

On the complexity of probabilistic trials for hidden satisfiability problems

Itai Arad^{*1}, Adam Bouland^{†2}, Daniel Grier^{‡2}, Miklos Santha^{§1,3}, Aarthi Sundaram^{¶1}, and Shengyu Zhang^{||4}

¹Center for Quantum Technologies, National University of Singapore, Singapore

²Massachusetts Institute of Technology, Cambridge, MA USA

³CNRS, IRIF, Université Paris Diderot 75205 Paris, France

⁴Department of Computer Science and Engineering, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong

June 11, 2016

Abstract

What is the minimum amount of information and time needed to solve 2SAT? When the instance is known, it can be solved in polynomial time, but is this also possible without knowing the instance? Bei, Chen and Zhang (STOC '13) considered a model where the input is accessed by proposing possible assignments to a special oracle. This oracle, on encountering some constraint unsatisfied by the proposal, returns only the constraint index. It turns out that, in this model, even 1SAT cannot be solved in polynomial time unless $P = NP$. Hence, we consider a model in which the input is accessed by proposing probability distributions over assignments to the variables. The oracle then returns the index of the constraint that is most likely to be violated by this distribution. We show that the information obtained this way is sufficient to solve 1SAT in polynomial time, even when the clauses can be repeated. For 2SAT, as long as there are no repeated clauses, in polynomial time we can even learn an equivalent formula for the hidden instance and hence also solve it. Furthermore, we extend these results to the quantum regime. We show that in this setting 1QSAT can be solved in polynomial time up to constant precision, and 2QSAT can be learnt in polynomial time up to inverse polynomial precision.

1 Introduction

SAT has been a pivotal problem in theoretical computer science ever since the advent of the Cook-Levin Theorem [7, 13] proving its NP-completeness. It has a wide array of applications in

*email: arad.itai@fastmail.com

†email: adam@csail.mit.edu

‡email: grierd@mit.edu

§email: miklos.santha@gmail.com

¶email: aarthims@gmail.com

||email: syzhang@cse.cuhk.edu.hk

operations research, artificial intelligence and bioinformatics. Moreover, it continues to be studied under various specialized models such as as random-SAT and building efficient SAT solvers for real-life scenarios. In the complexity theoretic setting, we know that while 3SAT is NP-complete [7, 13], 2SAT can be solved in linear time [12, 9, 3]. Given the fundamental nature of 2SAT, in this paper, we consider the following question:

What is the minimum amount of information needed to solve 2SAT in polynomial time?

More precisely, what happens if there is no direct access to the problem instance? Are there settings where one can *solve* 2SAT without ever *learning* the instance under consideration? We can also pose the same question for the quantum setting where the quantum analogue of SAT (QSAT) can be seen as a central problem in condensed matter physics. Complexity theoretically, we know that 2QSAT can be solved in linear time [2, 8] while 3QSAT is hard for QMA_1 [10], where QMA_1 is a quantum complexity class analogous to NP. We approach these questions through the “trial and error” model. In this model, one guesses a solution to an unknown constraint satisfaction problem and tests if it is valid. If so, then the problem is solved. Otherwise, the trial fails, and some information about what was wrong with the trial is revealed. This type of problem arises in a number of natural scenarios, in which one has incomplete or limited information about the problem they are trying to solve [4]. For example, the CSP may be instantiated by a complex biological or physical process to which one does not have access.

This approach to problem solving was first formalized by Bei, Chen and Zhang [4]. They considered several types of CSPs and analyzed the computational complexity in the “unknown input” setting. Specifically, they consider an oracle model where one can propose solutions to the CSP, and if the solution is not satisfying, then the oracle reveals the identity of a constraint which was violated. For example, if the CSP is an instance of Boolean satisfiability (SAT), then after an unsuccessful trial, one may learn that “clause 7 was violated”, but not anything further. In particular, literals present in clause 7 will not be revealed - only the label of violated clause is known. Furthermore, if there are several violated constraints, then the oracle reveals only one of them, in a possibly adversarial manner. In this paper, we will refer to this as the “arbitrary violated constraint” oracle.

This model gives extremely limited access to the instance. In fact, Ivanyos et al. [11] showed that even if the underlying CSP is a 1SAT instance, accessing it with the BCZ oracle, one cannot determine if it is satisfiable in polynomial time unless $P = NP$. This drastically increases the difficulty of deciding a trivial problem like 1SAT (assuming $P \neq NP$). Interestingly, if there is access to a SAT solver, then 1SAT (and even generic SAT) in this setting can be solved with polynomially many trials [4]. So in some sense, their model reveals a sufficient amount of *information* to solve the 1SAT instance. However, *decoding* this information requires superpolynomial time (assuming $P \neq NP$). In short the information needed to solve the problem is present, but it is not accessible to poly-time algorithms.

In this paper, we ask if there are any meaningful modifications of their model which allow us to solve simple CSPs like 1SAT and 2SAT in polynomial time. A natural starting point is to randomize the “arbitrary violated constraints” model. One obvious way to do that is to consider allowing randomized queries to the oracle. This however does not significantly decrease the complexity of the problems. A second approach to randomize is to let the oracle return a violated clause at random. Contrary to the previous approach, this model trivializes the problem, since by repeating the same trial many times the oracle will reveal all violated clause indices with high probability.

This in turn allows one to learn the entire instance, and therefore trivially, to solve 1SAT and 2SAT.

Motivated by these unfruitful approaches we consider a model which does not allow one to completely learn the underlying instance, but it still yields polynomial time algorithms for 1SAT and 2SAT. Specifically, in this model one can propose a probability distribution D over assignments, and the oracle reveals the index of the clause which is most likely to be violated by this trial. If there are multiple clauses with the same probability of violation under D , then the oracle can break ties arbitrarily. In particular, product distributions over the variables suffice for our application, so one merely specifies the probability p_i that each variable x_i is set to 1 in the assignment, to $1/\text{poly}$ precision. We show that in this model, there exist cases where one cannot learn the underlying 1SAT or 2SAT instance. However, despite this limitation, one can still solve in polynomial time 1SAT and a restricted version of 2SAT where clauses are not repeated. In the course of the algorithm for the restricted version of 2SAT, we actually learn an equivalent formula with the same set of satisfying assignments. Furthermore, we are able to extend this model to the quantum setting, and show that one can solve, in polynomial time, Quantum 1SAT (1QSAT) up to constant precision. We also show that in polynomial time we can learn Quantum 2SAT (2QSAT) up to inverse polynomial precision. This, however, seems insufficient to solve the hidden instance in polynomial time due to some subtle precision issues, which we discuss in Section 9.

Relation to prior work. As previously mentioned, Bei Chen and Zhang [4] introduced the trial and error model. They considered several examples of CSPs and analyzed their complexity under the unknown input model with the “arbitrary violated constraint” oracle. With regards to SAT, they showed an algorithm to solve hidden-SAT using polynomially many queries to the oracle (given access to a SAT oracle). Furthermore, they showed that one cannot efficiently learn generic SAT instances in this model, because it takes $\Omega(2^n)$ queries to the oracle to learn a clause involving all n variables of the instance.

Subsequently, Ivanyos *et al.* [11] characterized the complexity of classical CSPs in several hidden input models. In particular, they consider the “arbitrary violated constraint” model described above, as well as models which reveal more information such as the variables involved in the violated clause or the relation of the violated clause. They show a generic “transfer theorem” which classifies the complexity of hidden versions of CSPs given properties of the base CSP. In particular, their transfer theorem implies that the hidden version of 1SAT with arbitrary violated constraints cannot be solved in polynomial time unless $P = NP$. This indicates that the “arbitrary violated constraint” model is fairly restrictive.

In parallel, Bei, Chen and Zhang [5] considered a version of the trial and error model for linear programming. Suppose you have a linear program, and you are trying to determine whether or not it is feasible (By standard reductions this is as difficult as solving a generic LP). They consider a model in which one can propose a point, and the oracle will return the index of an arbitrary violated constraint (half-plane) in the linear program. They show that in this model, one requires exponentially many queries to the oracle to determine if an LP is feasible. However, they then consider a relaxation of this model, in which the oracle returns the index of the *worst-violated* constraint, i.e. the half-plane which is furthest (in Euclidean distance) from the proposed point. Surprisingly, they show (using a variant of the ellipsoid algorithm) that one can still solve linear programs in this model in polynomial time. Our model can be seen as an analogue of the “worst violated constraint” model of Bei, Chen and Zhang [5] for the case of hidden SAT (H-SAT).

Our Results. Our results can be broken into several sections. First, we consider a relaxation of the “arbitrary violated constraint” model of Bei, Chen and Zhang [4], in which the oracle reveals which subset of clauses are violated by each assignment¹. We show that in this model, there exist simple algorithms to learn H-1SAT or H-2SAT instances, and hence solve them in polynomial time. In some sense these models are almost “too easy” as they allow you to easily learn the instance (See Section 3).

We then explore the “worst violated constraint” model for the rest of the paper. We provide a toy example as to why this model is more powerful than the “arbitrary violated constraint” model of Bei, Chen and Zhang [4]. They showed that it requires $\Omega(2^n)$ time to learn a SAT clause involving all n variables. In Section 4 we show that it is possible to learn such a SAT clause in polynomial time in the “worst violated constraint” model. Among our main results is the analysis of the computational complexity of H-1SAT and that of H-2SAT.

Theorem (Informal statement). *Given a hidden SAT formula Φ on n variables and n clauses, it is possible to find a satisfying assignment for Φ in polynomial time if Φ is a (a) 1SAT formula or (b) 2SAT formula with no repeated clauses.*

Our algorithm for H-1SAT, in Section 5, works even when clauses are repeated multiple times in the instance, despite the fact that it’s not possible to *learn* the instance in this setting. This is in sharp contrast to the “arbitrary violated constraint” model, where even H-1SAT cannot be solved in polynomial time unless $P = NP$ [11]. The main difficulty in deriving our algorithm for H-1SAT comes from dealing with repeated clauses, which allow the oracle to obscure information about the instance. Unlike the H-1SAT case, the algorithm for H-2SAT discussed in Section 6, works by attempting to learn the instance; it either succeeds in learning an equivalent instance (in which case one can solve the problem using any 2SAT algorithm), or it accidentally stumbles upon a satisfying assignment in the meantime and aborts. The problem of solving H-2SAT with repeated clauses similar to H-1SAT is left for future work.

Following this we generalize these results to the quantum case. In this case the goal is to determine if a set of 1-qubit or a set of 2-qubit projectors is mutually satisfiable or not. We consider an analogue of this model in which one can propose a probability distribution over quantum states (i.e. a density matrix), and the oracle returns the index of the clause which is most likely to be violated. Our results for hidden QSAT (H-QSAT) show that

Theorem (Informal statement). *Given a H-QSAT instance H defined on n qubits with m projectors and $\epsilon > 0$, it is possible to*

- (a) *solve H to a precision ϵ in time $O(n^{\log(1/\epsilon)})$ if H is a 1QSAT instance and*
- (b) *learn each projector of H up to precision ϵ in time $O(n^4 + n^2 \log(1/\epsilon))$, if H is a 2QSAT instance as long as the interaction graph of H is not star-like.*

By star-like, we mean the interaction graph contains an edge that is incident to all other edges in the graph. At this point it is worth comparing the notions of *learning* and *solving* hidden instances both in the classical and quantum settings. The classical case is more straightforward where learning an instance means learning all the literals present in each clause, whereas solving means finding a satisfying assignment. For example, our algorithm for H-2SAT without repetitions

¹This is equivalent to a model in which the oracle reveals a random violated clause - by repeating each query many times one can learn the set of violated clauses with high probability.

learns the instance, while our algorithm for H-1SAT solves the instance without learning it. For hidden versions of 1SAT and 2SAT, learning the instance in polynomial time automatically triggers solving it in polynomial time as well.

However, in the quantum setting this simple relation between learning and solving breaks down. The continuous nature of QSAT means we can only learn a projector or find a satisfying assignment up to a specified precision ϵ . The latter is accomplished with our H-1QSAT algorithm in Section 8. However in the case of hidden 2QSAT learning the instance up to precision ϵ does not imply that one can solve the instance up to precision $\text{poly}(n, \epsilon)$ in polynomial time. This can be attributed to current algorithms for 2QSAT being very sensitive to precision errors. This issue of divergence between the notions of learning and solving H-2QSAT instances is further discussed in Section 9.

2 Notations and Preliminaries

Boolean Satisfiability. The Boolean satisfiability problem, generally referred to as SAT, is a constraint satisfaction problem defined on n variables $\mathbf{x} = \{x_1, \dots, x_n\}$ where we are given a formula represented as a conjunction of m clauses and each clause is a disjunction of *literals* (variables, x_j , or negated variables, \bar{x}_j). The problem is solved if we can find an assignment to the variables (i.e. $\forall i, x_i \in \{0, 1\}$) that sets the value of every clause to 1. In particular, if each clause involves at most k literals, then this problem is classified as kSAT. It is well known that while 2SAT can be solved in linear time [12, 9, 3], kSAT for $k \geq 3$ is NP-complete [7, 13]. A useful notion is that of *clause types* which is defined as the unordered set of literals present in the clause. Specifically, the clause type for $C_j = (x_a \vee \bar{x}_b \vee x_c)$ is denoted by $T(C_j) = \{x_a, \bar{x}_b, x_c\}$. So, all possible clause types for 2SAT would be $\{\{x_a, x_b\}, \{x_a, \bar{x}_b\}, \{\bar{x}_a, x_b\}, \{\bar{x}_a, \bar{x}_b\} \mid a, b \in [n] \text{ and } a \neq b\}$, where $[n]$ denotes the set $\{1, \dots, n\}$. From this definition, it is clear that 2SAT has $O(n^2)$ clause types and similarly, kSAT would have $\binom{2n}{k} = O(n^k)$ clause types. . Given a SAT formula ϕ , we say that the SAT formula ϕ' is *equivalent* to ϕ if for all assignments $\mathbf{x} \in \{0, 1\}^n$, \mathbf{x} satisfies ϕ if and only if it satisfies ϕ' . For any formula ϕ , $\text{SAT}(\phi) := \{\mathbf{x} \in \{0, 1\}^n \mid \phi(\mathbf{x}) = 1\}$.

Hidden SAT. While considering the unknown input version of SAT (resp. kSAT), the boolean formula is considered as hidden and accessible only via an oracle that accepts an assignment and reveals some form of violation information. In our case, this is the “worst violated oracle” which accepts a *probabilistic* assignment and reveals a clause that has the *highest probability* of being violated with ties being broken arbitrarily. A probabilistic assignment for a set of n variables is a function $\mathbf{a} : [n] \rightarrow [0, 1]$ such that $\text{Pr}[x_i = 1] = \mathbf{a}(i)$ and $\text{Pr}[x_i = 0] = \text{Pr}[\bar{x}_i = 1] = 1 - \mathbf{a}(i)$. For the sake of concise notation, these are usually written as $x_i = \mathbf{a}(i)$ and $\bar{x}_i = 1 - \mathbf{a}(i)$. This naturally translates to the notion of the probability of a clause C_j being violated which is defined as $\text{Pr}[C_j = 0] := \prod_{\ell \in T(C_j)} \text{Pr}[\ell = 0] = \prod_{\ell \in T(C_j)} (1 - \ell)$ which allows the oracle to calculate the probability for each clause being violated. Here we are using ℓ to refer both to the identity of a literal as well as to the probability that literal ℓ is set to true. Now, the problem H-SAT (resp. H-kSAT) consists of finding a satisfying assignment for a hidden SAT (resp. kSAT) formula by proposing probabilistic assignments to the “worst violated oracle”. One way we do this is also by *learning* an *equivalent formula* to the hidden instance and solve it to find a satisfying assignment. By learning we mean the process of using the information from a series of violations to determine what a clause in the hidden instance could be.

Note that it’s possible for an instance to contain clauses which will never be returned by the oracle. For instance, given clauses C_i and C_j , if $T(C_i) \subset T(C_j)$, then clause C_i will always be at

least as violated as C_j . Hence the oracle might never return clause C_j . For this reason we will say that C_i *obscures* C_j if $T(C_i) \subset T(C_j)$. An obscured clause might never be returned by the oracle.

The complexity of the algorithms in the following sections is in terms of the total running time where one query to the oracle takes unit time.

3 An easier model: an oracle which reveals all violated constraints

We begin by considering a simple generalization of the oracle model of Bei Chen and Zhang [4]. We will call this the “all violated constraints” model. In particular, suppose that you have a SAT instance which is hidden from you. Instead, you have access to an oracle which, given an assignment to the variables, identifies which clauses are violated by that assignment. In the model of Bei Chen and Zhang [4], the oracle only returns one violated clause, and the oracle’s response may be adversarial. Here, in contrast, the oracle reveals the identities of *all* violated clauses. This is a more natural model than the “arbitrary violated constraint” model, because in many real-life applications, you would expect to see a random violated constraint², or all constraints, not an adversarially chosen one. Therefore in many cases this is a more natural model than the one considered by Bei Chen and Zhang.

We begin by showing that this model is much more powerful than the “arbitrary violated constraint” model of BCZ. In particular, this model allows one to learn a hidden kSAT instance in $O(n^k)$ time. Therefore, by simply learning the underlying instance, one can solve H-1SAT or H-2SAT in polynomial time.

Theorem 1. *In the “all violated constraints” model, there is an algorithm which either learns an arbitrary kSAT instance on n variables and m clauses, or else finds a satisfying assignment to the instance, in time $O(mn^k)$, where the big- O notation hides a constant which depends on k .*

Proof Let x_1, x_2, \dots, x_n be the variables of your instance, and $C_1 \dots C_m$ be the clauses. In the following algorithm we will assume that all the assignments tried, fail to satisfy the instance; if they happen to satisfy it, then we have found a satisfying assignment and the algorithm aborts.

The type of a clause is the subset of literals contained in the clause. For instance, a clause could be of type $(x_1 \vee x_2 \vee \bar{x}_3)$ or of type $(\bar{x}_1 \vee x_3)$. In an instance of kSAT, there could be multiple clauses of the same type. We will show that for each clause type T involving k literals, we can learn which clauses of the instance are of type T in time $O(m)$. This implies the claim.

We will now show how to learn which clauses are of type $(x_1 \vee x_2 \vee \dots \vee x_k)$ in our instance. An analogous proof holds for other clause types. First propose the assignment $x_1 = x_2 = \dots = x_k = 0$ and $x_{k+1} = x_{k+2} = \dots = x_n = 0$. This returns some subset of violated clauses S . Next propose the assignment $x_1 = x_2 = \dots = x_k = 0$ and $x_{k+1} = x_{k+2} = \dots = x_n = 1$. This returns a subset S' of violated clauses. Now take the intersection of $S \cap S'$ (which can be done in $O(m)$ time).

We now claim that $S \cap S'$ contains all clauses C_j for which all the literals in C_j are the set $\{x_1, x_2, \dots, x_k\}$. To see this, first note that any clauses that are in $S \cap S'$ are clearly violated by both of the proposed assignments. Now consider any clause C_j which involves variables outside of $x_1 \dots x_k$. For example say clause C_j contains the literal x_{k+1} . Then C_j will be satisfied by one of the two proposed assignments; hence it will not be in $S \cap S'$.

²Note that a model which returns a random violated constraint is just as powerful to this one, since by repeating the experiment one could quickly learn the set of all violated constraints with high probability.

We therefore have that $S \cap S'$ contains all clauses on the literals $\{x_1 \dots x_k\}$. Some of these clauses are of type $(x_1 \vee x_2 \vee \dots \vee x_k)$, but others may involve subsets of these literals. Now to learn which of these are of type $(x_1 \vee x_2 \vee \dots \vee x_k)$, for each subset L of the literals $\{x_1 \dots x_k\}$ which is of size $k - 1$, perform the same experiment to learn which clauses S_L involve literals in L only. Then we have that $\cup_L S_L$ are all clauses on the literals $\{x_1 \dots x_k\}$ which involve $k - 1$ of those literals or fewer. (There are k such sets L , each of which takes $O(m)$ time). Hence $(S \cap S') \setminus (\cup_L S_L)$ is the set of all clauses of type $(x_1 \vee x_2 \vee \dots \vee x_k)$ as desired.

Therefore for each clause type, it takes $O(m)$ time to learn which clauses are of that type (where we have suppressed a constant depending on k). As there are $O(n^k)$ possible clause types in a kSAT instance, this implies one can learn kSAT for any fixed k in time $O(mn^k)$ as desired. \square

Corollary 1. *In the “all violated constraints” model, 2SAT can be solved in polynomial time.*

In some sense this model is too easy. We therefore turn our attention to a more restrictive model: the “worst violated constraint” model. Here we allow one to query probability distributions over assignments, and the oracle will return a clause which is the most likely to be violated by that distribution. The oracle may break ties arbitrarily. This can be seen as a CSP-version of the linear programming model of Bei Chen and Zhang [5].

4 Comparison to the Bei, Chen and Zhang Model: WIDESAT

Recall that in [4], a solution to a hidden formula can be found in $O(nm)$ time where n is the number of variables and m is the number of clauses. However, this does not mean that an equivalent instance of the hidden formula can be found in polynomial time. In fact, they show that any randomized algorithm necessarily requires exponentially many queries in order to generate a formula equivalent to the hidden one. Let a WIDESAT *clause* be any clause that contains all n distinct variables. It turns out that WIDESAT clauses are exactly the types of clauses that are difficult to learn in the model of Bei, Chen, and Zhang.³ In contrast, we show that similar WIDESAT instances can be learned in polynomial time in our model.

Proposition 1. *Given a hidden WIDESAT instance on n variables and m distinct clauses where $m \leq n$, we can learn an equivalent instance in $O(\binom{n}{m-1}2^m + n)$ time.*

Before we prove the proposition in full, let us gain some intuition as to why this task should be easier in our model. Suppose first that there was exactly 1 WIDESAT clause containing variables x_1, x_2, \dots, x_n . Propose the solution $(1, .5, .5, \dots, .5)$. If the oracle returns that the formula is satisfied, then x_1 must be the literal present in the clause, otherwise if it’s not satisfied, then \bar{x}_1 must be the literal present in the clause. Continuing for each separate variable completes the argument.

We will need the following lemma for the full generalization.

Lemma 1. *Suppose we have a set of m distinct binary strings of length n . Then there exists a subset of $m - 1$ indices such that each string restricted to those indices is unique.⁴*

³In fact, this is not hard to see. Suppose the hidden formula has exactly one WIDESAT clause. If the oracle always returns ‘YES’ to a proposed assignment, then each query only eliminates 1 out of exponentially many WIDESAT clause types.

⁴To see that $m - 1$ indices are necessary, consider any subset of size m of $\{e_i \mid 1 \leq i \leq n\}$ where e_i denotes the binary string with a 1 in position i and 0 everywhere else. It is clear that any $m - 2$ indices are insufficient to distinguish all strings.

Proof We proceed by induction. Consider any two strings in the set. Since all strings in the set are distinct, they must differ in at least one position (without loss of generality, in the first bit). Divide the set into two groups such that all strings in the first group start with a 0 and all strings in the second start with a 1. Notice that any two strings in differing groups are distinguished by the first index, but any two strings in the same group must still all be distinct when restricted to the last $n - 1$ bits. If the first group is of size k and second is of size $m - k$, then by induction we can distinguish the strings within the group with at most $(k - 1) + (m - k - 1) = m - 2$ indices. Adding the index we used for the first comparison, we arrive at the conclusion. \square

Proof of Proposition 1 Suppose that we have $m > 1$ distinct clauses. The strategy is as follows: for each subset of $m - 1$ variables, query the oracle with all possible 0-1 assignments to those variables, setting each variable not in the set to $1/2$. Using Lemma 1, we see that we will eventually query the subset of the variables which causes the oracle to return all m clauses. Therefore, we learn the presence of $m - 1$ variables exactly as they appear in the hidden formula. Once we have this information, it is a simple task to recreate each clause C_k . Set the $m - 1$ variables such that none satisfy clause C_k . Notice that all other clauses will be satisfied, so we have reduced this to our problem of WIDESAT with one clause, which we know how to learn in $O(n)$ time. \square

5 Hidden 1SAT

In this section, we will consider the problem of a hidden 1SAT instance Φ , possibly with repetitions. Our goal will be to determine whether or not Φ is satisfiable. A natural approach one might take to solve this problem would be to learn the identity of each clause in the instance Φ . Unfortunately, in the case that the 1SAT instance has repetitions, this is not possible.

Proposition 2. *There is no algorithm which, given an instance Φ which is unsatisfiable, learns all the literals present in Φ (even granted arbitrary numbers of queries to the oracle).*

Proof Consider the following two H-1SAT instances:

$$\Phi_1 : C_1 = x_1, C_2 = \bar{x}_1, C_3 = x_1, C_4 = \bar{x}_1$$

$$\Phi_2 : C_1 = x_1, C_2 = \bar{x}_1, C_3 = x_2, C_4 = x_2$$

Both of these instances are unsatisfiable. However, note that for any oracle query, it is possible for the oracle to give the same answer (i.e. clause index) for each query. To see this, if x_1 is more violated than \bar{x}_1 or x_2 , return clause C_1 . If \bar{x}_1 is more violated than x_1 or x_2 , then return clause C_2 . If x_2 is more violated than x_1 or \bar{x}_1 , then return clause C_3 if x_1 is more violated than \bar{x}_1 , otherwise return clause C_4 . One can easily check these oracle answers are consistent with either instance. Hence these instances are indistinguishable to adversarial oracle answers, so no algorithm can distinguish Φ_1 and Φ_2 . \square

Here the difficulty in learning an unsatisfiable instance does not lie in the repetition of clauses, but rather in determining for which i do both x_i and \bar{x}_i appear in Φ . This shows that no algorithm can learn the hidden 1SAT instance⁵. Hence if there is an algorithm to solve 1SAT in this hidden setting, then it must solve the instance despite the fact that it cannot deduce the underlying instance. Surprisingly, this turns out to be possible.

⁵Note, however, it is still possible that there exists an algorithm to learn the 1SAT instance when the instance is promised to be satisfiable.

Theorem 2. *Given a hidden 1SAT instance Φ on n variables and m clauses, it is possible to determine if Φ is satisfiable in time $O(mn^2)$.*

Proof Consider an ordering of the variables $x_1 \dots x_n$. The algorithm will work by inductively constructing a series of lists L_1, L_2, \dots, L_n . Each list L_i will contain a list of partial assignments to the variables $x_1 \dots x_i$. Each list will be of size at most m , with the exception of L_n which will be of size at most $2m$. Let us call a partial assignment p to $x_1 \dots x_i$ *good* if there exists an assignment p' to the variables $x_{i+1} \dots x_n$ such that the assignment $p \cup p'$ satisfies Φ . Correspondingly, call p *bad* if it cannot be extended to a satisfying assignment of Φ . (Note in the case of 1SAT, every partial assignment is either good or bad.) Our algorithm will guarantee that, if Φ is satisfiable, then at least one assignment in each list is “good”. Therefore, by constructing the list L_n , then trying all assignments in L_n , we will be guaranteed to find a satisfying assignment if one exists.

We now describe how to construct the lists $\{L_i\}_{i \in [n-1]}$ by induction. The base case of L_1 is trivial - just add both $x_1 = 0$ and $x_1 = 1$ to the list. We now show how to construct L_{i+1} given L_i . First, let \tilde{L}_{i+1} be all possible extensions of the assignments in L_i to the variable x_{i+1} . Clearly if one of the assignments in L_i was good, then one of the assignments in \tilde{L}_{i+1} is good. However, when $i + 1 < n$, the size of \tilde{L}_{i+1} could become too large - it is of size $2|L_i|$ which could at some point become larger than m . So we need to reduce the size of \tilde{L}_{i+1} so that it contains at most m partial assignments. To decide which partial assignments to keep, we will perform the following oracle queries: for each partial assignment $y \in \tilde{L}_{i+1}$, propose the following query q_y to the oracle: set $x_1 \dots x_{i+1}$ to 0 or 1 according to y , and set all other variables to value $1/2$. The oracle will return the identity of a clause C_j which is worst violated by this fractional assignment. Now partition the elements of \tilde{L}_{i+1} according to which clause C_j was returned by the query. This divides the elements of \tilde{L}_{i+1} into at most m equivalence classes. To construct L_{i+1} , simply pick one element from each equivalence class of \tilde{L}_{i+1} .

Clearly L_{i+1} has size at most m by construction. To complete the proof, we need to show that at least one element of L_{i+1} is good. First, by the induction hypothesis, at least one element of L_i is good. This implies at least one element $y^* \in \tilde{L}_{i+1}$ is good as well. Consider what happens when we perform the query q_{y^*} . Since y^* is good, q_{y^*} must satisfy all clauses involving the variables $x_1 \dots x_{i+1}$. If there are no clauses involving the remaining variables $x_{i+2} \dots x_n$, then q_{y^*} satisfies the instance, so the oracle will tell us this and we can terminate the algorithm. Otherwise, there is a clause involving some variable in $\{x_{i+2} \dots x_n\}$. When we query q_{y^*} , the worst violated clause will be some clause C_k involving a variable in $\{x_{i+2} \dots x_n\}$, which will be violated with probability $1/2$. So the equivalence class corresponding to C_k will contain a good assignment. Furthermore, since C_k involves one of the variables in $\{x_{i+2} \dots x_n\}$, it will never be returned as the worst violated clause for query $q_{y'}$ for any bad assignment $y' \in \tilde{L}_{i+1}$, because any bad assignment will violate a clause involving $\{x_1 \dots x_{i+1}\}$ by 1, while C_k will be violated only with probability $1/2$. Therefore the equivalence class corresponding to C_k will contain only good assignments. So by picking one assignment from each equivalence class, we will ensure L_{i+1} contains at least one good assignment, as claimed.

The time to construct each list is $O(mn)$, and the algorithm constructs n lists. Hence the algorithm runs in time $O(mn^2)$. \square

6 Hidden 2SAT without repetitions

In this section, we consider a hidden 2SAT formula Φ which is promised to contain no two clauses that are the same. Although Proposition 2 shows that we cannot always hope to learn Φ directly, it does not rule out the possibility of learning some Φ' such that $\text{SAT}(\Phi') = \text{SAT}(\Phi)$. In fact, this is exactly the approach we take.

Theorem 3. *Suppose Φ is a repetition-free H-2SAT instance on n variables. Then, we can generate Φ' such that $\text{SAT}(\Phi') = \text{SAT}(\Phi)$ in $O(\text{poly}(n))$.*

Before proving Theorem 3, we show that a satisfying assignment can be found when Φ is a satisfiable repetition-free 2SAT instance.

Theorem 4. *Suppose Φ is a hidden repetition-free 2SAT instance on n variables. Then it is possible to generate a satisfying assignment in time $O(n^2)$.*

Proof The idea is to attempt to learn each clause present in the formula. Suppose we wish to determine if the clause $(x_i \vee x_j)$ is present in Φ (an analogous procedure works to determine if a 1SAT clause x_i is in Φ). We can assume that the clause is unobscured because the presence of an obscured clause does not affect the set of satisfying assignments. Run the following procedure:

1. First query the oracle with the assignment $x_i = 0, x_j = 0, x_k = 0$ for $k \neq i, j$. If this is a satisfying assignment, then we are done. Otherwise, we know that there must exist a clause of type: (a) $(x_i \vee x_j)$; (b) $(x_i \vee x_k)$ for $k \neq i, j$; (c) $(x_j \vee x_k)$ for $k \neq i, j$; or (d) $(x_{k_1} \vee x_{k_2})$ for $k_1, k_2 \neq i, j$.
2. Now query the oracle with the assignment $x_i = 0, x_j = 0, x_k = 1$ for $k \neq i, j$. As before, if this is satisfying, we are done. Otherwise, we know that there must exist a clause of type: (a) $(x_i \vee x_j)$; (b) $(x_i \vee \bar{x}_k)$ for $k \neq i, j$; (c) $(x_j \vee \bar{x}_k)$ for $k \neq i, j$; or (d) $(\bar{x}_{k_1} \vee \bar{x}_{k_2})$ for $k_1, k_2 \neq i, j$.
3. We can now construct an explicit test for the presence of the clause $(x_i \vee x_j)$. We will propose two fractional assignments to the oracle. If $(x_i \vee x_j)$ is present, then the clause returned each time will be the same. If it is not present, then the returned clause will be different. Formally, query the oracle with the assignment $x_i = 0, x_j = 0, x_k = \frac{1}{4}$ for $k \neq i, j$ and then with the assignment $x_i = 0, x_j = 0, x_k = \frac{3}{4}$ for $k \neq i, j$. Table 1 shows the accompanying violations. It is clear that if $(x_i \vee x_j)$ is present in the formula, then it is returned on both assignments.

	$(x_i \vee x_j)$	$(x_i \vee x_k)$	$(x_j \vee x_k)$	$(x_{k_1} \vee x_{k_2})$	$(x_{k_1} \vee \bar{x}_{k_2})$	$(x_i \vee \bar{x}_k)$	$(x_j \vee \bar{x}_k)$	$(\bar{x}_{k_1} \vee \bar{x}_{k_2})$
1/4	1	3/4	3/4	9/16	3/16	1/4	1/4	1/16
3/4	1	1/4	1/4	1/16	3/16	3/4	3/4	9/16

Table 1: Violation of the clauses based on the fractional assignments of 1/4 and 3/4.

If it is not present, then from the table we can also see that one of the clauses known to exist from our first query must be returned on the 1/4 fractional assignment. However, one of the clauses known to exist from our second query must be returned on the 3/4 fractional assignment. Thus, the clause returned by the oracle changes when $(x_i \vee x_j)$ is not present.

Notice that the above procedure also works to detect all 1SAT and 2SAT clause types. Therefore, if we complete the above procedure with all $O(n^2)$ clause types without finding a satisfying assignment, then we have identified all unobscured clauses in the formula. It is clear that the conjunction of these clauses forms a formula Φ' such that $\text{SAT}(\Phi') = \text{SAT}(\Phi)$. Therefore, we can use any 2SAT algorithm which runs in time $O(n^2)$ on Φ' to find some satisfying assignment of Φ . \square

Proof of Theorem 3 First run the procedure in the proof of Theorem 4. Notice that it either learns all unobscured clauses or produces a satisfying assignment. Therefore, let us assume that we have some satisfying assignment, and without loss of generality, let us assume that it is the all ones assignment. This implies that each clause contains at least one positive literal.

Let us now give a procedure to find all those variables which must necessarily be set to 1; the 1SAT clauses corresponding to these variables will be added to Φ' . To determine whether or not x_i must be set to 1 in Φ , set all queries in Theorem 4 so that $x_i = 0$. If the formula Φ is unsatisfiable, then we know that x_i must be set to 1. If, however, all such variable assignments are still satisfiable, then all remaining clauses that are unobscured by the current assignment must be clauses on 2 variables. Furthermore, since all-ones is a satisfying assignment, these clauses must either be of the form $(x_i \vee x_j)$ or $(x_i \vee \bar{x}_j)$. Clearly, the above procedure ends after a polynomial number of steps. Notice that this corresponds to reaching a “branch point” in a typical 2SAT algorithm.

To simplify the exposition, suppose that after the above procedure the remaining variables are still on x_1, \dots, x_n and that we are querying the formula Φ . Let us first attempt to learn the clauses of the form $(x_i \vee \bar{x}_j)$. Once again, set $x_i = 0$ and $x_j = 1$ and use Theorem 4 to determine if the induced formula is satisfiable. If it is satisfiable, then clearly Φ did not contain $(x_i \vee \bar{x}_j)$. However, if it isn't satisfiable, then indeed $(x_i \vee \bar{x}_j)$ must be present in Φ because the assignment $x_i = 0, x_k = 1$ for all $k \neq i$ satisfies all other 2SAT clauses.

We now only left to learn the clause of the form $(x_i \vee x_j)$. Suppose we propose the assignment $x_i = 0, x_j = 0$, and $x_k = 1$ for all $k \neq i, j$. If the formula is not satisfied then the returned clause type must be one of (a) $(x_i \vee x_j)$; (b) $(x_i \vee \bar{x}_k)$; or (c) $(x_j \vee \bar{x}_k)$; as all such clauses are maximally violated (with probability 1). Now propose the assignment $x_i = 0, x_j = 0$, and $x_k = .5$ for all $k \neq i, j$. Clearly, if $(x_i \vee x_j)$ is not present, the oracle returns a clause that contains both a positive and negative literal. Having already learned clauses of this type, we know that $(x_i \vee x_j)$ must not be present in Φ . Likewise, if $(x_i \vee x_j)$ is present, the oracle returns a clause index which we have not previously learned, from which we can conclude that $(x_i \vee x_j)$ is indeed present. \square

While the above procedure may seem elementary, it acts as a stepping stone to tackle the harder problem of learning an unknown input instance of quantum 2SAT, which is introduced and discussed in the subsequent sections.

7 Quantum SAT Preliminaries

Notations. A quantum system of n qubits is described using a Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \dots \otimes \mathcal{H}_n$ where each \mathcal{H}_i is a two-dimensional Hilbert space of the i^{th} qubit. Vectors in \mathcal{H} are called *pure states* and they describe a state of the system. By adding a subscript i to the vector $|\alpha\rangle$ we indicate that $|\alpha\rangle_i$ is defined in the local Hilbert space \mathcal{H}_i of the i^{th} qubit. Similarly, $|\psi\rangle_{ij}$ denotes a 2-qubit state $|\psi\rangle$ in $\mathcal{H}_i \otimes \mathcal{H}_j$. In any local qubit space \mathcal{H}_i , we pick an orthonormal basis $|0\rangle, |1\rangle$ so that every 1-qubit state $|\alpha\rangle$ can be expanded as $|\alpha\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$. We define its orthogonal

state by $|\alpha^\perp\rangle := \alpha_1|0\rangle - \alpha_0|1\rangle$; ⁶ clearly, $\langle\alpha|\alpha^\perp\rangle = 0$. A standard geometrical representation of the state space of a single qubit is the Bloch sphere which is illustrated, for completeness, in Figure 1.

The Bloch sphere is a geometrical representation of the state space of a single qubit and as the name suggests is a sphere [14]. The antipodal points of the Bloch sphere corresponds to orthogonal states and in general, the north and south poles are usually indexed as $|0\rangle$ and $|1\rangle$ respectively as shown in Figure 1. The points on the surface of the sphere correspond to pure states and the interior of the sphere corresponds to mixed states - probabilistic mixtures of pure states. The center of the sphere is the completely mixed state ($\mathbb{I}/2$) which can be interpreted as an equal parts mixture of any state and its orthogonal state. The interested reader is referred to [14] for more details on the exact correspondence between quantum states and points on the Bloch Sphere.




Figure 1: The Bloch sphere with $|\psi\rangle$

A more general way to describe a quantum state is by its *density matrix*. Density matrices can be viewed as statistical ensembles of pure states that are described by vectors. A density matrix representation a single pure state $|\psi\rangle$ is given by the matrix $\rho = |\psi\rangle\langle\psi|$. General density matrices are given as a convex sum of density matrices of the pure states with the coefficient summing up to 1: $\sigma = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ where $\forall i, p_i \geq 0$ and $\sum_i p_i = 1$. Alternatively, they are defined as semi-definite operators whose trace is equal to 1. For instance, the density matrix $\frac{1}{2}\mathbb{I}$ can be written as $\frac{1}{2}\mathbb{I} = \frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1|$. The state of a quantum system can always be fully specified by a density matrix.

Observables in quantum mechanics are associated with Hermitian operators. The eigenvalues of such an operator correspond to the possible outcomes of a measurement. Given such a Hermitian operator A and a pure state $|\psi\rangle$, the expression $\langle\psi|A|\psi\rangle$ is the *expectation value* of A . It is the result we get if we measure A over many copies of the same state $|\psi\rangle$ and average the result. One can use the Chernoff bound to deduce that, with high probability, if we measure A over $\text{poly}(n)$ copies of a state $|\psi\rangle$, we obtain an approximation to $\langle\psi|A|\psi\rangle$ with an additive error of $1/\text{poly}(n)$.

The expectation value of A with respect to a state which is described by a density matrix ρ is given as $\text{Tr}(\rho A)$. Note that if ρ is given by $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ with $\sum_i p_i = 1$, then $\text{Tr}(\rho A) = \sum_i p_i \langle\psi_i|A|\psi_i\rangle$, which justifies the interpretation of ρ as a statistical ensemble of pure states. Like in the pure state case, using $\text{poly}(n)$ identical copies of ρ , one can estimate the expectation value $\text{Tr}(\rho A)$ up to an additive error of $1/\text{poly}(n)$.

Local Hamiltonians and Quantum SAT. While classically SAT is given as a CSP, quantum kSAT (kQSAT) is defined as a special case of the k -local Hamiltonian problem. A k -local Hamiltonian on n qubits is a Hermitian operator $H = \sum_{e=1}^m h_e$, where each h_e is a *local* Hermitian

⁶There are, of course, continuously many orthogonal states for every $|\alpha\rangle$, so here we simply choose one in a canonical way.

operator acting non-trivially on at most k qubits. Formally, it is written as $h_e = \hat{h}_e \otimes \mathbb{I}_{rest}$, where \hat{h}_e is defined on the Hilbert space of k qubits, and \mathbb{I}_{rest} is the identity operator on the Hilbert space of the rest of the qubits. When it is clear from the context, we often use h_e instead of \hat{h}_e , even while referring to its action on the local Hilbert space.

In physics, k -local Hamiltonians model the local interactions between particles in a many-body system and are the central tool for describing the physics of such systems. The *energy* of the system for every state $|\psi\rangle$ is defined by $E_\psi(H) := \langle \psi | H | \psi \rangle = \sum_e \langle \psi | h_e | \psi \rangle$. The lowest possible energy of the system is called the *ground energy* and is denoted by $E_0(H)$. It is easy to verify that $E_0(H)$ is the lowest eigenvalue of H . The corresponding eigenspace is called the *ground space* of the system, and its eigenvectors are called *ground states*. A central task in condensed matter physics is to understand the properties of the ground space, as it determines the low-temperature physics of the system.

There is a deep connection between the problem of approximating the ground energy of a local Hamiltonian and the classical problem of finding an assignment with minimal violations in a local CSP. In both cases, one tries to minimize a global function that is given in terms of local constraints. This connection is evident if we consider the special case when the local Hermitian operators h_e are given as local *projectors* Π_e . Then for any state $|\psi\rangle$, the local energy $\langle \psi | \Pi_e | \psi \rangle$ is a number between 0 and 1 that can be viewed as a measure to how much the state is ‘violating’ the quantum clause Π_e . When the local energy is 0, the state is inside the null space of the projector Π_e and is said to satisfy the constraint. The total energy of the system, $E_\psi = \langle \psi | H | \psi \rangle = \sum_e \langle \psi | \Pi_e | \psi \rangle$ then corresponds to the total violation of the state $|\psi\rangle$. When the ground energy of the system is 0, necessarily the ground space is the non-vanishing intersection of all the null spaces of the local projectors, and we say that the system is satisfiable. From a physical point of view, such a system is called *frustration-free*, since any ground state of the global system also minimizes the energy of every local term Π_e .

The quantum kQSAT problem is analogous to the classical kSAT problem. Whereas in the kSAT case we are asked to decide whether a k -local CSP is satisfiable or not, in the kQSAT problem we are asked to determine whether the ground energy of a k -local Hamiltonian made of projectors is 0 or not. Unlike the truth values of SAT clauses, however, the ground energy of a k -local Hamiltonian is a continuous function that is sensitive to any infinitesimal change in the form of the local projectors. To make the kQSAT problem more physically relevant, we define it using a promise: Given a k -local Hamiltonian of projectors over n qubits and a value $b > \frac{1}{n^\alpha}$ for some constant α , decide if the ground energy of H is 0 (the YES case) or the ground energy of H is at least b (the NO case). Bravyi [6] showed that kQSAT for $k \geq 4$ is QMA₁-complete while Gosset and Nagaï [10] showed that 3QSAT is also QMA₁-complete. The class QMA₁ stands for ‘Quantum Merlin Arthur’ with one-way error, and is the quantum generalization of the classical MA₁ class with one-way error. The differences are that the witness can be a quantum state over poly(n) qubits, and the verifier can be an efficient quantum machine. In Ref. [6] it was known that 2QSAT has an $O(n^4)$ classical algorithm, and is therefore in P. More recently linear time algorithms for the same problem have been constructed [2, 8].

As the Hamiltonian in a 2QSAT instance is a sum of 2-qubit projectors, every local projector is defined on a 4-dimensional Hilbert space and is of rank 1, 2 or 3. The non-zero subspace of each projector (the subspace on which it projects) is commonly referred to as the *forbidden space* of that projector and the orthogonal subspace is its *solution space*. Finally, we say that H has *no repetitions* if there does not exist any pair of different projectors $\Pi_e, \Pi_{e'}$ which act non-trivially on

the same set of qubits. In the case of repetition free 2QSAT, each projector can also be indexed by the qubit pairs it acts on and the instance can be written as $H = \sum_{(u,v) \in S} \Pi_{uv}$, where $S \subseteq [n] \times [n]$ and each Π_{uv} is non-zero. For any projector Π and a state $|\psi\rangle$, we say that $|\psi\rangle$ *satisfies* Π up to ϵ if $E_\psi(\Pi) := \langle \psi | \Pi | \psi \rangle \leq \epsilon^2$. The energy $E_\psi(\Pi)$ is the *violation energy* of $|\psi\rangle$ with respect to the projector Π . Notice that when the state of the system is described by a density matrix ρ , its violation energy with respect to Π is given by $E_\rho(\Pi) := \text{Tr}(\rho\Pi)$

Finally, a 2QSAT Hamiltonian H is said to have a *Star-like* configuration if there exists a pair of qubits u, v with $\Pi_{u,v} \neq 0$ such that *all* projectors involve either u or v .

Hidden QSAT. The hidden version of QSAT is defined analogously to the classical case. Our task is to decide whether a k -local Hamiltonian $H = \sum_e \Pi_e$ that is made of m k -local projectors over n qubits is frustration-free with $E_0 = 0$ (YES instance) or $E_0 > m \cdot 2\epsilon^2$ (NO instance). Here, $\epsilon > 0$ is some threshold parameter that can be assumed to be inverse polynomially small in n . Moreover, as in H-SAT, here we do not know the Hamiltonian itself; instead we can only send *quantum* states to a “worst violated oracle”, which will return the index e of the projector Π_e with the highest violation energy. Since we want to generalize the notion of a probabilistic assignment that is used in H-SAT, we allow ourselves to send the oracle qubits that hold a general quantum state ρ , which can only be described by a density matrix. Recall from the previous section that this can be regarded as an ensemble of pure quantum states. Then the oracle will return the the index e for which $\text{Tr}(\Pi_e \rho)$ is maximized. If the total energy of the proposed state is $\leq m \cdot \epsilon^2$ then the oracle will indicate that a satisfying assignment has been found.

8 Hidden Quantum 1SAT

The algorithm used to solve H-1SAT can be extended to solve the H-1QSAT problem as well. A 1-local projector defined on \mathbb{C}^2 is satisfiable if it is of rank at most 1 and can be viewed as setting the direction of the qubit on the Bloch sphere. Unlike the classical case, where we may view the 1SAT clauses as either the $|0\rangle\langle 0|$ or $|1\rangle\langle 1|$ projectors, here the projectors can point in any direction in the Bloch sphere. To handle the continuous nature of the Bloch Sphere, we consider discretizing it by using an ϵ -net that covers the whole sphere. This allows us to generalize the lists of 0 – 1 strings used in H-1SAT into lists of n -qubit product states where each qubit is assigned an element of the ϵ -net.

Given a 1-local projector $|\psi\rangle\langle\psi|$, its zero space is spanned by $|\psi^\perp\rangle$. We can divide the Bloch sphere into two hemispheres, one hemisphere containing states $|\phi\rangle$ having $|\langle\psi|\phi\rangle| \leq \frac{1}{2}$ and the other with states having $|\langle\psi|\phi\rangle| > \frac{1}{2}$. An n -qubit state $a = |a_1\rangle|a_2\rangle \dots |a_n\rangle$ is called *good* if for each qubit i , where $|\psi_i\rangle$ is its forbidden state, $|\langle\psi_i|a_i\rangle| \leq \frac{1}{2}$ and *bad* if $\forall i, |\langle\psi_i|a_i\rangle| > \frac{1}{2}$. For the n -qubit state $a = |a_1\rangle|a_2\rangle \dots |a_n\rangle$, let $a' := |a_1^\perp\rangle|a_2^\perp\rangle \dots |a_n^\perp\rangle$.

Now, we can sketch the H-1QSAT algorithm. Adapting the process described in Theorem 2 for an arbitrary n -qubit state a gives a list of n -qubit states, $L_{a/a'}$, where at least one state is *good*. This is formally stated in Lemma 2.

Lemma 2. *Let $a = |a_1\rangle \otimes \dots \otimes |a_n\rangle$ where $|a_i\rangle, |a_i^\perp\rangle$ is a basis for qubits i , for $i = 1, \dots, n$. Then one can produce a list, $L_{a/a'} \subset \bigotimes_{i=1}^n \{|a_i\rangle, |a_i^\perp\rangle\}$ of at most $2mn$ states such that, if the instance is satisfiable, there is at least one good n -qubit state in the list. The time taken to produce this list is $O(n^2m)$.*

Proof Construct the list inductively as in Theorem 2 where at stage k , $L_{k,a/a^\perp}$ contains at most m strings at least one of which is good for qubits $1, \dots, k$ if the instance is satisfiable, by constructing trials as follows. At stage k , replace $\{0, 1\}$ with $\{a_k, a_k^\perp\}$ for qubit k and the value of $\frac{1}{2}$ with $\frac{\mathbb{I}}{2}$, the completely mixed state, for qubits $k+1, \dots, n$ so that $L_{k,a/a^\perp}$ contains states from the set

$$\bigotimes_{i=1}^k \{|a_i\rangle, |a_i^\perp\rangle\} \otimes \left(\frac{\mathbb{I}}{2}\right)^{\otimes(n-k)}.$$

This almost finishes the process except for one caveat when there exists no projector on qubits $k+2, \dots, n$ while constructing $L_{k+1,a/a^\perp}$ from $L_{k,a/a^\perp}$. This situation also occurs at the last step while constructing L_n . In both cases, while proposing a good state, all violations are $\leq \frac{1}{2}$ and any clause ID returned by the oracle involves a qubit in $1, \dots, k+1$. This same clause could also be violated with probability $> \frac{1}{2}$ when a bad string is proposed which will incorrectly be put in the same equivalence class as the good one. Then, picking just one representative from C_j is insufficient and the size of the lists cannot be compressed. To fix this, we add the following checks:

1. If $k+1 = n$, just double the number of strings on the list, assuming that there is a clause involving n , i.e. the last qubit that is assigned values.
2. Repeat the algorithm n times by placing a different qubit at the last position each time. Let L_{a/a^\perp} be the union of all the lists found in this manner and is of size $2m \cdot n$. For a non-empty instance, at least one of the trials in the list will be such that there is a clause on the qubit in the last position and will hence contain a good string.

Hence, L_{a/a^\perp} with $2mn$ strings contains at least one good string and by repeating the classical process n times, we get an $O(mn^2)$ time procedure for this. \square

However, this only gives us an assignment that violates each projector by $\leq \frac{1}{4}$ while we require assignments that violate each projector by $\leq \epsilon^2$. The key observation involves constructing two lists $L_{a/a'}$ and $L_{b/b'}$ where $b \neq a, a'$ and picking a state from each list. Consider the case when both states are good. Let the states on qubit i from each list be $|a_i\rangle$ and $|b_i\rangle$ respectively. Each state defines a hemisphere R_{i,a_i} and R_{i,b_i} containing all the states that are bad with respect to the forbidden state for qubit i , $|\psi_i\rangle$. Then, $|\psi_i\rangle$, should be contained in $R_{i,a_i b_i} := R_{i,a_i} \cap R_{i,b_i}$. The optimal choice for b_i , given a_i , would be one where $|R_{i,a_i b_i}| \leq \frac{|R_{i,a_i}|}{2}$. Then, similar to performing a binary search on the Bloch Sphere, repeating this process $\log_2\left(\frac{1}{\epsilon}\right)$ times, will give a region consisting of good approximations to the forbidden state (See Figures 2 (a) and (b) for illustrations).

Theorem 5. *Let $\epsilon > 0$. Given a H-1QSAT on n qubits containing m projectors, there exists an $O((2mn)^{2 \log \frac{1}{\epsilon}} \cdot mn^2)$ time algorithm, with the property that*

- (a) *for a frustration free instance, it outputs an assignment where for each projector, the forbidden state is violated with probability $\leq \epsilon^2$ and*
- (b) *for a NO instance, the algorithm outputs UNSAT.*

Proof Initially, with no information, for each qubit i , $R_i =$ Bloch sphere. Now the algorithm executes the following steps:

- Start by picking an arbitrary state, say $\bar{a} = |0\rangle^{\otimes n}$, and construct $L_{|0\rangle^{\otimes n}/|1\rangle^{\otimes n}}$ as per the procedure in Lemma 2. For each $a \in L_{|0\rangle^{\otimes n}/|1\rangle^{\otimes n}}$:
 - a defines the region R_{i,a_i} in this branch of the iteration.

- For $i = 1, \dots, n$ pick a basis $\{|b_i\rangle, |b_i^\perp\rangle\}$ such that their equator bisects R_{i,a_i} .
- Set $\bar{b} = b_1 \dots b_n$, construct $L_{\bar{b}/\bar{b}'}$ and for each $b \in L_{\bar{b}/\bar{b}'}$:
 - * The tuples (a, b) define the region $R_{i,a_i b_i}$ in this branch.
 - * Repeat the process to find \bar{c} to bisect each $R_{i,a_i b_i}$;
 - * Find a new region $R_{i,a_i b_i c_i}$ for each $c \in L_{\bar{c}/\bar{c}'}$.
 - * Continue the recursion up to $\log_2(\frac{1}{\epsilon})$ depth and let the last list be L_{z/z^\perp} .
 - * Propose $|\phi^\perp\rangle = \bigotimes_{i=0}^n |\varphi_i^\perp\rangle$ where $\forall i, |\varphi_i\rangle \in R_{i,a_i b_i \dots z_i}$ to the oracle. Output $|\phi^\perp\rangle$ if the oracle returns YES, otherwise continue.
- Output UNSAT if none of the trials satisfy the instance.

This algorithm essentially creates a recursion tree with each new string created where the width of the recursion at each point is $2mn$ and the depth is $\log_2(\frac{1}{\epsilon})$. This leads to $(2mn)^{\log_2 \frac{1}{\epsilon}}$ trials to be proposed at the end and the number of lists created is also $(2mn)^{\log_2 \frac{1}{\epsilon}}$, each at a cost of $O(mn^2)$. Hence, the total running time of the algorithm is $O((2mn)^{\log \frac{1}{\epsilon}} \cdot mn^2)$.




Figure 2: Shown here are four stages of the algorithm for a qubit i , with forbidden state $|\psi_i\rangle$, starting with picking string a, b and c followed by the last string z . (a) After picking $|a_i\rangle$ for qubit i , the hemisphere orthogonal to it is R_{i,a_i} ; (b) On choosing $|b_i\rangle$, the interesting region is the quadrant $R_{i,a_i b_i}$; (c) After $|c_i\rangle$ is determined, $|\psi_i\rangle$ should be in the hatched $1/8^{th}$ region of the sphere; (d) Continuing the process of picking strings for $\log_2 \frac{1}{\epsilon}$ steps and picking the final state $|z_i\rangle$ shows that $|\psi_i\rangle$ should be present in $R_{i,a_i b_i \dots z_i}$.

Figure 2 shows exactly how the algorithm given in Theorem 5 proceeds. We consider a qubit i and the strings a, b, c, \dots, z that are picked in one branch of the recursion tree of the algorithm. To argue the correctness of this algorithm, we analyze region $R_{i,a_i b_i \dots z_i}$ obtained at the leaf of the recursion tree. Let the forbidden state for qubit i be $|\psi_i\rangle$. At the beginning, let $\forall i, |R_i| = 1$ (the complete Bloch sphere) and the only guarantee for each list is that there is at least one good string in it. Tracing the path in the recursion tree to the leaf, let us assume that each step of the recursion picks a good string i.e. a, b, \dots, z are all good strings. For a and $\forall i$, the forbidden state $|\psi_i\rangle$ is in the opposite hemisphere to $|a_i\rangle$ which reduces the size of the region to $|R_{i,a_i}| = 1/2$ as shown in Figure 2(a). Taking (a, b) at the next iteration, the region for each qubit is the overlap of two hemispheres $R_{i,a_i} \cap R_{i,b_i}$ and by construction, since b_i bisects R_{i,a_i} , the overlaps of the hemispheres also bisect R_{i,a_i} setting $|R_{i,a_i b_i}| = 1/4$ (Figure 2(b)). As this pattern continues, each step of the iteration halves the region for qubit i and we are left with regions of size at most ϵ at the end of the branch as shown in Figures 2(c) and 2(d). If the instance is satisfiable, the state proposed will satisfy each projector up to ϵ resulting in the oracle to return YES. Of course, when one of the strings chosen is bad, the proposal $|\phi_j^+\rangle$ for some qubit j will end up having a large inner product with the forbidden state $|\psi_j\rangle$ and will result in the oracle returning the ID of the projector involving j . This concludes the proof. \square

9 Hidden Quantum 2SAT

This section deals with a 2QSAT instance that is hidden and can only be accessed by a worst-violation oracle. We show how learn the underlying local Hamiltonian to precision ϵ by finding 2-local projectors Π'_e such that $\|\Pi_e - \Pi'_e\| \leq \epsilon$ for every projector Π_e . This yields an approximate local Hamiltonian $H' = \sum_e \Pi'_e$ whose ground energy is at most $m\epsilon$ away from the ground energy of the original Hamiltonian $H = \sum_e \Pi_e$. If ϵ is set such that $m\epsilon$ is much smaller than the promise gap of the initial Hamiltonian H (which merely requires $\epsilon < 1/\text{poly}$), then the Hamiltonian H' will have a promise gap as well. This is stated in Theorem 6.

Theorem 6. *Given a H-2QSAT problem $H = \sum_{(u,v)} \Pi_{uv}$ on n qubits, and precision ϵ . If the interaction graph for H is not Star-like, then there is an $O(n^4 + n^2 \log(\frac{1}{\epsilon}))$ algorithm that can find an approximation $H' = \sum_{(u,v)} \Pi'_{uv}$ where $\forall (u, v), \|\Pi'_{uv} - \Pi_{uv}\| \leq \epsilon$*

The algorithm proceeds by:

- (1) Identifying two pairs of qubits $(i, j) \neq (k, \ell)$ on which two projectors are defined, Π_{ij} and $\Pi_{k\ell}$, and finding a constant approximation for these projectors;
- (2) Improving the constant approximation of the two projectors recursively so that the approximation improves by a factor of 2 in each iteration and
- (3) Using the ϵ -approximation of a projector to identify the rest of the independent projectors and approximating them to ϵ -precision.

Using the H' output by the above algorithm in a procedure which could find a good approximation to the ground energy of H' would completely solve H-2QSAT. At this time, though, existing 2QSAT algorithms [2, 6, 8] are not robust to such errors and seem to require $\frac{1}{\exp(n)}$ precision. Our algorithm for H-2QSAT does allow one to learn the projectors to exponential precision, since the

dependence on ϵ in Theorem 6 is merely logarithmic. However, in this parameter regime our algorithm is somewhat unrealistic, as this would require the oracle to be able to distinguish between values that are exponentially close together⁷. If our oracle were constrained to be implementable in polynomial time by an experimenter, acting on polynomially many copies of the proposed state ρ , then one could only learn the instance up to error $\epsilon = \frac{1}{\text{poly}(n)}$. A natural open question is to determine whether one can still solve 2QSAT when one only knows the individual clauses to inverse polynomial precision; we believe this is a fundamental question about the nature of 2QSAT, which is left for future work.

To prove Theorem 6, we require the notion of converting energy violations to distances between states and vica-versa as presented below.

Energies and Distances. Along with the energy of a state with respect to a local term, another useful measure of violation is the distance between quantum states which is helpful in understanding how precise the solution is. We will often switch between distances of 2-qubit states and their violation energy. It is therefore important to relate these two measures.

Assume we have a projector $\Pi_\psi = |\psi\rangle\langle\psi|$ and a state $\rho_\alpha := |\alpha\rangle\langle\alpha|$. Recall that the violation energy is given by $\text{Tr}(\Pi_\psi\rho_\alpha) = |\langle\alpha|\psi\rangle|^2$. The *Frobenius distance* is defined as

$$\|\alpha - \psi\| := \sqrt{\text{Tr}[(|\alpha\rangle\langle\alpha| - |\psi\rangle\langle\psi|)(|\alpha\rangle\langle\alpha| - |\psi\rangle\langle\psi|)^\dagger]}. \quad (1)$$

It now follows that $\|\alpha - \psi\|^2 = 2 - 2|\langle\alpha|\psi\rangle|^2 = 2 - 2\text{Tr}(\Pi_\psi\rho_\alpha)$. Therefore, the violation energy is related to the Frobenius distance as

$$\text{Tr}(\Pi_\psi\rho_\alpha) = 1 - \frac{1}{2}\|\alpha - \psi\|^2. \quad (2)$$

For the sake of clarity, we currently assume that all the projectors in H' are of rank 1 and discuss the entire algorithm for this case. To generalize to cases where the rank of projectors is > 1 requires only a slight modification that does not affect the running time significantly and will be sketched after describing the algorithm in full. Now, we prove a series of lemmas that will enable us to prove Theorem 6.

Before delving into the details of the implementation for Step 1, we construct a procedure **Test-I**, which checks if there is a projector between particles i, j at a constant distance from $|\psi\rangle\langle\psi|$. The test returns YES if and only if there is a projector at i, j which is close to $|\psi\rangle\langle\psi|$ along with the ID of the projector and NO otherwise. The test is defined for two fixed constants $0 < \nu < 1$ and $0 < \epsilon < 1$, and a ν -net \mathcal{T}_ν in the space of rank-1 projectors on two qubits.

Test-I($\mathbf{i}, \mathbf{j}, |\psi\rangle\langle\psi|$)

- For all possible pairs (k, ℓ) which are different from (i, j) :
 - For all $|\alpha\rangle\langle\alpha| \in \mathcal{T}_\nu$, construct the trial state $\rho_\alpha^{k\ell}$, propose it to the oracle and receive a projector ID as violation where

$$\rho_\alpha^{k\ell} := |\psi\rangle\langle\psi|_{ij} \otimes [(1 - \epsilon)|\alpha\rangle\langle\alpha|_{k\ell} + \epsilon|\alpha^\perp\rangle\langle\alpha^\perp|_{k\ell}] \otimes \left(\frac{\mathbb{I}}{2}\right)^{\otimes(n-4)} \quad (3)$$

- If for all trials the oracle answers the *same* ID, output YES together with the ID of the projector, otherwise output NO.

⁷This seems to give the oracle too much power - because having the ability to distinguish exponentially close quantum states would allow one to solve PP-hard problems [1]. In contrast all of the problems considered are in NP due to the presence of classical, $\text{poly}(n)$ size witnesses.

The correctness of **Test-I** is proved in the following claim.

Claim 1. *Assume that $\epsilon + \frac{1}{2}\nu^2 < 1/4$, and further assume that there is at least one projector in the system that is defined on qubits $k \neq \ell$, which are different from i, j . Then, if Test-I outputs YES, there exists a projector $|\psi'\rangle\langle\psi'|$ on the (i, j) qubits, whose ID is the one that was output and*

$$\| |\psi'\rangle\langle\psi'| - |\psi\rangle\langle\psi| \| \leq \sqrt{2\epsilon + \nu^2}. \quad (4)$$

Conversely, if there exists a projector $|\psi'\rangle\langle\psi'|$ on (i, j) such that

$$\| |\psi'\rangle\langle\psi'| - |\psi\rangle\langle\psi| \| < \sqrt{2\epsilon}, \quad (5)$$

then the test will report it.

Proof Assume we have a projector $|\phi'\rangle\langle\phi'|$ on qubits k, ℓ , which are different than i, j . Let $|\alpha\rangle\langle\alpha|$ be the closest member of \mathcal{T}_ν to $|\phi'\rangle\langle\phi'|$, and let us calculate the violation energy for this particular assignment. Note that $\text{Tr}(|\alpha^\perp\rangle\langle\alpha^\perp| \cdot |\phi'\rangle\langle\phi'|) = 1 - \text{Tr}(|\alpha\rangle\langle\alpha| \cdot |\phi'\rangle\langle\phi'|)$ and so the violation due to this assignment is

$$\begin{aligned} & (1 - \epsilon) \text{Tr}(|\alpha\rangle\langle\alpha| \cdot |\phi'\rangle\langle\phi'|) + \epsilon[1 - \text{Tr}(|\alpha\rangle\langle\alpha| \cdot |\phi'\rangle\langle\phi'|)] \\ & = (1 - 2\epsilon) \text{Tr}(|\alpha\rangle\langle\alpha| \cdot |\phi'\rangle\langle\phi'|) + \epsilon \end{aligned}$$

By assumption, $\| |\alpha\rangle\langle\alpha| - |\phi'\rangle\langle\phi'| \| \leq \nu$, and therefore by Eq. (2),

$$1 - \frac{1}{2}\nu^2 \leq \text{Tr}(|\alpha\rangle\langle\alpha| \cdot |\phi'\rangle\langle\phi'|) \leq 1,$$

which implies that

$$1 - \epsilon - \frac{1}{2}\nu^2 \leq \max(k, \ell) \text{ violation energy} \leq 1 - \epsilon.$$

Similarly, by looking at the state from the ν -net that is closest to $|\phi'^\perp\rangle\langle\phi'^\perp|$, we deduce that

$$\epsilon \leq \min(k, \ell) \text{ violation energy} \leq \epsilon + \frac{1}{2}\nu^2.$$

If we got the same answer for all tests then it cannot be due to one of the completely mixed states since there the violation is always $1/4$, and we know that for the maximal α violation is at least $1 - \epsilon - \frac{1}{2}\nu^2 > 1/4$. It cannot also be due to the (k, ℓ) projector since we know that the minimal violation energy there is at most $\epsilon + \frac{1}{2}\nu^2 < 1/4$, i.e., less violated than the mixed projectors. Therefore, it must be the projector at (i, j) .

Moreover, the violation of (i, j) must be at least as big as the maximal (k, ℓ) violation:

$$\text{Tr}(|\psi\rangle\langle\psi| \cdot |\psi'\rangle\langle\psi'|) \geq 1 - \epsilon - \frac{1}{2}\nu^2.$$

Therefore,

$$\| |\psi\rangle\langle\psi| - |\psi'\rangle\langle\psi'| \| = \sqrt{2 - 2\text{Tr}(|\alpha\rangle\langle\alpha| \cdot |\psi'\rangle\langle\psi'|)} \leq \sqrt{2\epsilon + \nu^2}.$$

For the other direction, note that if $\| |\psi\rangle\langle\psi| - |\psi'\rangle\langle\psi'| \| < \sqrt{2\epsilon}$ then by Eq. (2), its violation must satisfy $\text{Tr}(|\psi\rangle\langle\psi| \cdot |\psi'\rangle\langle\psi'|) > 1 - \epsilon$, which is bigger than both the violations of the completely mixed state and the maximal violation of (k, ℓ) . \square

We call a 2-qubit state $|\psi\rangle_{ij}$ δ -good for $\delta = \sqrt{2\epsilon + \nu^2}$ if it is returned during a call to **Test-I** $(i, j, |\psi\rangle\langle\psi|)$. The idea is that if we find an δ -good state for (i, j) , it would be a constant approximation for the projector on (i, j) . Suppose no ϵ -good state is found for (i, j) then we can conclude that there is *no projector* on (i, j) or the interaction graph for H has a *Star-like* configuration.

Step 1

- Set $\nu^2 := \frac{1}{16}$, $\epsilon := \frac{1}{32}$ and $\eta := \frac{1}{4}$
- Repeat for all pairs $i \neq j$ until a δ -good approximation for some Π_{ij} is found.

- For all $|\psi\rangle \in \mathcal{T}_\eta$, perform **Test-I**($i, j, |\psi\rangle\langle\psi|$). If the test is positive, set $|\psi\rangle\langle\psi|$ as the δ -good approximation for Π_{ij} .
- Repeat the above process for all qubit pairs (k, ℓ) that are disjoint from (i, j) till a δ -good approximation for $\Pi_{k\ell}$ is found.

Lemma 3. *If there exist two projectors in H' acting on disjoint pairs of qubits then Step 1 will always succeed in finding independent pairs $(i, j), (k, \ell)$ and projectors $|\psi^{(0)}\rangle\langle\psi^{(0)}|_{ij}, |\phi^{(0)}\rangle\langle\phi^{(0)}|_{k\ell}$ such that $\| |\psi^{(0)}\rangle\langle\psi^{(0)}|_{ij} - |\psi'\rangle\langle\psi'|_{ij} \| \leq \frac{1}{\sqrt{8}}$ and $\| |\phi^{(0)}\rangle\langle\phi^{(0)}|_{k\ell} - |\phi'\rangle\langle\phi'|_{k\ell} \| \leq \frac{1}{\sqrt{8}}$. Moreover, Step 1 can be executed in $O(n^4)$ time.*

Proof Set the parameters according to Step 1. Then from Claim 1, **Test-I** succeeds in finding states that are $\sqrt{2\epsilon + \nu^2}$ -good = $\frac{1}{\sqrt{8}}$ -good approximations. For each **Test-I**($i, j, |\psi\rangle$), iterating over all possible disjoint pairs (c, d) gives us $\binom{n-2}{2}$ pairs to check. The space of 2-qubit states being in $\mathbb{C}^2 \otimes \mathbb{C}^2$ is a space of dimension 4 and for any $\gamma > 0$, a γ -net over $\mathbb{C}^2 \otimes \mathbb{C}^2$ will contain $O\left(\frac{1}{\gamma^4}\right)$ points. So, for $\nu, \eta \in O(1)$, we check only $O(1)$ states in \mathcal{T}_ν for each qubit pair (c, d) and run **Test-I** for $O(1)$ states in \mathcal{T}_η giving a total of $O(n^2)$ trials proposed. To find the second projector on (k, ℓ) we repeat the above process for at most $O(n^2)$ pairs. Hence, Step 1 can be executed in $O(n^4)$ time and the result follows directly. \square

At the end of Step 1 we have two projectors $|\psi^{(0)}\rangle\langle\psi^{(0)}|_{ij}$ and $|\phi^{(0)}\rangle\langle\phi^{(0)}|_{k\ell}$ which are a constant approximation of their hidden counterparts Π'_{ij} and $\Pi'_{k\ell}$ i.e. at a distance $\delta \leq \sqrt{2\epsilon + \nu^2}$. For this step, we construct the procedure **Test-II** to improve their value to $|\psi^{(1)}\rangle\langle\psi^{(1)}|_{ij}$ (resp. $|\phi^{(1)}\rangle\langle\phi^{(1)}|_{k\ell}$) such that it is at a distance $\leq \frac{\delta}{2}$ from Π'_{ij} (resp $\Pi'_{k\ell}$). Then, Step 2 basically repeats this test c times to improve the value to $\Pi_{ij}^{(c)}$ to a distance $\leq \frac{\delta}{2^c}$ from Π'_{ij} and when $c = O(\log n)$, this will get us polynomially close to Π'_{ij} .

From Step 1, we know that Π'_{ij} lies somewhere in a radius of δ around $|\psi^{(0)}\rangle\langle\psi^{(0)}|_{ij}$ and similarly for qubits (k, ℓ) . Let \mathcal{B}_{ij} be the ball of radius δ around $|\psi^{(0)}\rangle\langle\psi^{(0)}|_{ij}$ and correspondingly consider $\mathcal{B}_{k\ell}$. The states will be enumerated over $\mathcal{T}_{\nu'}^{\mathcal{B}}$ which is the ν' -net restricted to some ball \mathcal{B} in the space of rank-1 projectors on 2 qubits.

Set $\nu' = \frac{\nu}{2}$, $\eta' = \frac{\eta}{2}$ and $\epsilon' = \frac{\epsilon}{4}$. The test is defined for values of $\nu', \eta', \epsilon' > 0$ as

Test-II(i, j, k, ℓ)

- For (i, j) , repeat over all $|\psi\rangle\langle\psi| \in \mathcal{T}_{\eta'}^{\mathcal{B}_{ij}}$ and perform **Test-I**($i, j, |\psi\rangle\langle\psi|$) over $\mathcal{T}_{\nu'}^{\mathcal{B}_{k\ell}}$ with parameter ϵ' . If the test outputs YES, set $|\psi^{(1)}\rangle_{ij} = |\psi\rangle$.
- For (k, ℓ) , repeat over all $|\phi\rangle\langle\phi| \in \mathcal{T}_{\eta'}^{\mathcal{B}_{k\ell}}$ and perform **Test-I**($k, \ell, |\phi\rangle\langle\phi|$) over $\mathcal{T}_{\nu'}^{\mathcal{B}_{ij}}$ with parameter ϵ' . If the test outputs YES, set $|\phi^{(1)}\rangle_{k\ell} = |\phi\rangle$

The correctness of **Test-II** is determined by Claim 2.

Claim 2. *If there exists projectors $|\psi^{(l)}\rangle\langle\psi^{(l)}|_{ij}, |\phi^{(l)}\rangle\langle\phi^{(l)}|_{k\ell}$ on qubit pairs $(i, j), (k, \ell)$ that have been approximated to a distance δ , then **Test-II**(i, j, k, ℓ) will give us projectors $|\psi^{(l+1)}\rangle\langle\psi^{(l+1)}|_{ij}$ and $|\phi^{(l+1)}\rangle\langle\phi^{(l+1)}|_{k\ell}$ such that*

$$\left\| |\psi^{(l+1)}\rangle\langle\psi^{(l+1)}|_{ij} - |\psi'\rangle\langle\psi'|_{ij} \right\| \leq \frac{\delta}{2} \quad \text{and} \quad \left\| |\phi^{(l+1)}\rangle\langle\phi^{(l+1)}|_{k\ell} - |\phi'\rangle\langle\phi'|_{k\ell} \right\| \leq \frac{\delta}{2}. \quad (6)$$

Test-II(i, j, k, l) requires $O\left(\left(\frac{\delta}{\eta'}\right)^4 \times \left(\frac{\delta}{\nu'}\right)^4\right)$ trials where δ is the radius of \mathcal{B}_{ij} and \mathcal{B}_{kl} .

Proof Clearly, as **Test-I** works on the complete space of 2 qubit rank 1 projectors, it will also work on the restricted ball of size δ . Since the existence of nontrivial projectors Π'_{ij} and Π'_{kl} has already been determined from Step 1, we are sure to find another state over the new η' -net that approximates the projectors according to the new values. Then, setting the parameters as mentioned in Step 2, the bound follows directly from Claim 1. For the number of trials, since (k, ℓ) is fixed for (i, j) and vica-versa, the trials only iterate over the number of states in the δ -ball of a γ -net over the space of 2-qubit states which contains $O\left(\frac{\delta^4}{\gamma^4}\right)$ states. Substituting for the values of γ gives the required number of trials. \square

Step 2

To approximate the projectors on qubit pairs $(i, j), (k, \ell)$ to polynomial accuracy, collect parameters η, ν and ϵ from Step 1. Set the counter $c = 0$.

- Set $\nu' = \frac{\nu}{2}, \eta' = \frac{\eta}{2}, \epsilon' = \frac{\epsilon}{4}$ and $\delta = \sqrt{2\epsilon + \nu^2}$.
- Let \mathcal{B}_{ij} be the ball of radius δ around $|\psi^{(c)}\rangle\langle\psi^{(c)}|_{ij}$ and correspondingly \mathcal{B}_{kl} .
- Run **Test-II**(i, j, k, ℓ) which output states $|\psi^{(c+1)}\rangle\langle\psi^{(c+1)}|_{ij}$ and $|\phi^{(c+1)}\rangle\langle\phi^{(c+1)}|_{kl}$
- Update the parameters such that $\nu = \nu', \eta = \eta'$ and $\epsilon = \epsilon'$.
- Increment the counter and repeat the process till $c = O(\log n)$.

A crucial part of the analysis for Step 2 is to show that it can be executed in polynomial time. This is ensured by showing that the number of trials for each iteration of **Test-II** actually remains a constant independent of n .

Lemma 4. *Given projectors $|\psi^{(1)}\rangle\langle\psi^{(1)}|_{ij}, |\phi^{(1)}\rangle\langle\phi^{(1)}|_{kl}$ on qubit pairs $(i, j), (k, \ell)$ that have been approximated to a distance δ , Step 2 will successfully find projectors $|\bar{\psi}\rangle\langle\bar{\psi}|_{ij}$ and $|\bar{\phi}\rangle\langle\bar{\phi}|_{kl}$ such that $\| |\bar{\psi}\rangle\langle\bar{\psi}|_{ij} - |\psi'\rangle\langle\psi'|_{ij} \| \leq \frac{1}{\text{poly}(n)}$ and $\| |\bar{\phi}\rangle\langle\bar{\phi}|_{kl} - |\phi'\rangle\langle\phi'|_{kl} \| \leq \frac{1}{\text{poly}(n)}$. Additionally, this step can be executed in $O(\log n)$ time.*

Proof We start the first iteration in Step 2 with the parameters $\delta, \frac{\nu}{2}, \frac{\eta}{2}$ and proceed in each iteration by halving these parameters. In effect, at the t^{th} iteration, the parameters used are $\frac{\delta}{2^{t-1}}, \frac{\eta}{2^t}$ and $\frac{\nu}{2^t}$. The costliest operation in executing Step 2 involves **Test-II** being performed at each iteration. From Claim (2) the t^{th} iteration of **Test-II** can be executed in time expressed via the parameters δ_t, ν_t, η_t as

$$O\left(\frac{\delta_t^8}{\eta_t^4 \nu_t^4}\right) = O\left(\frac{\delta^8}{2^{8(t-1)}} \frac{2^{4t}}{\eta^4} \frac{2^{4t}}{\nu^4}\right) = O\left(\frac{\delta^8}{\eta^4 \nu^4} 2^8\right) \in O(1)$$

where the last inclusion holds as δ, η and ν start as constants. With $O(\log n)$ iterations, this leads to Step 2 being executed in $O(\log n)$ time. Considering the accuracy of the states output, it is clear that in iteration t the projectors output are $\frac{\delta}{2^t}$ close to the projectors in the hidden instance. For $t = O(\log n)$ this translates to a distance of $\frac{\delta}{2^{O(\log n)}} \leq \frac{\delta}{O(n^c)}$ for some constant $c > 0$ and this in turn is written as $\frac{1}{\text{poly}(n)}$ for $\delta \in O(1)$ and the result follows. \square

Step 3

To approximate the remaining projectors, do the following:

- Pick a pair of qubits (u, v) that is independent from at least one of the projectors approximated so far.
- Perform Step 1 to approximate $\Pi_{uv}^{(0)}$ to constant accuracy (if it exists). Otherwise, pick another pair of qubits.
- Let (x, y) be independent of (u, v) such that Π_{xy} has been approximated to $\frac{1}{\text{poly}(n)}$ accuracy. Use $\Pi_{xy}^{(0)}, \dots, \Pi_{xy}^{(O(\log n))}$ to approximate Π_{uv} to $\frac{1}{\text{poly}(n)}$ accuracy as per Step 2.
- Repeat for all possible independent qubit pairs⁸.

Dealing with higher rank projectors. As mentioned earlier, the tests have been clearly designed assuming the presence of only rank 1 projectors. To generalize **Test-I** for projectors of rank > 2 , we don't stop after finding just one state $|\psi\rangle$ that succeeds the test. By continuing to iterate over all 2-qubit states, we can find a constant number of states that span the forbidden subspace and then use any one of the linear algebra techniques to find a basis for that space whose dimensions would give the rank of the projector. This would also approximate the basis up to constant accuracy at the end of **Test-I** as the states that will be returned from the test can be shown as having a low distance (or high violation energy) with at least one of the non-zero components of the high rank projector. Then, repeating **Test-II** for each basis element will successfully approximate each of them to $\frac{1}{\text{poly}(n)}$ accuracy. Note that each of these changes do not significantly affect the runtime of the algorithm or the number of trials proposed.

Now we proceed to the proof of Theorem 6.

Proof of Theorem 6 As outlined previously, putting together the 3 steps gives us the required algorithm. Consider **Test-I** (i, j) contains the projector $|\psi\rangle$ and all the states of the form ρ_{kl}^α used for the test. From Claim 1, we know that $\|\alpha_{kl} - \Pi'_{kl}\| \leq \frac{1}{\sqrt{8}}$ for some (k, l) and some α . Then, any state that is output by **Test-I** should be closer to Π'_{ij} to have a larger overlap with it. In case of (i, j) being disjoint from (k, l) , there is no problem to ensure this.

However, when no projectors independent of (i, j) exist, consider another projector (i, k) . Now, the states used in **Test-I** would be of the form ρ_k^α . Let the reduced density matrix on i with respect to $|\psi\rangle\langle\psi|$ be ρ_i . Then the error threshold used for any state output by **Test-I** in this case would be related to $\max_\alpha \text{Tr}(\rho_i |\alpha\rangle_k \Pi'_{ik}) \leq \text{Tr}(\rho_i \text{Tr}_k(\Pi'_{ik}))$. It is possible for the latter value to be almost 0 in the case that ρ_i almost satisfies Π'_{ik} (e.g. the product state on i satisfies Π'_{ik}). This would lead to an inaccurate error threshold and affect the accuracy of the states output by **Test-I**. This explains the necessity of H' not having a *Star-like* configuration.

For the running time argument, from Lemmas 3 and 4, the running time for finding a projector up to constant accuracy is $O(n^2)$ and to improve the accuracy to $\beta \ll 1$ takes $O(\log \frac{1}{\beta})$ time. Step 3 essentially repeats the combination of (Step 1 + Step 2) for n^2 pairs of qubits and results in an overall running time of $O(n^2(n^2 + \log \frac{1}{\beta}))$. Setting $\beta < \frac{1}{n^c}$ for some constant c , makes the overall runtime $O(n^4 + n^2 \log n) = O(n^4)$. Similarly, the correctness also follows from combining Lemmas 3 and 4. \square

⁸To learn a projector Π_{ik} when Π_{ij} has already been found, set qubit j to the mixed state and choose a different Π_{mn} to use in Steps 1 and 2 for improving the accuracy of Π_{ik}

Theorem 6 clearly excludes the pathological cases of *Star-like* configurations. These are discussed below.

Learning even one of the projectors in the case of the interaction graph being a Star⁹ is impossible. However, the intermediate case when there is exactly one edge in the graph that is not independent of the other edges, seems to have an intermediate albeit slightly unnatural solution. In fact, to explicitly learn some projectors and then solve the instance requires the oracle to distinguish between exponentially small values. In particular, the following lemma holds.

Lemma 5. *Given a H-2QSAT problem $H' = \sum_{(u,v)} \Pi'_{uv}$ on n qubits, $\epsilon \geq \frac{1}{n^\beta}$ for a constant β and a function $f(n) \in \exp(n)$. If there is exactly one edge (i, k) that does not have any independent projector, then there is an $O(n^4 + n^2 \log f(n))$ time algorithm that can find an approximation $H = \sum_{(u,v)} \Pi_{uv}$ where*

$$\forall (i, j), \neq (i, k) \|\Pi'_{ij} - \Pi_{ij}\| \leq \frac{1}{f(n)}$$

and we can find a 2 qubit state $|\Phi\rangle_{ik}$ in $O(\epsilon^4)$ time, such that $\langle \Phi | \Pi'_{ik} | \Phi \rangle \leq \epsilon^2$.

Proof Sketch. For every qubit pair except (i, k) , use Step 1 to find a constant approximation to the hidden projector and repeating Step 2 for $\text{poly}(n)$ iterations, find exponentially close approximations to the hidden projectors. Using these approximations, we can find an $n - 2$ -qubit (resp. $n - 1$ -qubit) state that almost satisfies all projectors except Π'_{ik} following any $O(n + m)$ algorithm to solve 2QSAT [2, 8]. The $n - 1$ qubit state includes either i or k but not both in the satisfiable case. Then, iterating over all 2-qubit (resp. 1-qubit) states on an ϵ -net, and proposing the complete n qubit state to the oracle, will let us find a state that has low overlap with Π'_{ik} . \square

For the remaining case of the Star graph, at the present time, we do not have an algorithm to learn the projectors to any level of accuracy. This is due to the interference of the center of the star with all the projectors skewing the error thresholds used in this type of algorithm. Of course, the brute force technique to find the ground state by iterating over an ϵ -net of all n -qubit states with at most pairwise entanglement leads to an exponential number of trials to be proposed but the power of the oracle doesn't change. Hence, an obvious trade-off between the power of the oracle and the number of trials proposed exists although both techniques currently lead to unnatural algorithmic techniques for the pathological cases.

10 Acknowledgements

Research was supported by the Singapore Ministry of Education and the National Research Foundation by the Tier 3 Grant MOE2012-T3-1-009, by the European Commission IST STREP project Quantum Algorithms (QALGO) 600700, the French ANR Blanc Program Contract ANR-12-BS02-005. A.B. was supported in part by the NSF Graduate Research Fellowship under grant no. 1122374 and by the NSF Alan T. Waterman award under grant no. 1249349. S.Z.'s research was supported in part by RGC of Hong Kong (Project no. CUHK419413).

⁹A star graph is one where no two edges of the graph are independent.

References

- [1] Daniel S. Abrams and Seth Lloyd. Nonlinear quantum mechanics implies polynomial-time solution for NP-complete and $\#P$ problems. *Phys. Rev. Lett.*, 81:3992–3995, 1998.
- [2] Itai Arad, Miklos Santha, Aarthi Sundaram, and Shengyu Zhang. Linear time algorithm for quantum 2SAT. *CoRR*, abs/1508.06340, 2015. To appear in the 43rd International Colloquium on Automata, Languages and Programming.
- [3] Bengt Aspvall, Michael F. Plass, and Robert Endre Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Inf. Process. Lett.*, 8(3):121–123, 1979. Erratum: *Information Processing Letters* 14(4): 195 (1982).
- [4] Xiaohui Bei, Ning Chen, and Shengyu Zhang. On the complexity of trial and error. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *STOC*, pages 31–40. ACM, 2013.
- [5] Xiaohui Bei, Ning Chen, and Shengyu Zhang. Solving linear programming with constraints unknown. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *ICALP (1)*, volume 9134 of *Lecture Notes in Computer Science*, pages 129–142. Springer, 2015.
- [6] Sergey Bravyi. Efficient algorithm for a quantum analogue of 2-SAT. In Kazem Mahdavi, Deborah Koslover, and Leonard L. Brown, editors, *Contemporary Mathematics*, volume 536. American Mathematical Society, 2011.
- [7] S. A. Cook. The complexity of theorem proving procedures. In *Proceedings of the Third Annual ACM Symposium*, pages 151–158, New York, 1971. ACM.
- [8] Niel de Beaudrap and Sevag Gharibian. A linear time algorithm for quantum 2-SAT. *CoRR*, abs/1508.07338, 2015. To appear in 31st Conference on Computational Complexity.
- [9] Shimon Even, Alon Itai, and Adi Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM J. Comput.*, 5(4):691–703, 1976.
- [10] David Gosset and Daniel Nagaj. Quantum 3-SAT is QMA1-complete. *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, 0:756–765, 2013.
- [11] Gábor Ivanyos, Raghav Kulkarni, Youming Qiao, Miklos Santha, and Aarthi Sundaram. On the complexity of trial and error for constraint satisfaction problems. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *ICALP*, volume 8572 of *Lecture Notes in Computer Science*, pages 663–675. Springer, 2014.
- [12] M. R. Krom. The decision problem for a class of first-order formulas in which all disjunctions are binary. *Mathematical Logic Quarterly*, 13(1-2):15–20, 1967.
- [13] L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.
- [14] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.