)*, volume 111 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:21, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.TQC.2018.2.

3 Eric Allender, Michael Bauland, Neil Immerman, Henning Schnoor, and Heribert Vollmer. The complexity of satisfiability problems: Refining schaefer's theorem. *Journal of Computer and System Sciences*, 75(4):245–254, June 2009. doi:10.1016/j.jcss.2008.11.001.

4 Matthew Amy, Andrew N. Glaudell, Shaun Kelso, William Maxwell, Samuel S. Mendelson, and Neil J. Ross. Exact synthesis of multiqubit clifford-cyclotomic circuits. In *Reversible Computation - 16th International Conference, RC 2024, Toruń, Poland, July 4-5, 2024, Proceedings*, volume 14680 of *Lecture Notes in Computer Science*, pages 238–245. Springer, 2024. doi:10.1007/978-3-031-62076-8_15.

5 Alberto Baiardi, Michał Lesiuk, and Markus Reiher. Explicitly correlated electronic structure calculations with transcorrelated matrix product operators. *Journal of Chemical Theory and Computation*, 18(7):4203–4217, 2022. doi:10.1021/acs.jctc.2c00167.

6 Sergey Bravyi. Efficient algorithm for a quantum analogue of 2-sat, 2006. arXiv:quant-ph/0602108.

7 Sergey Bravyi, Arvid J. Bessen, and Barbara M. Terhal. Merlin-arthur games and stoquastic complexity, 2006. arXiv:quant-ph/0611021.

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

9 Sergey Bravyi and Mikhail Vyalyi. Commutative version of the local hamiltonian problem and common eigenspace problem. *Quantum Info. Comp.*, 5:187–215, May 2005. doi:10.26421/QIC5.3-2.

10 J. H. Busnaina, Z. Shi, Jesús M. Alcaine-Cuervo, Cindy X. Yang, I. Nsanzineza, E. Rico, and C. M. Wilson. Native three-body interactions in a superconducting lattice gauge quantum simulator, 2025. arXiv:2501.13383.

11 Samuel R. Buss and Louise Hay. On truth-table reducibility to sat. *Information and Computation*, 91(1):86–102, March 1991. doi:10.1016/0890-5401(91)90075-d.

12 Chris Cade, Marten Folkertsma, and Jordi Weggemans. Complexity of the guided local hamiltonian problem: Improved parameters and extension to excited states, 2024. arXiv:2207.10097.

13 Jin-Yi Cai, Thomas Gundermann, Juris Hartmanis, Lane A. Hemachandra, Vivian Sewelson, Klaus Wagner, and Gerd Wechsung. The boolean hierarchy i: Structural properties. *SIAM Journal on Computing*, 17(6):1232–1252, December 1988. doi:10.1137/0217078.

14 Ricardo Rivera Cardoso, Alex Meiburg, and Daniel Nagaj. Quantum sat problems with finite sets of projectors are complete for a plethora of classes, 2025. arXiv:2506.07244.

15 Andrew M. Childs, David Gosset, and Zak Webb. The bose-hubbard model is qma-complete. *Theory Comput.*, 11:491–603, 2015. doi:10.4086/TOC.2015.V011A020.

16 Alexander Y. Chuang, Huan Q. Bui, Arthur Christianen, Yiming Zhang, Yiqi Ni, Denise Ahmed-Braun, Carsten Robens, and Martin W. Zwierlein. Observation of a halo trimer in an ultracold bose-fermi mixture, 2024. arXiv:2411.04820.

17 R. Cruz-Torres, D. Nguyen, F. Hauenstein, A. Schmidt, S. Li, D. Abrams, H. Albataineh, S. Alsalmi, et al. Probing few-body nuclear dynamics via ^{3}H and ^{3}He $(e, e^{'}p)$pn cross-section measurements. *Phys. Rev. Lett.*, 124:212501, 2020. doi:10.1103/PhysRevLett.124.212501.

18 Toby Cubitt and Ashley Montanaro. Complexity classification of local hamiltonian problems. *SIAM Journal on Computing*, 45(2):268–316, 2016.

19 Sevag Gharibian and Jamie Sikora. Ground state connectivity of local hamiltonians. *ACM Trans. Comput. Theory*, 10(2):8:1–8:28, 2018. doi:10.1145/3186587.

20 Brett Giles and Peter Selinger. Exact synthesis of multiqubit clifford+t circuits. *Phys. Rev. A*, 87:032332, March 2013. doi:10.1103/PhysRevA.87.032332.

21 David Gosset, Jenish C. Mehta, and Thomas Vidick. QCMA hardness of ground space connectivity for commuting Hamiltonians. *Quantum*, 1:16, July 2017. doi:10.22331/q-2017-07-14-16.

22 David Gosset and Daniel Nagaj. Quantum 3-sat is QMA_1-complete. *SIAM Journal on Computing*, 45(3):1080–1128, 2016. doi:10.1137/140957056.

23 Sean Hallgren, Daniel Nagaj, and Sandeep Narayanaswami. The local hamiltonian problem on a line with eight states is qma-complete. *Quantum Inf. Comput.*, 13(9-10):721–750, 2013. doi:10.26421/QIC13.9-10-1.

24 Ying hao Chen. 2-local hamiltonian with low complexity is qcma, 2019. arXiv:1909.03787.

25 Lane A. Hemachandra. The strong exponential hierarchy collapses. *Journal of Computer and System Sciences*, 39(3):299–322, December 1989. doi:10.1016/0022-0000(89)90025-1.

26 Sandy Irani and Jiaqing Jiang. Commuting local hamiltonian problem on 2d beyond qubits, 2023. arXiv:2309.04910.

27 Stephen P. Jordan, Hirotada Kobayashi, Daniel Nagaj, and Harumichi Nishimura. Achieving perfect completeness in classical-witness quantum merlin-arthur proof systems. *Quantum Info. Comput.*, 12(5-6):461–471, 2012. doi:10.26421/QIC12.5-6-7.

28 Julia Kempe, Alexei Kitaev, and Oded Regev. *The Complexity of the Local Hamiltonian Problem*, pages 372–383. Springer, 2004. doi:10.1007/978-3-540-30538-5_31.

29 A Yu Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191, 1997.

30 Alexei Y. Kitaev, Alexander Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate studies in mathematics*. American Mathematical Society, 2002. URL: https://bookstore.ams.org/gsm-47/.

31 Yi-Kai Liu, Matthias Christandl, and F. Verstraete. Quantum computational complexity of the n-representability problem: Qma complete. *Phys. Rev. Lett.*, 98:110503, 2007. doi:10.1103/PhysRevLett.98.110503.

32 Alex Meiburg. Quantum constraint problems can be complete for BQP, QCMA, and more, 2021. arXiv:2101.08381.

33 Roberto Oliveira and Barbara M. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Info. Comput.*, 8(10):900–924, 2008. doi:10.26421/QIC8.10-2.

34 C.H. Papadimitriou and M. Yannakakis. The complexity of facets (and some facets of complexity). *Journal of Computer and System Sciences*, 28(2):244–259, April 1984. doi:10.1016/0022-0000(84)90068-0.

35 Francesco Petiziol, Mahdi Sameti, Stefano Carretta, Sandro Wimberger, and Florian Mintert. Quantum simulation of three-body interactions in weakly driven quantum systems. *Phys. Rev. Lett.*, 126:250504, June 2021. doi:10.1103/PhysRevLett.126.250504.

36 Asad Raza, Jens Eisert, and Alex B. Grilo. Complexity of geometrically local stoquastic hamiltonians, 2024. arXiv:2407.15499.

37 Dorian Rudolph. Towards a universal gateset for QMA_1, 2024. arXiv:2411.02681.

38 Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, STOC '78, pages 216–226, New York, NY, USA, 1978. Association for Computing Machinery. doi:10.1145/800133.804350.

39 Norbert Schuch and Frank Verstraete. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature physics*, 5(10):732–735, 2009. doi:10.1038/nphys1370.

40 Jordi Weggemans, Marten Folkertsma, and Chris Cade. Guidable Local Hamiltonian Problems with Implications to Heuristic Ansatz State Preparation and the Quantum PCP Conjecture. In *19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024)*, volume 310 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:24, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.TQC.2024.10.

41 Tzu-Chieh Wei, Michele Mosca, and Ashwin Nayak. Interacting boson problems can be qma hard. *Phys. Rev. Lett.*, 104:040501, 2010. doi:10.1103/PhysRevLett.104.040501.

42 Pawel Wocjan, Dominik Janzing, and Thomas Beth. Two qcma-complete problems. *Quantum Inf. Comput.*, 3(6):635–643, 2003. doi:10.26421/QIC3.6-7.

43 Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. doi:10.1145/3402029.

A Notation and background

A.1 Notation

For a bitstring x, let $|x|$ denote the number of bits in x.

A *promise problem* $A = (A_{yes}, A_{no})$ is a computational problem consisting of two non-intersecting sets $A_{yes}, A_{no} \subseteq \{0,1\}^*$ where given an instance $x \in \{0,1\}^*$ (promised to be in one of the two sets), one is tasked to determine if $x \in A_{yes}$ (x is a yes-instance) or $x \in A_{no}$ (x is a no-instance).[16] If $A_{yes} \cup A_{no} = \{0,1\}^*$, then A is called a *language*. For an instance x, we let $n = |x|$ denote the size of x.

For some complexity classes, we specify the gate set used. Here, we use the Clifford-cyclotomic gate sets $\mathcal{G}_m$ defined in Ref. [4]. Specifically, we only consider those that are a power of two. These are: $\mathcal{G}_2 := \{X, \text{CNOT}, \text{Toffoli}, H \otimes H\}$, $\mathcal{G}_4 := \{X, \text{CNOT}, \text{Toffoli}, \zeta_8 H\}$, and for $l \geq 3$, $\mathcal{G}_{2^l} := \{H, \text{CNOT}, T_{2^l}\}$. Here, $T_{2^l} = \text{diag}(1, \zeta_{2^l})$ where $\zeta_{2^l} = e^{2\pi i/2^l}$ is a primitive 2^l-th root of unity.

In all quantum circuits considered here, we let $U_0 = I$ be a dummy unitary used for convenience. The same is true for classical circuits Q and classical reversible circuits R. For circuits that decide computational problems, we let *ans* denote the qubit that when measured provides this decision. We accept the instance if the qubit is measured and yields outcome “1”, and reject otherwise. Usually, *ans* is the first ancilla qubit of the circuit.

For a circuit U_n that decides an instance x with $|x| = n$, we denote U_x as the circuit where the instance x is encoded into it and the inputs are only ancilla qubits in the $|0\rangle$ state.

A.2 Classes with perfect completeness

This version of the paper assumes familiarity with basic complexity classes; for a detailed introduction, we refer the reader to the full version of the paper [14]. Here, we only discuss a variation of probabilistic complexity classes with *perfect completeness*.

These are the classes where the acceptance probability of yes-instances is equal to one, and they are one of the two types of classes with *one-sided error*. Although these classes appear to be similar to their two-sided error variation, quantum complexity classes with one-sided error require a more precise treatment as they are not known to be independent

[16] The asterisk over the set is known as the *Kleene star* and is used to represent strings of any finite size.

of the gate set used. Indeed, the Solovay-Kitaev theorem [29] used to resolve this issue for classes with two-sided error only works for approximate equivalence of universal gate sets and not perfect equivalence. Thus, for classes with one-sided error (with some exceptions), one must specify the gate set used by the quantum circuits. This is not the case for classical complexity classes as it is known that every classical circuit using gate set $\mathcal{G}$ can be perfectly simulated by another circuit using a universal gate set $\mathcal{G}'$.

Given this discussion, we can then define one-sided error classes as follows:

▶ **Definition 23** (Classes with perfect completeness). *Let $\mathcal{C}$ be a complexity class with two-sided error. The variant of this class with perfect completeness is defined in a similar way to $\mathcal{C}$ except for the following differences:*

1. *For a promise problem A, the acceptance probability must be exactly 1 when $x \in A_{yes}$.*
2. *If $\mathcal{C}$ is a quantum complexity class, the gate set $\mathcal{G}$ used by the quantum circuits $\{U_n\}$ must be specified.*

The class is generally denoted as $\mathcal{C}_1$, or $\mathcal{C}_1^{\mathcal{G}}$ if it is a quantum complexity class.

This sensibility to the gate set in quantum complexity classes is the reason why, in Theorems 1 and 2, we explicitly state that LCT-QSAT and SLCT-QSAT are complete for BQP_1 with the particular choice of gate set $\mathcal{G}_8$. It also presents other complications. To see this, consider BQP. It is evident that $\mathsf{BQP}_1^{\mathcal{G}} \subseteq \mathsf{BQP}$ for any arbitrary gate set $\mathcal{G}$, and also that $\mathsf{P} \subseteq \mathsf{BQP}$. However, is it true that $\mathsf{P} \subseteq \mathsf{BQP}_1^{\mathcal{G}}$? Fortunately, one can show that for the Clifford+T gate set (i.e. $\mathcal{G}_8$) used in this paper, the class $\mathsf{BQP}_1^{\mathcal{G}_8}$ follows the intuitive containment of classes.

Interestingly, Jordan *et al.* [27] showed that if the circuits that decide a QCMA problem consist of gates with a succinct representation (e.g. $\mathcal{G}_8$), the acceptance probability of yes-instances can be amplified additively to be exactly 1. In other words, they showed that $\mathsf{QCMA} \subseteq \mathsf{QCMA}_1^{\mathcal{G}_8}$, concluding that $\mathsf{QCMA}_1^{\mathcal{G}_8} = \mathsf{QCMA}$. This explains why in Theorem 4 we state that the problem WITNESSED SLCT-QSAT is QCMA-complete. To this day, it remains an open question whether a similar scheme can also work for BQP and QMA. In the case of QMA, it seems this is not the case as one can show that there exists an oracle for which $\mathsf{QMA} \neq \mathsf{QCMA}_1$ [1]. However, a similar claim was made about QCMA and QCMA_1.

A.3 k-QSAT & the Circuit-to-Hamiltonian transformation

Here, we introduce QUANTUM k-SAT (denoted here as k-QSAT) as presented by Gosset and Nagaj in Ref. [22]. We present relevant parts of the proofs showing that k-QSAT is contained in QMA_1 for any constant k, and QMA_1-hard for $k \geq 6$. While Bravyi's [6] original work demonstrates hardness for $k \geq 4$, we choose to present this slightly weaker result for brevity, but also to introduce our clock encoding and notation useful for the rest of this paper.

As we are working to prove the inclusion and hardness of this problem for a class requiring perfect completeness, it is necessary to specify the gate set used by the quantum circuits. For reasons discussed below, we choose $\mathcal{G}_8$. In addition, we also have to be wary that all operations can be performed with perfect accuracy using gates from this set and all measurements are in the computational basis. For this purpose, Gosset and Nagaj introduce the following set of projectors.

▶ **Definition 24** (Perfectly measurable projectors). *Let $\mathcal{P}$ be the set of projectors such that every matrix element in the computational basis is of the form $\frac{1}{4}(a + ib + \sqrt{2}c + i\sqrt{2}d)$ for all $a, b, c, d \in \mathbb{Z}$.*

The (promise) problem k-QSAT can be defined as follows.

▶ **Definition 25** (k-QSAT). *Given an integer n and an instance x consisting of a collection of projectors $\{\Pi_i\} \subset \mathcal{P}$ where each Π_i acts nontrivially on at most k qubits, the problem consists on deciding whether (1) there exists an n-qubit state $|\psi_{sat}\rangle$ such that $\Pi_i|\psi_{sat}\rangle = 0$ for all i, or (2) for every n-qubit state $|\psi\rangle$, $\Sigma_i \langle\psi|\Pi_i|\psi\rangle \geq 1/poly(n)$. We are promised that these are the only two cases. We output "YES" if (1) is true, or "NO" otherwise.*

One can think of this problem as being presented with a list of constraints or *clauses* (the projectors Π_i) and tasked with distinguishing between the following cases: (1) there exists a state that satisfies all constraints (a *satisfying state*), or (2) any possible state induces a violation of the constraints greater than $1/poly(n)$. The promise sets the conditions for classifying instances as either $x \in A_{yes}$ or $x \in A_{no}$.[17]

A.3.1 In QMA_1

Suppose we are presented with a witness state $|\psi_{\text{wit}}\rangle$ and a k-QSAT instance composed of projectors $\{\Pi_i\}$. The quantum algorithm that decides whether this state satisfies all projectors Π_i consists of simply measuring the eigenvalues of all projectors on this state. Then, if all measured eigenvalues are 0, we conclude that all projectors are satisfied by the state and output "YES". Otherwise, we reject.

Specifically, we measure the eigenvalue of a projector Π_i by applying the unitary

$$V(\Pi_i) = \Pi_i \otimes X + (I - \Pi_i) \otimes I, \tag{2}$$

to the witness and an additional ancilla qubit in the state $|0\rangle$, followed by a measurement of the ancilla in the computational basis. Here, X denotes the Pauli-X gate. The probability that $|\psi_{\text{wit}}\rangle$ does not satisfy projector Π_i (obtain outcome "1") is given by

$$p_i = \langle\psi_{\text{wit}}|\, \Pi_i \,|\psi_{\text{wit}}\rangle\,. \tag{3}$$

Defining the acceptance probability as the probability that all measurements produce outcome "0", and assuming $V(\Pi_i)$ can be implemented perfectly with gate set $\mathcal{G}$, one can show that this algorithm meets the completeness and soundness conditions of QMA_1, concluding that k-QSAT is contained in this class.

As mentioned, to support this claim, it is necessary to demonstrate that $V(\Pi_i)$ can be implemented perfectly using gate set $\mathcal{G}_8$. This follows from the fact that the projectors Π_i are from the set $\mathcal{P}$ together with a theorem by Giles and Selinger [20].

A.3.2 QMA_1-hard

Now, we discuss elements of the proof demonstrating that k-QSAT is QMA_1-hard when $k \geq 6$ and for any gate set $\mathcal{G}$ that is universal for quantum computation.

The idea is to demonstrate that any instance x of an arbitrary promise problem in QMA_1 can be transformed or *reduced* in polynomial time into an instance x' of k-QSAT, where the answer to both problems is the same for all instances. Furthermore, we also need to show that all projectors of the resulting k-QSAT instance act on at most 6 qubits.

[17] Without the promise, the problem seems to become harder, as it requires distinguishing between the case where the projectors are satisfiable, and the case where they are not but the violation induced by some states could be exponentially close to zero. Without the promise, the problem is most likely not contained in QMA_1.

Let $U_x = U_L \dots U_1$ with $U_i \in \mathcal{G}$ and $L = poly(n)$ be the QMA_1 verification circuit where given an instance x of a problem $A = (A_{yes}, A_{no})$, U_x decides whether $x \in A_{yes}$ or $x \in A_{no}$. The input to the circuit consists of the p-qubit witness state $|\psi_{\text{wit}}\rangle$, and a q-qubit ancilla register D (referred to as the *data* register) initialized to the state $|0\rangle^{\otimes q}$, where p and q are two polynomials in $n = |x|$. Additionally, let the answer be obtained by measuring the ancilla qubit *ans* in the computational basis. The goal of the reduction is to engineer a set of 6-local projectors such that they are uniquely satisfied by the state encoding the evaluation of the circuit U on $|\phi_0\rangle := |0\rangle^{\otimes q} \otimes |\psi_{\text{wit}}\rangle$ at all steps of the computation. This state is appropriately known as the (computational) *history state* and is given by

$$|\psi_{hist}\rangle := \frac{1}{\sqrt{L+1}} \sum_{t=0}^{L} U_t \dots U_0 \, |\phi_0\rangle_D \otimes |C_t\rangle_C \, . \tag{4}$$

Here, we have introduced a *clock* register C acting on a new (not yet specified) Hilbert space used to keep track of the current step in the computation. This history state can be defined in many ways depending on the implementation of the states $|C_t\rangle$. In this paper, we choose a clock encoding acting on $\mathcal{H}_{clock} = (\mathbb{C}^3)^{\otimes L+1}$, consisting of the *ready* state $|r\rangle$, the *active* state $|a\rangle$, and the *dead* state $|d\rangle$. The clock progresses as

$$\begin{aligned} |C_0\rangle &= |a_0 r_1 r_2 \dots r_L\rangle \, , \\ |C_1\rangle &= |d_0 a_1 r_2 \dots r_L\rangle \, , \\ &\vdots \\ |C_L\rangle &= |d_0 d_1 d_2 \dots a_L\rangle \, . \end{aligned} \tag{5}$$

The projectors that allow us to build the required 6-QSAT instance act on both of these Hilbert spaces and are given by

$$\begin{aligned} P_{init}^{(i)} &:= |1\rangle\langle 1|_i \otimes |a\rangle\langle a|_0 \, , \\ P_{out}^{(i)} &:= |0\rangle\langle 0|_i \otimes |a\rangle\langle a|_L \, , \\ P_{prop,U}^{(i)} &:= \frac{1}{2}\left[I^{\otimes 2} \otimes |ar\rangle\langle ar| + I^{\otimes 2} \otimes |da\rangle\langle da| - U \otimes |da\rangle\langle ar| - U^\dagger \otimes |ar\rangle\langle da| \right], \end{aligned} \tag{6}$$

which receive an index to specify its action on a given particle. Moreover, $P_{prop,U}$ acts on clock qudits $i-1$ and i. Observe that P_{init} and P_{out} act on a single data and clock particle, while $P_{prop,U}$ acts on two data qubits and two clock particles. As each clock particle can be represented by two qubits, albeit a bit wastefully, it is evident that these projectors are at most 6-local (on qubits). Other clock encodings may lead to different locality.[18]

Each projector in Equation (6) penalizes states that do not meet certain requirements. (Initialization) P_{init} requires that when clock particle 0 is in the state $|a\rangle$, data qubit i is initialized to $|0\rangle$. (Computational propagation) $P_{prop,U}$ requires that as clock particles i and $i+1$ transition from $|ar\rangle$ to $|da\rangle$, U is applied to two qubits of the data register. (Readout)

[18] In Ref. [6], Bravyi employs a four-state clock encoding, $2L+1$ clock basis states, and an additional propagation projector. This allows interactions between either two clock particles at a time or one clock particle and two data qubits, resulting in 4-local projectors. However, this comes at a cost of increased clock particle dimensionality.

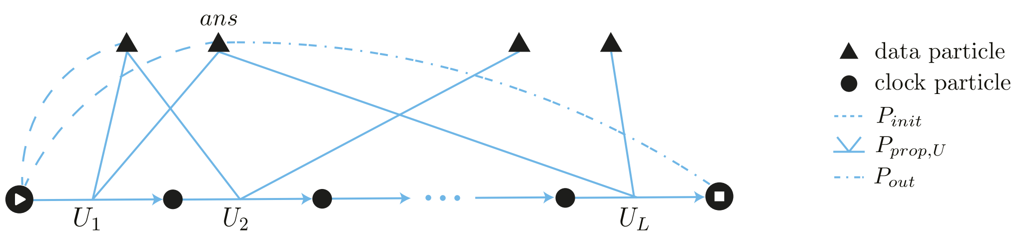

Figure 4 Representation of a 6-QSAT instance which encodes a QMA_1 verification circuit $U = U_L \dots U_1$. For simplicity, we let U act on four data qubits: two ancilla qubits (those present in P_{init} clauses), and two for the witness state (uninitialized ones). The ancilla measured at the end of the computation is labeled *ans*. The leftmost and rightmost clock particles are marked with "start" and "stop" icons, indicating the action of P_{start} and P_{stop} clauses, respectively. The P_{clock} clauses are shown as arrows on top of $P_{prop,U}$ lines, representing the clock progression.

Finally, P_{out} requires that when clock qudit L is in the state $|a\rangle$, data qubit i is in the state $|1\rangle$.[19] Aside from these projectors, one also has to define

$$\begin{aligned} P_{start} &:= |r\rangle\langle r|_0 \,, \\ P_{stop} &:= |d\rangle\langle d|_L \,, \\ P^{(i)}_{clock} &:= |r\rangle\langle r|_i \otimes (I - |r\rangle\langle r|)_{i+1} + |a\rangle\langle a|_i \otimes (I - |r\rangle\langle r|)_{i+1} + |d\rangle\langle d|_i \otimes |r\rangle\langle r|_{i+1} \,, \end{aligned} \tag{7}$$

which are at most 4-local projectors requiring that the clock states have the form described in Equation (5). Furthermore, the six types of projectors of Equations (6) and (7) are of the form given in Definition 24 and are hence projectors from $\mathcal{P}$, as required. Finally, using these projectors, the instance that encodes the verifier circuit $U = U_L \dots U_1$ is given by

$$\begin{aligned} H_{init} &:= \sum_{b \in ancilla} P^{(b)}_{init}, \\ H_{prop} &:= \sum_{t=1}^{L} P^{(t)}_{prop,U_t} \\ H_{out} &:= P^{(ans)}_{out}, \\ H_{clock} &:= P_{start} + P_{stop} + \sum_{c \in C} P^{(c)}_{clock}. \end{aligned}$$

We illustrate this instance in Figure 4. The set of projectors that define this k-QSAT instance are the individual terms of the sum. They are often grouped into positive semi-definite terms as above for historical reasons. Briefly, the H_{init} term requires that all ancilla qubits from register D are initialized to $|0\rangle$, leaving the data qubits for the witness state "free" or uninitialized. H_{prop} defines a clock register of $L+1$ particles and requires that as time progresses from $t-1$ to t, U_t is applied to the data qubits. H_{out} requires that at the end of the computation *ans* is measured to be "1". Finally, H_{clock} requires that we obtain a running clock register progressing as shown in Equation (5). Together, H_{init}, H_{prop}, and H_{clock} require that if there exists a state satisfying all of their projectors, the state must mimic the evaluation of the quantum circuit $U = U_L \dots U_1$ on the state $|\phi_0\rangle$. This is the history state

[19] Unlike Bravyi [6] and Meiburg [32], we define P_{out} so it is satisfied when the logical qubit is in the state $|1\rangle$, and not $|0\rangle$.

of Equation (4) with the clock encoding of Equation (5). Moreover, if the verification circuit U accepts yes-instances with certainty, the history state also satisfies H_{out} and is thus the unique ground state of the 6-local Hamiltonian $H = H_{init} + H_{prop} + H_{out} + H_{clock}$.

This concludes the transformation of the circuit into local Hamiltonians. Completing the proof that 6-QSAT is QMA_1-hard requires showing that, if $x \in A_{yes}$, then x' has a frustration-free ground state, and if $x \in A_{no}$, then the ground state energy of H is not too low. Proving these is beyond the scope of this section.