

A Linear Time Algorithm for Quantum 2-SAT

Niel de Beaudrap^{*1} and Sevag Gharibian^{†2}

1 Department of Computer Science, University of Oxford, Oxford, UK

niel.debeaudrap@cs.ox.ac.uk

2 Department of Computer Science, Virginia Commonwealth University,
Richmond, USA

sgharibian@vcu.edu

Abstract

The Boolean constraint satisfaction problem 3-SAT is arguably the canonical NP-complete problem. In contrast, 2-SAT can not only be decided in polynomial time, but in fact in deterministic linear time. In 2006, Bravyi proposed a physically motivated generalization of k -SAT to the quantum setting, defining the problem “quantum k -SAT”. He showed that quantum 2-SAT is also solvable in polynomial time on a classical computer, in particular in deterministic time $O(n^4)$, assuming unit-cost arithmetic over a field extension of the rational numbers, where n is the number of variables. In this paper, we present an algorithm for quantum 2-SAT which runs in linear time, *i.e.* deterministic time $O(n + m)$ for n and m the number of variables and clauses, respectively. Our approach exploits the transfer matrix techniques of Laumann *et al.* [QIC, 2010] used in the study of phase transitions for random quantum 2-SAT, and bears similarities with both the linear time 2-SAT algorithms of Even, Itai, and Shamir (based on backtracking) [SICOMP, 1976] and Aspvall, Plass, and Tarjan (based on strongly connected components) [IPL, 1979].

1998 ACM Subject Classification F.2.1 Numerical Algorithms and Problems

Keywords and phrases quantum 2-SAT, transfer matrix, strongly connected components, limited backtracking, local Hamiltonian

Digital Object Identifier 10.4230/LIPIcs.CCC.2016.27

1 Introduction

Boolean constraint satisfaction problems lie at the heart of theoretical computer science. Among the most fundamental of these is k -SAT, in which one is given a formula ϕ on n variables, consisting of a conjunction $\phi(x) = C_1 \wedge C_2 \wedge \dots \wedge C_m$ of m clauses, each of which is a disjunction of k literals, *e.g.* $(x_h \vee \bar{x}_i \vee x_j)$ for $1 \leq h, i, j \leq n$. The problem is to determine whether there exists an assignment $x \in \{0, 1\}^n$ which simultaneously satisfies all of the constraints C_i , *i.e.* for which $\phi(x) = 1$. While 3-SAT is NP-complete [6, 22, 16], 2-SAT admits a number of polynomial time algorithms (*e.g.* [7, 20, 10, 2, 23]), the fastest of which require just linear time [10, 2].

In 2006, Bravyi [3] introduced k -QSAT, a problem which generalizes k -SAT, as follows. In place of clauses C_i , acting on k -bit substrings of n bit strings $x \in \{0, 1\}^n$, one considers

^{*} Supported by Centrum Wiskunde & Informatica, European Commission FET-Proactive project Quantum Algorithms (QALGO) 600700, and the UK Quantum Technology Hub project NQIT.

[†] Supported by a Government of Canada NSERC Banting Postdoctoral Fellowship, the Simons Institute for the Theory of Computing at UC Berkeley, and NSF grant CCF-1526189. This project was initiated while SG was visiting Ronald de Wolf and Centrum Wiskunde & Informatica, whom SG thanks for their hospitality.



orthogonal projectors $\bar{\Pi}_i$ which act on k -qubit subsystems of an n -qubit system $|\psi\rangle \in \mathcal{H}^{\otimes n}$, where $\mathcal{H} := \mathbb{C}^2$. (A sketch of how k -SAT can be embedded into k -QSAT is given in Section 2.) These projectors extend to act on states $|\psi\rangle$ by defining $\Pi_i = \bar{\Pi}_i \otimes I$, so that Π_i acts as the identity on all tensor factors apart from those qubits on which $\bar{\Pi}_i$ is defined. One then considers $|\psi\rangle$ to “satisfy” the 2-QSAT instance if $\Pi_i |\psi\rangle = 0$ for all i . This formulation may be motivated, *e.g.*, by problems in many-body physics [9, 4]. While 3-QSAT is complete for QMA₁ [3, 13] (a quantum generalization¹ of NP), 2-QSAT is solvable in deterministic polynomial time [3], using $O(n^4)$ field operations over $\mathbb{C}$.

Given the existence of linear time algorithms for classical 2-SAT, this raises the natural question: *Can 2-QSAT also be solved in linear time?* Our main result in this paper is as follows.

► **Theorem 1.1.** *There exists a deterministic algorithm SOLVE_Q which, given an instance of 2-QSAT, outputs a representation of a satisfying assignment if one exists (presented as a list of one- and two-qubit unit vectors to be taken as a tensor product), and rejects otherwise.*

- *SOLVE_Q halts in time $O(n + m)$ on inputs on n qubits with m projectors (assuming unit-cost operations over $\mathbb{C}$).*
- *Furthermore, SOLVE_Q can produce its output using $O((n + m)M(n))$ bit operations, where $M(n)$ is the asymptotic upper bound on the cost of multiplying two n bit numbers;*
- *If the projectors are all product projectors, the algorithm SOLVE_Q requires only $O(n + m)$ bit operations regardless of what computable subfield $\mathbb{F} \subset \mathbb{C}$ the projector coefficients range over.*

In particular, the setting of product constraints above includes classical 2-SAT: in this case the bit-complexity of our algorithm matches optimal 2-SAT algorithms [10, 2].

► **Remark.** For general instances of 2-QSAT, the $O((m + n)M(n))$ bit-complexity of our algorithm compares favourably with the complexity of extracting a satisfying assignment using Bravyi’s 2-QSAT algorithm, which requires $O(n^4 M(n))$ bit operations if one uses similar algebraic algorithms to ours. In “Significance and open questions” below, we discuss the question of field-operation-complexity vs. bit-complexity, as well as whether our algorithm is tight in terms of bit complexity.

Techniques employed

The origin of this work is the observation that Bravyi’s 2-QSAT algorithm can be thought of as an analogue of Krom’s 2-SAT algorithm [20], which involves computing the transitive closure of directed graphs. Krom’s algorithm repeatedly applies a fixed inference rule for each pair of clauses sharing a variable. The repeated application of the inference rule leads to an $O(n^3)$ time to determine satisfiability and an $O(n^4)$ time to compute a satisfying assignment. Bravyi’s algorithm has the same runtimes, measured in terms of the number of field operations.

This work aims to develop a quantum analogue of Aspvall, Plass, and Tarjan’s (APT) linear time 2-SAT algorithm [2], which reduces 2-SAT to computing the strongly connected components of a directed graph. Note that classically $(\alpha \vee \beta)$ is equivalent to $(\bar{\alpha} \Rightarrow \beta)$ and $(\bar{\beta} \Rightarrow \alpha)$, for literals α and β . APT constructs an *implication graph* G of a 2-SAT instance

¹ Here, Quantum Merlin Arthur (QMA) is the quantum analogue of Merlin-Arthur (MA) in which the proof and verifier are quantum, and QMA₁ is QMA with perfect completeness. Unlike the classical setting, in which MA is known to admit perfect completeness [27, 12], whether QMA=QMA₁ remains open (see *e.g.* [15]).

ϕ , with vertices labelled by literals x_i and $\bar{x}_i$ for each i , and edges $\bar{\alpha} \rightarrow \beta$ and $\bar{\beta} \rightarrow \alpha$ for each clause $(\alpha \vee \beta)$. Then, they show that ϕ is satisfiable if and only if x_i and $\bar{x}_i$ are not in the same strongly connected component of G for any i [2]. As the strongly connected components of G can be computed in linear time [25], this yields a linear time algorithm for 2-SAT.

In the quantum setting, not all n -qubit states can be described by assignments to individual qubits (*e.g.*, entangled states). Fortunately, Chen *et al.* [4] show that we may reduce any instance of 2-QSAT to an instance which is satisfiable if and only if there is a satisfying state, in which qubits have separate assignments (see Section 2 for details). In this setting, there is a natural analogue of the equivalence $(x_i \vee x_j) \equiv (\bar{x}_i \Rightarrow x_j) \wedge (\bar{x}_j \Rightarrow x_i)$ in terms of so-called “transfer matrices” (*e.g.* [3, 21]). For any rank-1 quantum constraint $\Pi_{ij} \in L(\mathbb{C}^2 \otimes \mathbb{C}^2)$ on qubits i and j , there exists a *transfer matrix* $T_{ij} \in L(\mathbb{C}^2)$, such that for any assignment $|\psi_i\rangle$ to qubit i such that $T_{ij}|\psi_i\rangle \neq 0$, the state on qubit j for which the constraint Π_{ij} is satisfied is given by $T_{ij}|\psi_i\rangle$.² (Conversely, for any $T_{ij} \in L(\mathbb{C}^2)$, there is a unique rank-1 orthogonal projector $\Pi_{ij} \in L(\mathbb{C}^2 \otimes \mathbb{C}^2)$ whose nullspace is spanned by $|\psi_i\rangle \otimes T_{ij}|\psi_i\rangle$ for $|\psi_i\rangle$ ranging over $\mathbb{C}^2$.) This suggests a quantum analogue G of an implication graph: For each possible assignment $|\psi\rangle$ to a qubit i , we define a vertex $(i, |\psi\rangle)$, and include a directed edge $(i, |\psi\rangle) \rightarrow (j, |\phi\rangle)$ if there is a transfer matrix T_{ij} (corresponding to some constraint Π_{ij}) such that $T_{ij}|\psi\rangle = c|\phi\rangle$ for some $c \neq 0$. We then ask if for each qubit i , there is a vertex $(i, |\psi_i\rangle)$ which cannot reach any $(i, |\psi'_i\rangle)$ where $|\psi_i\rangle \not\propto |\psi'_i\rangle$. If there are such paths $(i, |\psi_i\rangle) \rightarrow \dots \rightarrow (i, |\psi'_i\rangle)$ for all $|\psi_i\rangle$, this is analogous to x_i and $\bar{x}_i$ being in a common strong component in the APT algorithm.

As it stands, this approach has a shortcoming: In the quantum regime, each qubit has a *continuum* of possible assignments (rather than two), which may generate unbounded orbits in an APT-style algorithm. However, by applying techniques of Laumann *et al.* [21] from the study of phase transitions in random 2-QSAT, we may in some cases reduce the set of possible assignments for a qubit i to one or two. Consider the *interaction graph* G' of a 2-QSAT instance, in which vertices correspond to qubits, and two vertices are connected by an (undirected) edge if the corresponding qubits i and j are subject to a constraint Π_{ij} . Suppose $C = (v_1, \dots, v_t, v_1)$ is a cycle in G' , with transfer matrices $T_{v_i v_{i+1}}$ arising from each constraint $\Pi_{v_i v_{i+1}}$, and compute $T_C := T_{v_t v_1} \cdots T_{v_2 v_3} T_{v_1 v_2}$. If T_C has a non-degenerate spectrum, then the only possible satisfying assignments for v_1 are eigenvectors of T_C [21] (see also Lemma 2.2). In effect, computing T_C “simulates” uncountably many (!) traversals $(i, |\alpha\rangle) \rightarrow \dots \rightarrow (i, |\beta\rangle)$ in G ; restricting to the eigenvectors of T_C corresponds to ignoring vertices in G which are infinitely far from the top of any topological order of G . If we hence describe cycles C with non-degenerate T_C as *discretizing*, this suggests the approach of finding a discretizing cycle at each qubit i , and using it to reduce the number of possible states on i to one or two. This simple principle is the starting point of our work.

Despite this simplicity, some obstacles must be addressed to obtain a linear-time algorithm. In the setting of random 2-QSAT [21], every cycle C is a discretising cycle with probability one, as there is zero probability that either a transfer matrix is singular, or that a product of them has a degenerate spectrum. This allows one to quickly reduce the space of assignments possible for a qubit. In contrast, in our setting (*i.e.*, worst case analysis), we cannot assume

² The usual convention is to describe quantum states by unit vectors in $\mathbb{C}^2$, albeit up to equivalence under multiplication by $z \in \mathbb{C}$ for $|z| = 1$. However, vectors produced via transfer matrices might not be normalised. As we are not explicitly concerned with the probabilities of any measurement outcomes obtained from quantum processes, we represent quantum states by vectors which are equivalent up to multiplication by arbitrary (non-zero) scalar factors.

such a distribution of transfer matrices arising from a 2-QSAT instance. For instance, any constraint Π_{ij} corresponding to a product operator (*e.g.*, a classical 2-SAT constraint) has a singular transfer matrix, which when multiplied with other singular matrices may give rise to a singular cycle matrix. Even if a discretising cycle C does exist using some of the edges $jk, k\ell, \dots$, we may have to traverse those edges multiple times to discover C , which is worrisome for a linear-time algorithm. Furthermore, we must address the case in which there are no discretising cycles at all to get a discrete algorithm started. In order to demonstrate a *linear*-time algorithm for 2-QSAT in the spirit of APT, these problems must be carefully addressed.

Our approach to resolve these issues is as follows. In an instance of 2-QSAT in which all transfer matrices are non-singular, we show that discretising cycles are easy to find if they exist, and that the absence of discretising cycles allows one to easily obtain a satisfying state. If, on the other hand, singular transfer matrices are present, the corresponding product constraints $\Pi_{ij} = |\alpha\rangle\langle\alpha|_i \otimes |\beta\rangle\langle\beta|_j$ *themselves* impose a different discretising influence: If $|\alpha^\perp\rangle$ and $|\beta^\perp\rangle$ are states orthogonal to $|\alpha\rangle$ and $|\beta\rangle$ respectively, then at least one of the assignments $(i, |\alpha^\perp\rangle)$ or $(j, |\beta^\perp\rangle)$ is required for a satisfying assignment. This leads us to adopt an approach of “trial assignments” which is highly reminiscent of another linear-time 2-SAT algorithm due to Even-Itai-Shamir [10], which attempts to reduce to an instance of 2-QSAT with fewer product constraints by determining partial assignments satisfying Π_{ij} . (For simplicity, we also adopt the approach of trial assignments for qubits whose state-space have been reduced by discretizing cycles.) This leads us to our algorithm SOLVE_Q (Figure 1, in Section 4), which combines elements of both the Even-Itai-Shamir [10] and Apsvall-Plass-Tarjan [2] linear-time 2-SAT algorithms as described above.

Our approach can be summarised as follows. Following Chen *et al.* [4], we first preprocess our input instance Π and either determine that Π is unsatisfiable, or obtain a new 2-QSAT instance Π' which is satisfiable by a product state if Π is satisfiable at all. From this point on, our algorithm uses the central notion of a *chain reaction* (CR) (see Section 3): this roughly models the idea that given an assignment $|\psi_i\rangle$ to qubit i , following a sequence of transfer operators according to the implication graph of Π' deterministically results in assignments to a subset of other qubits in the instance. In particular, what we are interested in is finding *conflict-free* CRs, which are CRs that terminate without reassigning a value to a qubit j which conflicts with a previous assignment for j . To exploit conflict-free CRs, we first show a Set-and-Forget Theorem (Theorem 3.6), which essentially says the following: if Π' is satisfiable, then any choice of assignments to a subset S which is prescribed by a conflict-free CR, is also consistent with a global satisfying assignment. Thus, given such a conflict-free CR, we can remove the qubits in S and all constraints acting on it from Π' , reducing to a smaller 2-QSAT instance Π'' which is satisfiable if and only if Π' is. Hence, the problem of deciding Π is reduced to the task of repeatedly finding conflict-free CRs. To show that the discovery of conflict-free CRs may be done in *linear* time, we use three key ideas. First, for any product constraint in the graph, there are two associated CRs C_1 and C_2 ; we show that at least one of these must be conflict-free, or Π is not satisfiable. Second, once all product constraints have been exhausted, our next source of conflict-free CRs is the notion of discretizing cycles. In general, it is *not* true that running a depth-first search in the constraint graph of Π'' will yield a discretizing cycle, even if such a cycle exists! However, we show that if all constraints are *entangled*, then a single depth-first search (per connected component of the interaction graph) indeed suffices to find the discretizing cycle. Third and finally, if no discretizing cycles exist in Π'' , then we show that it is easy to find a conflict-free CR. The resulting algorithm, SOLVE_Q , is presented in Figure 1.

Previous work

There is a long history of polynomial time solutions for classical 2-SAT [24, 7, 20, 10, 2, 23], ranging from time $O(n^4)$ to $O(n + m)$. As we mention above, the most relevant of these to our setting are the algorithms of Even, Itai, and Shamir [10] (based on limited backtracking) and Aspvall, Plass, Tarjan [2] (based on strongly connected component detection).

In contrast, little work has been performed in the quantum setting. Until recently, Bravyi's algorithm was the only explicitly articulated algorithm for 2-QSAT, and requires $O(n^4)$ field operations and $O(n^4 M(n))$ bit operations. Other work on 2-QSAT instead concerns either the structure of the solution space of instances of 2-QSAT [21, 9, 4], or bounds on counting complexity [14, 8].

Propagation of assignments using transfer matrices is present already in Bravyi [3], and the results of Laumann *et al.* [21] allow us to restrict the possibly satisfying states on single qubits by finding discretising cycles. We incorporate these into efficient discrete algorithms for testing possible assignments, and provide a cost analysis in terms of field operations and bit operations. In contrast to the random 2-QSAT setting of [21], we do not assume any particular distribution on constraints.

Note: Very recently, Arad *et al.* [1] independently and concurrently presented an algorithm for 2-QSAT, which also runs in $O(n + m)$ time using unit-cost field operations. The overall structure of our algorithm appears similar to theirs, though our treatment of the key issue of 2-QSAT instances with only entangled constraints appears to use different techniques (in particular, Ref. [1] appears to be based on results of Ji, Wei, Zeng [14] which modify the instance itself, whereas we use ideas of [21] to tackle the existing instance via the concept of discretizing cycles). As well as obtaining an upper bound on field operations matching Ref. [1], we also include an analysis of the bit complexity of our algorithm SOLVE_Q , and in particular indicate how our algorithm matches the asymptotic bit complexity of the best algorithms on classical instances of 2-SAT.

Significance and open questions

From a complexity theoretic perspective, just as k -SAT and MAX- k -SAT are canonical NP-complete problems, Quantum k -SAT and its optimization variant, k -LOCAL HAMILTONIAN [19], are canonical QMA₁- and QMA-complete problems for $k \geq 3$ and $k \geq 2$ respectively [3, 13, 19, 17], thus making their study central to quantum complexity theory. From a many-body physics perspective, quantum k -SAT deals with the study of ground states of *frustration-free* local Hamiltonian systems. Such systems include Kitaev's well-known Toric code Hamiltonian [18] (which is 4-local), whose ground space encodes logical qubits of a topological quantum error correcting code. Our work can hence be viewed as aiming to understand which classical techniques for k -SAT can be generalized to explore the ground spaces of such frustration-free systems.

Bit complexity. We now discuss the number of field operations used by our algorithm, $O(m + n)$, versus the number of bit operations, $O((m + n)M(n))$, in Theorem 1.1. There is no such distinction in the complexity of existing 2-SAT algorithms: As bits have only a finite range of values, traversing a chain of implications in the implication graph poses no precision issues. In the quantum setting, however, such a traversal involves computing products of $O(n)$ transfer matrices over some field extension of the rationals. Trial assignments resulting from these products may require $O(n)$ bits per entry to represent; testing whether two possible assignments are equivalent may involve multiplying pairs of n -bit integers. This is the source

of the $M(n)$ term in the bit complexity estimate of Theorem 1.1. To compare, similar considerations applied to Bravyi’s 2-QSAT algorithm gives an upper bound of $O(n^4 M(n))$ bit operations.

It is not obvious that a faster runtime in terms of bit complexity should be possible in general. As we show in Section 6, it is simple to construct a 2-QSAT instance with $m \in O(n)$ and whose unique product state solution requires $\Theta(n^2)$ bits to write down. Thus, among algorithms which explicitly output the entire solution, our algorithm is optimal up to log factors, taking time $O(nM(n)) \in \tilde{O}(n^2)$ for $M(n) \in O(n \log(n) 2^{O(\log^*(n))})$ [11]. Furthermore, as we also show in Section 6, for any algorithm **A** for 2-QSAT which produces the marginal of a satisfying solution (if one exists) on a single qubit in reduced terms³, there is a linear-time reduction from multiplication of n bit integers to the problem solved by **A**. It follows that such an algorithm **A** must run in time $\Omega(M(n))$. As discussed in Section 6, this implies that unless $M(n) \in O(n)$, there is no general algorithm for 2-QSAT with linear bit complexity if the output is required to be in reduced form.

Linear bit complexity. Theorem 1.1 gives a setting in which our algorithm does have linear bit complexity – when all constraints are product operators. This special case still has essentially quantum features, such as satisfiable instances requiring two-qubit entanglement (which our algorithm treats using techniques described in Section 2), and phase-transitions for satisfiability and counting complexity in randomly sampled instances which match those of 2-QSAT rather than classical 2-SAT [8]. It also includes the classical 2-SAT instances, for which our algorithm has optimal bit-complexity.

Open questions. Our algorithm uses results of Chen *et al.* [4], which shows that any satisfiable instance of 2-QSAT has a solution which is “almost” a product state (our algorithm finds such solutions). In the degenerate case, however, there may also exist satisfying states with long-range entanglement, which may also be of interest to find. As our aim here is to study the optimal computational complexity of 2-QSAT, as opposed to seeking particular types of solutions, we leave this as an open question. We also ask: Is the bit-complexity of $O((n+m)M(n))$ for producing explicit assignments optimal? Is there an $O(M(n))$ upper bound for producing representations of marginals of satisfying assignments?

Organization of this paper

In Section 2, we give notation, definitions, and the basic framework for our analysis (including transfer matrices). Section 3 presents a series of lemmas and theorems to demonstrate how to overcome the obstacles presented in this introduction, and which form the basis of a proof of correctness for our algorithm SOLVE_Q . Section 4 states SOLVE_Q . Section 5 sketches bounds on the runtime of SOLVE_Q in terms of the field operations and bit operations. Additional technical details are deferred to the full version of this paper. Section 6 discusses lower bounds on the bit complexity of 2-QSAT.

2 Preliminaries

We begin by setting notation, stating definitions, and laying down the basic framework for our algorithm, including details on transfer matrices.

³ N.B. Our algorithm SOLVE_Q is not such an algorithm, as the output may include cancellable factors in its representation.

Notation

The notation $:=$ denotes a definition and $[n] := \{1, \dots, n\}$. The vector space of (possibly non-normalised) single-qubit pure states is denoted $\mathcal{H} := \mathbb{C}^2$. For a string $x = x_1 x_2 \dots x_n \in \{0, 1\}^n$, we write $|x\rangle := |x_1\rangle \otimes \dots \otimes |x_n\rangle$. For a vector space $\mathcal{X}$ over $\mathbb{C}$, we write $L(\mathcal{X})$ for the set of linear operators on $\mathcal{X}$. The nullspace of an operator A is denoted $\ker(A)$. For vectors $|\psi\rangle$ and $|\phi\rangle$, we write $|\psi\rangle \propto |\phi\rangle$ if $|\psi\rangle = c|\phi\rangle$ for *non-zero* $c \in \mathbb{C}$; if we wish to also allow $c = 0$, we write $|\psi\rangle \propto^* |\phi\rangle$ instead. The latter two definitions extend straightforwardly to matrices. Given $|\psi\rangle \in \mathcal{H}$, we write $|\psi^\perp\rangle$ for the unique vector (up to scalar factors) which is orthogonal to $|\psi\rangle$.

2.1 Quantum 2-SAT

We now present a formal definition of quantum k -SAT (or k -QSAT).

► **Definition 2.1** (Quantum k -SAT [3]). Let $n \geq k$ be an integer, and $\{\Pi_i\}_{i=1}^m \subset L(\mathcal{H}^{\otimes k})$ be a set of k -local orthogonal projection operators (*i.e.*, of the form $I \otimes \tilde{\Pi}_i$ for k -qubit projectors $\tilde{\Pi}_i$) with coefficients over some number field $\mathbb{F}$.

Decision problem. Does there exist a state $|\psi\rangle \in \mathcal{H}^{\otimes n}$ such that $\Pi_i |\psi\rangle = 0$ for all $i \in [m]$?

Search problem. Produce a description of such a state $|\psi\rangle$ if it exists.

For precision reasons, we require in particular that the coefficients are drawn from a number field (a finite-degree field extension $\mathbb{F} = \mathbb{Q}[\omega]$). We suppose that $\mathbb{F}$ is also specified as part of the input by means of a minimal polynomial $p \in \mathbb{Q}[x]$ for which $\mathbb{F} \cong \mathbb{Q}[x]/p$, together with a specification of how $\mathbb{F}$ embeds into $\mathbb{C}$ [5]. (More details are given in the full version, where the runtime of the algorithm is carefully analyzed.) In the literature for 2-QSAT, one is usually more interested in how the structure of the placement of the projectors Π_i affects the solution space, rather than the complexity of the specification of $\mathbb{F}$ or the coefficients. We therefore suppose that there is some constant K which bounds from above the size of the specification of $\mathbb{F}$, and of the coefficients of the operators Π_i .

We next sketch how a 2-SAT instance ϕ can be embedded into 2-QSAT (cases $k > 2$ are similar). For each clause C on boolean variables (x_a, x_b) , we define an operator $\Pi_C \in L(\mathcal{H}^{\otimes 2})$ of the form $\Pi_C := |c_a\rangle\langle c_a| \otimes |c_b\rangle\langle c_b|$, where $c_a = 1$ if the variable x_a is negated in C , and $c_a = 0$ otherwise; we fix c_b similarly. Then Π_C is satisfied by $|x_a x_b\rangle \in \mathcal{H}^{\otimes 2}$ if and only if C is satisfied by $x_a x_b \in \{0, 1\}^2$. We extend Π_C to an operator on $\mathcal{H}^{\otimes n}$ by taking its tensor product with $I_2 \in L(\mathcal{H}^{\otimes n-2})$ on all tensor factors i apart from $a, b \in [n]$. Performing this for all clauses yields an instance of 2-QSAT, $\{\Pi_C\}$, in which all of the projectors are product operators (as mentioned in Section 1), and which imposes the same constraints on standard-basis vectors $|x\rangle$ as the clauses C impose on $x \in \{0, 1\}^n$. Furthermore, as each Π_C is positive semidefinite and diagonal, any $|\psi\rangle$ for which $\Pi_C |\psi\rangle = 0$ for all clauses C must be a linear combination of vectors $|x\rangle$ which also satisfy $\Pi_C |x\rangle = 0$ for all C . Thus this instance of 2-QSAT is satisfiable if and only if the original instance of 2-SAT is, in which case there is a bijection between the solution space of the 2-SAT instance and a basis for the solution-space of the 2-QSAT instance.

Finally, for a given 2-QSAT instance, we denote by G its (potentially infinite) *implication graph* (defined in Section 1), and by G' its *interaction graph*, whose vertices are labelled by qubits, and with a distinct edge between vertices i, j for each projector acting on them.

Reduction to cases satisfied by product states

We mainly consider product-state solutions to instances of 2-QSAT, in spite of instances (such as those described in “Significance and open questions” in Section 1) in which no product state can be a solution. A paradigmatic example is given by a single constraint $\Pi_* = I_4 - |\Psi^-\rangle\langle\Psi^-|$, where $|\Psi^-\rangle := (|01\rangle - |10\rangle)/\sqrt{2}$; the unique satisfying assignment is the entangled state $|\Psi^-\rangle$. Chen *et al.* [4] nevertheless show that all instances of 2-QSAT are “almost” product-satisfiable in the following sense: The only pairs of qubits (i, j) which are entangled for all satisfying states are those for which the sum of all constraints on (i, j) is an operator S_{ij} of rank 3 (as with Π_* above). We may treat such pairs by imposing the unique assignment $|\psi_{ij}\rangle \in \ker(S_{ij})$, and considering what restrictions this imposes on other qubits k as a result of constraints on (i, k) or (j, k) . If we find no conflicts as a result of all such assignments, we obtain a sub-problem which is either unsatisfiable, or satisfiable by a product state. (We describe this reduction in more detail in Section 4.)

Reduction to rank-1 instances

We may require that all constraints have rank 1 (but possibly with multiple constraints on pairs of qubits), by decomposing projectors Π_{ij} of higher rank into rank-1 projectors $\Pi_{ij,1}, \Pi_{ij,2}, \dots$, for which $\Pi_{ij} = \sum_k \Pi_{ij,k}$. By the preceding reduction to product-satisfiable constraints, there will then be at most two independent constraints acting on any pair (i, j) .

2.2 Transfer matrices

A central tool in this work is the *transfer matrix*, which for product states generalizes the equivalence between $(x_i \vee x_j)$ and $(\bar{x}_i \Rightarrow x_j) \wedge (\bar{x}_j \Rightarrow x_i)$ for bits. Consider a rank-1 constraint $\Pi_{ij} = |\phi\rangle\langle\phi|$ on qubits i and j , where $|\phi\rangle$ has Schmidt decomposition $|\phi\rangle = \alpha|a_0\rangle|b_0\rangle + \beta|a_1\rangle|b_1\rangle$. Then, the *transfer* matrices $\mathsf{T}_{\phi,ij}, \mathsf{T}_{\phi,ji} \in \mathsf{L}(\mathbb{C}^2)$ from i to j and from j to i are respectively given by:

$$\mathsf{T}_{\phi,ij} = \beta|b_0\rangle\langle a_1| - \alpha|b_1\rangle\langle a_0|, \quad \mathsf{T}_{\phi,ji} = \beta|a_0\rangle\langle b_1| - \alpha|a_1\rangle\langle b_0|. \quad (1)$$

(When the state $|\phi\rangle$ is clear from context, we simply write T_{ij} and T_{ji} .) Given any assignment $|\psi_i\rangle \in \mathbb{C}^2$ on qubit i , the transfer matrix $\mathsf{T}_{\phi,ij}$ prescribes which single-qubit states $|\psi_j\rangle$ on j are required to satisfy Π_{ij} , via the constraint $|\psi_j\rangle \propto^* \mathsf{T}_{\phi,ij}|\psi_i\rangle$. If $\mathsf{T}_{\phi,ij}|\psi_i\rangle \neq 0$, then $|\psi_j\rangle$ is uniquely determined (up to equivalence by a scalar factor). This is guaranteed when $|\phi\rangle$ has Schmidt rank 2, as $\mathsf{T}_{\phi,ij}$ then has full rank. On the other hand, if $\mathsf{T}_{\phi,ij}|\psi_i\rangle = 0$, then Π_{ij} is satisfied for any assignment on j , so that j remains unconstrained. This situation may only occur if $|\phi\rangle$ is a product constraint, so that $\mathsf{T}_{\phi,ij}$ has a nullspace of dimension 1. This generalises the effect in the classical setting, that assigning $x_i := 1$ satisfies the constraint $C = (x_i \vee x_j)$, regardless of the value of x_j : the corresponding constraint and transfer matrix are $|\phi\rangle = |00\rangle$ and $\mathsf{T}_{\phi,ij} = -|1\rangle\langle 0|$, respectively.

Walk and cycle matrices

We take the closure of the transfer matrices, under composition along walks in the graph. For any walk $W = (v_1, v_2, \dots, v_k)$ in a graph $G = (V, E)$, multiplying the transfer matrices $\mathsf{T}_{v_{k-1}v_k} \cdots \mathsf{T}_{v_2v_3} \mathsf{T}_{v_1v_2}$ yields a new transfer matrix T_W , which we call the *walk matrix* of W (or *path matrix*, if W is a path). For such a walk W , define $W^R := (v_k, v_{k-1}, \dots, v_2, v_1)$. If a transfer matrix T_W has singular value decomposition $\mathsf{T}_W = s_0|l_0\rangle\langle r_0| + s_1|l_1\rangle\langle r_1|$, one may

show by induction on the length of W that

$$\mathsf{T}_{W^R} = \pm(s_0 |r_1\rangle\langle\ell_1| + s_1 |r_0\rangle\langle\ell_0|), \quad (2)$$

where the sign depends on whether W has odd or even length. In particular, this implies that $\mathsf{T}_W \mathsf{T}_{W^R} = \pm s_0 s_1 I$. Thus $\mathsf{T}_W \mathsf{T}_{W^R} \propto^* I$ for all walks W , with a proportionality factor of zero if and only if T_W is singular. In particular, walk operators can sometimes be composed to represent “cancellation” of edges: For walks $U_1 = W'W$ and $U_2 = W^R W''$, if T_W is invertible, we have $\mathsf{T}_{W'W''} \propto \mathsf{T}_{W''} \mathsf{T}_{W^R} \mathsf{T}_W \mathsf{T}_{W'}$ = $\mathsf{T}_{U_1 U_2}$, representing a form of composition of walks in which repeated edges $(ij)(ji)$ cancel.

For $C = (v, u_1, u_2, \dots, u_k, v)$ a cycle in G , the *cycle matrix* of C at v is just the walk operator T_C arising from the walk from v to itself along C . We consider the cycles C and (e.g.) $C' = (u_1, u_2, \dots, u_k, v, u_1)$ to be distinct as walks; in particular, C and C' may give rise to distinct cycle matrices $\mathsf{T}_{C'} \not\propto \mathsf{T}_C$, which in any case represent operators on the state-spaces of distinct qubits.

Walk operators (and cycle operators in particular) allow us to more easily express long-range constraints implicit in the original projectors Π_{ij} (as one may show by induction):

► **Lemma 2.2** (Inconsistency Lemma). *Let $W = (v, v_1, v_2, \dots, v_\ell, w)$ be a walk in G' with walk operator T_W , and let $|\Psi\rangle \in \mathcal{H}^{\otimes n}$ be a product of single-qubit states $|\psi_v\rangle$ for each $v \in [n]$. If $|\psi_w\rangle \not\propto^* \mathsf{T}_W |\psi_v\rangle$, then at least one constraint Π_{ij} corresponding to an edge in W is not satisfied by $|\Psi\rangle$.*

3 Efficient reductions via trial assignments in 2-QSAT

As outlined in Section 2, we consider rank-1 instances of 2-QSAT which either have a product solution or are unsatisfiable. In this section, we describe a means to incorporate transfer matrices into an efficient algorithm for 2-QSAT via the notion of a *chain reaction*: An EIS-style subroutine for trial assignments.

As in Section 1, we define the *implication graph* of a 2-QSAT instance to be an (infinite) directed graph $G = (V, E)$, where V is the set of pairs $(i, |\psi\rangle)$ for qubits i and (distinct) states $|\psi\rangle \in \mathcal{H}$. There is a directed edge $(i, |\psi\rangle) \rightarrow (j, |\phi\rangle)$ if and only if there is a constraint Π_{ij} with transfer matrix T_{ij} such that $\mathsf{T}_{ij} |\psi\rangle \propto |\phi\rangle$. A “chain reaction” is a depth-first exploration of the nodes of G :

► **Definition 3.1** (Chain reaction (CR)). For a qubit i and state $|\psi_i\rangle \in \mathcal{H}$, to *induce a chain reaction (CR) at i with $|\psi_i\rangle$* means to “partially traverse” G , starting from $(i, |\psi_i\rangle)$ and keeping a record of the vertices $(u, |\psi_u\rangle)$ seen for each u . This traversal is governed by a depth-first search (DFS) in the interaction graph G' , as follows. For each vertex $(u, |\psi_u\rangle)$ visited and each edge $\{u, v\}$ in G' , compute $\mathsf{T}_{uv} |\psi_u\rangle$. If this vector is non-zero, let $|\psi_v\rangle := \mathsf{T}_{uv} |\psi_u\rangle$, and traverse to $(v, |\psi_v\rangle)$ in G . For any vertex $(v, |\psi_v\rangle)$ visited by the CR, we say that the CR *assigns* $|\psi_v\rangle$ to v . In the sequence of vertices in G visited by the CR, we may refer to instances of vertices $(v, |\psi\rangle)$ for a given $v \in V$ as the *first assignment*, the *second assignment*, etc. made to v by the CR.

Edges of G' (and walks in G') which are traversed by the depth-first search (DFS) governing a chain reaction, are also said to be traversed by the chain reaction (CR) itself.

The role of CRs in our analysis is to reveal constraints imposed by transfer matrices in an efficient manner. Specifically, if the DFS in G' which governs the CR encounters a cycle, it will visit a vertex v in G' twice, and so makes “assignments” to v more than once. If these

assignments do not match, we say the CR has a *conflict*. If no such conflicts occur, the CR is called *conflict-free*. (In either case, it does not continue the traversal of the CR from the second, third, *etc.* assignments.) We formalise the intuitive significance of conflicts as follows:

► **Lemma 3.2 (Conflict Lemma).** *If a CR induced at v with $|\psi_v\rangle \in \mathcal{H}$ has a conflict, then no product state $|\Psi\rangle \in \mathcal{H}^{\otimes n}$ for which the state of v is $|\psi_v\rangle$ is a satisfying assignment.*

Proof. A conflict in the CR indicates the presence of two walks W_1 and W_2 in the interaction graph G' , from v to some vertex w , for which $T_{W_1} |\psi_v\rangle \not\propto T_{W_2} |\psi_v\rangle$. It follows from the Inconsistency Lemma (Lemma 2.2) that any product state in which v takes the state $|\psi_v\rangle$ is not satisfying. ◀

With the concept of CRs in hand, we can present the key ideas used by our algorithm. First, conflict-free CRs yield partial assignments, which preserve the satisfiability of the instance defined on the remaining unassigned qubits. Second, if a 2-QSAT instance is satisfiable, then a conflict-free CR can be found efficiently. Our algorithm (presented in Figure 1) essentially operates by repeatedly finding conflict-free CRs, and removing the qubits given assignments by each CR, until either a conflict is detected (in which case we reject), or no unassigned qubits remain (in which case we accept).

3.1 Using conflict-free chain reactions to remove qubits

The main result in this Section is Theorem 3.6 (Set-and-Forget Theorem), which is essentially the converse of Lemma 3.2, and allows us to reduce instances of 2-QSAT by providing partial solutions obtained from a CR induced on a single qubit.

We begin by proving a correspondence between CRs and walk operators, in the sense that if there is a walk $W = (v, v_1, v_2, \dots, w)$ in G' , and if $|\psi_v\rangle \notin \ker(T_W)$, a CR induced at v with a state $|\psi_v\rangle$ should assign $T_W |\psi_v\rangle$ to w . The obstacle here is that the CR might not traverse any of the edges of W before assigning a state to w ; we must then relate W to other walks in G' . We do so as follows.

► **Lemma 3.3 (Unique Assignment Lemma).** *Suppose there exists a state $|\psi\rangle$ and a walk W in G' from v to w such that $T_W |\psi\rangle \propto |\phi\rangle$. Then, for any conflict-free CR induced on v with $|\psi\rangle$, w is assigned $|\phi\rangle$.*

Proof. We show that there is a walk $\tilde{W}$ in G' which is followed by the CR, for which $T_{\tilde{W}} |\psi\rangle \propto |\phi\rangle$. Suppose $W = (v, u_\ell, \dots, u_1, u_0)$ for $u_0 := w$. For each $i \geq 0$, let W_i denote the segment $(v, u_\ell, \dots, u_i)$ of the walk W . Let m be the smallest integer such that the CR traverses W_m . If $m = 0$, then we may take $\tilde{W} = W$ is the walk followed by the CR from v to w . Otherwise, we show a reduction to “deform” W , to obtain walks $W', W'', \dots$, and a decreasing sequence $m > m' > m'' > \dots$, for which the CR follows the walks $W_m, W'_{m'}, W''_{m''}$, *etc.* These walks have successively shorter “tails” of edges which are not followed by the CR: the final such walk $\tilde{W}$ is then one which is completely followed by the CR.

Given that $m > 0$, let $|\psi_m\rangle = T_{W_m} |\psi\rangle$. By hypothesis, the CR does not traverse the edge (u_m, u_{m-1}) , either because $T_{u_m u_{m-1}} |\psi_m\rangle = 0$, or because of an assignment on u_{m-1} . The former implies $T_W |\psi\rangle = 0 \not\propto |\phi\rangle$, contrary to hypothesis. Then there is a walk $W'_{m-1} = (v, u'_r, \dots, u'_m, u_{m-1})$ in G' , which is followed by the CR to make the assignment to u_{m-1} . (Note that the assignments to u_{m-1} made by both W and W'_{m-1} are proportional to one another, as otherwise the CR would have detected a conflict when attempting to traverse edge (u_m, u_{m-1}) during its breadth-first search.) We extend the walk W'_{m-1} to a walk $W' = (v, u'_r, \dots, u'_m, u_{m-1}, \dots, u_1, w)$. The CR has traversed W' at least as far as the vertex

u_{m-1} , missing out fewer edges at the end than it does for W . Furthermore, as the CR is conflict-free, we have $T_{u_1 w} T_{u_2 u_1} \cdots T_{u_m u_{m-1}} |\psi_m\rangle \propto T_{W'} |\psi\rangle$, so that $|\phi\rangle \propto T_W |\psi\rangle \propto T_{W'} |\psi\rangle$ by construction.

Repeating the reduction above yields a walk $\tilde{W}$ in G' which is completely followed by the CR, for which $T_{\tilde{W}} |\psi\rangle \propto |\phi\rangle$ by induction. Then $|\phi\rangle$ is the assignment made to w by the CR. $\blacktriangleleft$

Note that the above result holds regardless of which walk W we consider from v to w , so long as $T_W |\psi\rangle \neq 0$. Thus a conflict-free CR induced at v depends on a consistency between all walk operators, from v to any other given w , relative to the initial assignment $|\psi_v\rangle$. For the case $w = v$, we then have:

► **Lemma 3.4** (Circuit Lemma). *Let W be a closed walk starting and ending at v . If $|\psi_v\rangle$ is not an eigenvector of T_W , then inducing a CR at v with $|\psi_v\rangle$ yields a conflict.*

Proof. By definition, the CR assigns $|\psi_v\rangle$ to v . If the CR is conflict-free, then either $T_W |\psi_v\rangle = 0$ or $T_W |\psi_v\rangle \propto |\psi_v\rangle$, by Unique Assignment (Lemma 3.3). Thus, if $|\psi_v\rangle$ is not an eigenvector of T_W , such a CR will have a conflict. $\blacktriangleleft$

Lemma 3.3 also allows us to decouple the set of vertices given assignments by a CR, from the rest:

► **Lemma 3.5** (Unilateral Lemma). *For any state $|\psi\rangle$ and vertex v , suppose that a CR C_1 induced at v with $|\psi\rangle$ is conflict-free. Let A denote the set of vertices given an assignment by C_1 , and $|\psi_a\rangle$ denote the assignment made by C at a given $a \in A$. Then, for any constraint Π_{ab} for $a \in A$ and $b \in V \setminus A$ and for any $|\phi\rangle \in \mathcal{H}$, $\Pi_{ab}(|\psi_a\rangle \otimes |\phi\rangle) = 0$.*

Proof. For $a \in A$, the CR C_1 must discover a walk $W = (v, v_1, v_2, \dots, v_\ell)$ for $v_\ell := a$, such that for any sub-walk $W_i = (v, v_1, \dots, v_i)$ for $1 \leq i \leq \ell$, we have $T_{W_i} |\psi\rangle \neq 0$. The assignment made to a by C_1 is then $|\psi_a\rangle := T_W |\psi\rangle$ by construction. Conversely, as $b \notin A$, it follows by the Unique Assignment (Lemma 3.3) that all walks W_* in G' from v to w satisfy $T_{W_*} |\psi\rangle = 0$: this holds in particular for the walk $W' = (v, v_1, \dots, a, b)$. Then $T_{ab} |\psi_a\rangle = 0$, which is to say that $\Pi_{ab}(|\psi_a\rangle \otimes |\phi\rangle) = 0$ for all $|\phi\rangle$. $\blacktriangleleft$

The Unilateral Lemma allows us to treat conflict-free CRs as “set-and-forget” subroutines, in which we establish partial assignments on a set of qubits which we may remove from an instance $\mathcal{P} = \{\Pi_{ij}\}_{ij \in E}$ of 2-QSAT, obtaining a simpler, equivalent instance $\mathcal{P}' \subset \mathcal{P}$. Formally, we have the following.

► **Theorem 3.6** (Set-and-Forget Theorem). *Let $\mathcal{P} = \{\Pi_{ij}\}_{ij \in E}$ be an instance of 2-QSAT with interaction graph $G' = (V, E)$. Suppose that C is a conflict-free CR induced at $v \in V$ with $|\psi_v\rangle \in \mathcal{H}$, and let A denote the set of vertices given assignments by C . Let $\mathcal{P}'$ be a 2-QSAT instance obtained from $\mathcal{P}$ by removing all constraints acting on A . Then $\mathcal{P}$ is satisfiable by product states if and only if $\mathcal{P}'$ is.*

Proof. For a given $a \in A$, let $|\psi_a\rangle$ denote the assignment made by C to a . By construction, the states $|\psi_a\rangle$ jointly satisfy all constraints between vertices in A ; and by the Unilateral Lemma (Lemma 3.5), the states $|\psi_a\rangle$ also unilaterally satisfy constraints between vertices in A and vertices in $V \setminus A$. If $\mathcal{P}'$ is satisfiable by a state $|\Phi\rangle = \bigotimes_{v \in V \setminus A} |\phi_v\rangle$, then $\mathcal{P}$ is satisfiable by $|\Psi\rangle = [\bigotimes_{a \in A} |\psi_a\rangle] \otimes |\Phi\rangle$. For the converse, suppose that $\mathcal{P}$ is satisfiable by some state $|\Psi'\rangle = \bigotimes_{v \in V} |\psi'_v\rangle$ (which may not agree with the assignments made by C). Define $|\Psi\rangle = [\bigotimes_{a \in A} |\psi_a\rangle] \otimes [\bigotimes_{v \in V \setminus A} |\psi'_v\rangle]$. Again, $|\Psi\rangle$ satisfies all constraints acting on vertices $a \in A$, and by construction it also satisfies all constraints internal to $V \setminus A$. Then $|\Psi\rangle$ also satisfies $\mathcal{P}$, and its restriction to $V \setminus A$ satisfies $\mathcal{P}'$. $\blacktriangleleft$

3.2 How to find conflict-free chain reactions efficiently

The Set-and-Forget Theorem (Theorem 3.6) provides us with the following approach to find a product assignment for an instance $\mathcal{P}$ of 2-QSAT: (i) pick an unassigned vertex v , (ii) find $|\psi_v\rangle$ such that the CR induced at v with $|\psi_v\rangle$ is conflict-free, and (iii) use this CR to produce a partial assignment, reducing to an instance $\mathcal{P}'$ with fewer qubits. It remains to attempt to find such a state $|\psi_v\rangle$, or determine that none exist, from the continuum $\mathcal{H}$ of single-qubit states.

As we describe in Section 1, and as shown by the Circuit Lemma (Lemma 3.4), it suffices for us to restrict our search for $|\psi\rangle$ to the eigenvectors of T_W for a closed walk W , *e.g.* a cycle. Define a *discretizing cycle* as a directed cycle C (starting and ending at some vertex v) with cycle matrix $T_C \not\propto^* I$. For such cycles, the Circuit Lemma allows us to narrow down our search for $|\psi_v\rangle$ to the eigenvectors of T_C , of which there are at most two. This raises two questions: (1) How to find discretizing cycles efficiently, and (2) how to deal with variables which are not on any discretizing cycle.

As noted in Section 1, product operators complicate the task of detecting discretising cycles, but also provide a second way to narrow the search for assignments $|\psi_v\rangle$ leading to conflict-free CRs.

► **Lemma 3.7 (Product Constraint Lemma).** *In a product-satisfiable instance of 2-QSAT with a rank-1 product constraint projecting onto a state $|\phi_{uv}\rangle = |\gamma_u\rangle \otimes |\gamma_v\rangle$, at least one of the CRs at vertex u or v with states $|\gamma_u^\perp\rangle$ or $|\gamma_v^\perp\rangle$, respectively, is conflict-free.*

Proof. Suppose that the instance is product satisfiable, but that a CR starting at qubit u with state $|\gamma_u^\perp\rangle$ has a conflict. Then by the Conflict Lemma (Lemma 3.2), for any satisfying product state $|\psi\rangle = \bigotimes_{v \in V} |\psi_v\rangle$, we have $|\psi_u\rangle \not\propto |\gamma_u^\perp\rangle$. By construction, we have $|\psi_v\rangle \propto T_{uv} |\psi_u\rangle = |\gamma_v^\perp\rangle \neq 0$. Thus a CR induced at v with $|\gamma_v^\perp\rangle$ will be conflict-free (as otherwise $|\psi\rangle$ cannot be a satisfying assignment, again by the Conflict Lemma). ◀

Using Lemma 3.7 together with the Set-And-Forget Theorem (Theorem 3.6), we may find a partial assignment satisfying any given product constraint; repeating this for all product constraints will either (i) reveal that the original 2-QSAT instance is unsatisfiable, (ii) yield a satisfying assignment for the entire instance, or (iii) yield an equivalent instance of 2-QSAT in which all constraints are projectors onto *entangled* states.

Let us call an instance of 2-QSAT *irreducible* if it has a connected interaction graph G' , and all of its constraints are rank-1 projectors onto entangled states. In such an instance of 2-QSAT, all transfer matrices are invertible. A conflict-free CR induced at any vertex will yield assignments for every other vertex; thus, a single discretizing cycle suffices to determine whether or not the instance is satisfiable. We show that when a discretising cycle is present in such an instance of 2-QSAT, it is easily found:

► **Lemma 3.8.** *Suppose G' is an interaction graph of an irreducible instance of 2-QSAT, which contains a discretizing cycle C . Let $T \subset G'$ be a tree which contains all of the vertices of C . Then there is at least one edge e in C , such that the (unique) cycle in the graph $T \cup \{e\}$ is a discretizing cycle.*

Proof. In the tree T , there is a unique path P_{vw} from any given vertex $v \in V$ to any other connected vertex w . Furthermore, by the irreducibility of the 2-QSAT instance, $T_{P_{vw}}$ is non-singular in each case. Suppose that $C = (v_1, v_2, \dots, v_\ell, v_1)$ is a discretizing cycle in the implication graph G . Consider the closed walk from v_1 to itself in T , given by $W = P_{v_1 v_2} P_{v_2 v_3} \cdots P_{v_\ell v_1}$. By induction, we may show that the truncated walk $W' =$

$P_{v_1 v_2} P_{v_2 v_3} \cdots P_{v_{\ell-1} v_\ell}$ satisfies $T_{W'} \propto T_{P_{v_1 v_\ell}} \propto T_{P_{v_\ell v_1}}^{-1}$ for each ℓ : thus $T_W \propto I$. However, $T_C = T_{v_\ell v_1} \cdots T_{v_2 v_3} T_{v_1 v_2} \not\propto I$ by hypothesis. Then there is an edge vw in C for which $T_{vw} \not\propto T_{P_{vw}}$. Then the unique cycle C' in $T \cup \{vw\}$ contains the path P_{vw} from v to w , as well as the edge vw , and has cycle matrix $T_{C'} = T_{vw} T_{P_{vw}} \propto T_{vw}^{-1} T_{P_{vw}} \not\propto I$. $\blacktriangleleft$

► **Theorem 3.9** (Cycle Discovery Theorem). *Suppose G' is the interaction graph of an irreducible instance of 2-QSAT, and contains a discretizing cycle C . Then a depth-first search from any vertex $v \in V$, in which each edge is traversed at most once, suffices to discover a discretizing cycle C' .*

Proof. Consider a DFS starting from any vertex $v \in V$. Define a tree $T \subset G$, in which each edge e traversed by the DFS is included if and only if e is traversed for the first time some vertex is visited. As the DFS reaches each vertex w , it also computes the path operator $T_{P_{vw}}$ for the path taken from v to w . Each time the DFS traverses an edge uw from some vertex u to a vertex w which it has previously visited, it tests whether $T_{P_{vu}} \propto T_{uw} T_{P_{vw}}$. If so, it continues the DFS from w . Otherwise the cycle C' consisting of $P_{vu}^R P_{vw}$ followed by wu is discretizing, as $T_{C'} \propto T_{uw} T_{P_{vu}} T_{P_{vw}}^{-1} \not\propto I$. Conversely by Lemma 3.8, if G has a discretizing cycle, the DFS must eventually traverse such an edge. $\blacktriangleleft$

Implicit in Theorem 3.9 is a linear-time algorithm for finding discretising cycles in an irreducible instance of 2-QSAT, when one is present. It remains to describe how to treat irreducible instances which have no discretizing cycles. The absence of any means of discretising the state-space of any qubit in such an instance actually represents freedom of choice in this case; while this is implicit in Refs. [3, 21, 9], we prove it here for the sake of completeness.

► **Lemma 3.10** (Free Choice Lemma). *In an irreducible instance of 2-QSAT with no discretizing cycles, any choice of single-qubit state $|\psi_v\rangle$ for some v in the component gives rise to a conflict-free CR.*

Proof. Let G' be the interaction graph. Consider a CR induced at v with $|\psi_v\rangle$, and consider the paths P_{vw} to each vertex w , by which the CR makes its first assignment $|\psi_w\rangle := T_{P_{vw}} |\psi_v\rangle$ to w . If P'_{vw} is another walk from v to w , we have $T_{P_{vw}}^R T_{P'_{vw}} \propto I$, from the hypothesis that there are no discretising cycles; then $T_{P'_{vw}} \propto T_{P_{vw}}$. Thus, regardless of the choice of $|\psi_v\rangle$, a consistent assignment $T_{P'_{vw}} |\psi_v\rangle$ is computed every time the CR traverses an edge to visit w . $\blacktriangleleft$

4 A linear-time 2-QSAT algorithm

We finally present our 2-QSAT algorithm in Figure 1, whose correctness follows immediately by combining the results of Section 3. Following [10], we implement CRs (corresponding to their trial assignments) in parallel to ensure a linear bound on run-time; this is expanded upon in Section 5.

Preprocessing stage to impose input constraints

For conciseness, we present SOLVE_Q in Figure 1 with restrictions on the inputs it takes. As we indicate in Section 2, following Chen *et al.* [4], these restrictions ensure that the instance presented to SOLVE_Q is either satisfiable by a product state or unsatisfiable. These restrictions can be imposed through a pre-processing phase, as follows. For each pair $\{u, v\}$ subject to multiple constraints, sum the projectors to obtain positive semidefinite operator S_{uv} . Then perform the following:

Input: An instance of 2-QSAT consisting of rank-1 projectors $\mathcal{P} = \{\Pi_{ij}\}$ with interaction graph $G' = (V, E)$, with at most two parallel edges (u, v) per distinct $\{u, v\} \subset V$.

1. *Discretize on product constraints* – While there exists a projector $\Pi_{ij} = |\phi_{ij}\rangle\langle\phi_{ij}|$ such that $|\phi_{ij}\rangle = |\gamma_i\rangle \otimes |\gamma_j\rangle$ is a product state: simulate CRs at each $v \in \{i, j\}$ with $|\gamma_v^\perp\rangle$, in parallel.
 - a. If conflicts arise in both CRs, halt and output “UNSAT”.
 - b. Fix the assignments for the first conflict-free CR that terminates, remove the set A of vertices that it visited from G' , and go to Step 1.
2. *Discretize on cycles* – While there exists $v \in V$: search for a discretizing cycle $C \subseteq G'$ in the same connected component of v .
 - If such a cycle C is found at a vertex u : Let T_C be its cycle matrix, and S denote the set of eigenvectors of T_C . Simulate CRs at u with each $|\psi_u\rangle \in S$, in parallel.
 - a. If conflicts arise in both CRs, halt and output “UNSAT”.
 - b. Fix the assignments for the first conflict-free CR that terminates, remove the set A of vertices that it visited from G' , and go to Step 2.
 - If no such cycle is found: Induce a CR at v with $|\psi_v\rangle := |0\rangle$. Fix assignments made by the CR, remove the set A of vertices that it visits from G' , and go to Step 2.
3. *Normalize* – For each qubit v , compute whether the assignment $|\psi_v\rangle$ is normalised: if not, compute a normalised version $|\psi_v\rangle := |\psi_v\rangle / \sqrt{\langle\psi_v|\psi_v\rangle}$.

Output: “UNSAT”, or unit vectors $|\psi_v\rangle \in \mathcal{H}$ for each $v \in V$ which jointly satisfy $\mathcal{P}$.

■ **Figure 1** An algorithm for 2-QSAT, denoted SOLVE_Q .

1. If any pair $\{u, v\}$ has $\text{rank}(S_{uv}) = 4$, halt with output UNSAT (as $\ker(S_{uv})$ contains no states).
2. For each pair $\{u, v\}$ with $\text{rank}(S_{uv}) = 2$, replace the constraints on $\{u, v\}$ with $\Pi_{uv,1} = |\eta_1\rangle\langle\eta_1|$ and $\Pi_{uv,2} = |\eta_2\rangle\langle\eta_2|$, for linearly independent columns $|\eta_1\rangle, |\eta_2\rangle$ of S_{uv} .
3. For each pair $\{u, v\}$ with $\text{rank}(S_{uv}) = 3$, record the unique state $|\psi_{uv}\rangle$ which spans $\ker(S_{uv})$ as a joint assignment to (u, v) , and remove the constraints on $\{u, v\}$. If $|\psi_{uv}\rangle = |\psi_u\rangle \otimes |\psi_v\rangle$, record $|\psi_u\rangle$ and $|\psi_v\rangle$ as assignments to u and v respectively. (If any qubit is subject to conflicting assignments, halt with output UNSAT.)
4. For each pair $\{u, v\}$ given an assignment $|\psi_{uv}\rangle$ in the preceding step:
 - If $|\psi_{uv}\rangle = |\psi_u\rangle \otimes |\psi_v\rangle$, induce CRs (sequentially) at u with $|\psi_u\rangle$ and at v with $|\psi_v\rangle$.
 - If not, and there are non-product constraints Π_{iu} or Π_{iv} for any i , halt with output UNSAT (as any state of i is compatible only with product states on $\{u, v\}$). Otherwise, for each $\Pi_{iu} = |\gamma_i\rangle\langle\gamma_i| \otimes |\gamma_u\rangle\langle\gamma_u|$ or $\Pi_{iv} = |\gamma_i\rangle\langle\gamma_i| \otimes |\gamma_v\rangle\langle\gamma_v|$, induce a CR (sequentially) at i with $|\gamma_i^\perp\rangle$.

For any CR induced, halt (with output UNSAT) either if the CR has a conflict, or if it makes an assignment to some other qubit w which has been given a different assignment as a result of a rank-3 constraint. If no conflict is detected, record the assignments, and remove the set A of qubits given assignments from G' .

This preprocessing phase involves much the same subroutines as SOLVE_Q itself, and does not contribute to the asymptotic run-time. (We include these steps in our detailed runtime analysis in the full version.)

5 Runtime analysis

We briefly sketch the runtime analyses for SOLVE_Q in terms of field operations over $\mathbb{C}$ and bit operations, and discuss an optimization for the setting of product state constraints. A more in-depth treatment is given in the full version. We assume a random-access machine, so that memory access takes unit time. The constraints Π_i are specified as 4×4 matrices with coefficients from a finite-degree field extension $\mathbb{F}:\mathbb{Q}$, whose specification is also part of the input; arithmetic operations over such number fields can be performed efficiently [5]. From this representation we extract the basis vectors $|\eta_i\rangle$ for the image of Π_i by taking columns of Π_i , and omit normalisation: SOLVE_Q then uses $|\eta_i\rangle$ to represent Π_i . Vectors are only normalised as the final step of the algorithm.

Field operations

SOLVE_Q requires $O(n + m)$ operations over $\mathbb{C}$, for n and m the number of variables and clauses, respectively. As each vector $|\eta_i\rangle$ is in $\mathbb{C}^4$, operations on them (such as determining if $|\eta_i\rangle$ is a product constraint in Step 1) require $O(1)$ field operations. Following EIS [10], we simulate CRs in parallel by interleaving their steps, terminating both simulations as soon as one of them is found to be conflict-free. In the preprocessing phase and in Step 1b, this ensures that the number of vertices and edges removed (upon completion of a conflict-free CR) is proportional to the number of vertices and edges visited during the parallel CRs. Hence, the total number of edge-traversals of SOLVE_Q is $O(m)$. Finally, by Step 2, the instance has been simplified to a disjoint union of irreducible instances. Theorem 3.9 ensures that if a discretizing cycle exists in any of the components, it can be found by a depth-first search; moreover, a single conflict-free CR suffices to assign satisfying states to all vertices in each component.

Bit operations

The bit-complexity of SOLVE_Q differs from the field-operation complexity, for the simple reason that multiplying k transfer matrices yields a path matrix with $O(k)$ -bit entries. Thus, operations such as determining the eigenvectors of such matrices, or whether $|\psi\rangle \propto |\phi\rangle$ for vectors in the image of these matrices, can take time $O(M(k))$, where $M(k)$ is the time to multiply two k -bit integers. This follows from the fact that computing $\sqrt{D} \in \mathbb{Z}$ for a perfect square $D \in \mathbb{Z}$ can be performed in $O(M(n))$ time using Newton's method (see *e.g.* Theorem 9.28 of Ref. [26]); and that equality testing over $\mathbb{Q}$ is bounded by $O(M(n))$, for rationals $r, s \in \mathbb{Q}$ with n bit representations as ratios. (To test whether $\frac{a}{b}$ and $\frac{c}{d}$ are equal, one tests whether $ad - bc = 0$.) Since the number of times we might need to compute eigenvectors or decide proportionality may scale as $m + n$, the runtime of $O((m + n)M(n))$ follows.

It may be necessary for SOLVE_Q to represent its output using further field extensions $\mathbb{E}:\mathbb{F}$, for instance, when solving the characteristic polynomial $\det(\lambda I - T_C)$ of a cycle matrix T_C , if the discriminant $D = (\text{Tr } T_C)^2 + 4(\det T_C)$ is not a perfect square in $\mathbb{F}$. However, by the Set and Forget Theorem 3.6, any extension required by a CR will be independent of the CRs involved in the assignments made by other CRs; furthermore, the extensions involved in each CR is only quadratic, and specifically by a square root $\sqrt{D}$ of an element $D \in \mathbb{F}$.

The approach taken to the quadratic extensions by SOLVE_Q is unconventional. Specifically, unless $D \in \mathbb{Q}$, we do not evaluate whether or not $\sqrt{D}$ is in $\mathbb{F}$ before defining the (possibly trivial) “extension” $\mathbb{E} = \mathbb{F}[\sqrt{D}]$. That is, we allow representations of number fields

$\mathbb{F}_k := \mathbb{Q}[\omega_1, \omega_2, \dots, \omega_k]$ in which $\omega_j = \sqrt{\alpha_j}$ for some $\alpha_j \in \mathbb{F}_{j-1}$ (possibly including the case $\omega_1 = \sqrt{s}$ for $s \in \mathbb{Q}$), and where it may come to pass that $\omega_j \in \mathbb{F}_{j-1}$. This prevents us from easily presenting coefficients in a normal form: crucially however, it is still possible for us to perform equality tests and arithmetic operations in time $O(M(n))$, for $\alpha \in \mathbb{F}_k$ expressed as $\frac{1}{\mu}f(\omega_1, \dots, \omega_k)$ for $\mu \in \mathbb{Z}$ and $f \in \mathbb{Z}[x]$ with coefficients of size $O(n)$, provided that k is bounded by a constant. (In the case of SOLVE_Q , we bound $k \leq 3$.)

Thus while the output of SOLVE_Q may not be reduced, it nevertheless presents exact, normalised, satisfying states by means of tensor factors. Complete details are to be found in the full version.

Reduced complexity of 2-QSAT for product constraints

Using a simple optimization which exploits product constraints, SOLVE_Q can in fact accept inputs over any field extension $\mathbb{F}:\mathbb{Q}$ (algebraic or otherwise), and solve them with $O(n+m)$ bit operations provided that all projectors are product operators. This requires only that arithmetic operations and equality testing against 0 can be performed in $\mathbb{F}$ in $O(1)$ time on inputs with representations of size $O(1)$. Specifically: the transfer matrix of a product constraint $\Pi_{uv} = |\gamma_u\rangle\langle\gamma_u| \otimes |\gamma_v\rangle\langle\gamma_v|$ is $T_{uv} \propto |\gamma_v^\perp\rangle\langle\gamma_u|$, whose image is spanned by $|\gamma_v^\perp\rangle$. For any assignment $|\psi_u\rangle$ to u , if $T_{u,v}|\psi_u\rangle \neq 0$, we can set v to $|\gamma_v^\perp\rangle$ (which by assumption on the input requires $O(1)$ bits), as opposed to the potentially more complex vector $T_{u,v}|\psi_u\rangle \propto |\gamma_v^\perp\rangle$. Thus, in Step 1, the complexity of the assignments made by a CR are no more complex than the vectors of the projectors Π_{uv} in the input, so that all algebraic operations may be performed in $\Theta(1)$ time rather than $O(M(n))$ time. In particular, for classical 2-SAT instances, we recover an $O(m+n)$ upper bound on the bit-complexity of SOLVE_Q , matching the asymptotic performance of the APT and EIS algorithms [2, 10].

6 On lower bounds for bit complexity

Most investigations into 2-QSAT are presented in terms of unit-cost operations over some algebraic number field $\mathbb{F}$. As a result, no restrictions are usually put on how the output of a classical solution to 2-QSAT is represented. To consider lower bounds on the bit-complexity of presenting a solution to 2-QSAT, it becomes necessary to consider what restrictions to impose on the output, as without such restrictions the notion of what form the output may take becomes ill-defined. We impose the restriction of outputs which are *rationalised*, as follows. Let $\mathbb{F} = \mathbb{Q}[\omega]$ be an algebraic number field, so that ω is an algebraic number whose minimal polynomial p is a monic polynomial over $\mathbb{Z}$. An element $\alpha \in \mathbb{F}$ is presented in *rationalised form* by an expression of the form $f(\omega)/D = \alpha$, where $D > 0$ is an integer and $f \in \mathbb{Z}[x]$ is a polynomial such that $\deg(f) < \deg(p)$. Despite the unconventional representation described in Section 5, this is one constraint which the output of SOLVE_Q respects.

There are further restrictions which one might consider, such as the output state vectors being normalised (which SOLVE_Q satisfies), and that they be *reduced*: that the coefficients $\alpha = f(\omega)/D$ satisfy $\gcd(f, D) = 1$. Consider, for instance, an algorithm which produces its output in *minimal form*: each state that it outputs is normalised, in reduced rationalised form, and involves the minimal field extension $\mathbb{F}:\mathbb{Q}$ necessary to do so, represented as $\mathbb{F} = \mathbb{Q}[\omega]$ where the minimal polynomial of ω is a monic polynomial over the integers. While SOLVE_Q does not compute outputs in minimal form (*e.g.*, it may fail to produce outputs in reduced form), we show that the multiplication time $O(M(n))$ for n bit integers is a relevant lower

bound for algorithms which do, suggesting that the role of $M(n)$ in the upper bound of SOLVE_Q is not merely accidental.

► **Lemma 6.1.** *There exist instances of 2-QSAT on n vertices and $m \in O(n)$ clauses, such that exhibiting a requested tensor factor of a satisfying solution, in minimal form, requires $\Omega(M(n))$ bit operations in the worst case.*

Proof. Let M and N be positive, odd n -bit integers, with binary expansions $M = \sum_{t=0}^{n-1} 2^t M_t$ and $N = \sum_{t=0}^{n-1} 2^t N_t$, where $M_i, N_i \in \{0, 1\}$ for each $0 \leq i < n$. We construct an instance of 2-QSAT whose unique product state solution is one in which one of the qubits q is assigned a state

$$|\psi_q\rangle := \frac{M}{\sqrt{D}} |0\rangle + \frac{2^n + MN}{\sqrt{D}} |1\rangle = \frac{M\sqrt{D}}{D} |0\rangle + \frac{(2^n + MN)\sqrt{D}}{D} |1\rangle, \quad (3)$$

where $D = M^2 + (2^n + MN)^2$. Either the middle or the right-hand expression in Eqn. (3) is in rationalised and normalised form, depending on whether D is a perfect square. As M , $2^n + MN$, and D are coprime, that rationalised expression is in reduced form, if $\mathbb{F} = \mathbb{Q}[\sqrt{D}]$. If D is neither a perfect square nor square-free, it may be that $\sqrt{D}$ is represented as $\delta\sqrt{D'}$ in $\mathbb{Q}[\sqrt{D}]$, where $D = D'\delta^2$. In this case, by hypothesis, a representation of $|\psi_q\rangle$ in reduced form would be identical (up to signs) to

$$|\psi_q\rangle = \frac{M\sqrt{D'}}{D'\delta} |0\rangle + \frac{(2^n + MN)\sqrt{D'}}{D'\delta} |1\rangle. \quad (4)$$

In any case, the minimal form representation would provide a specification of the extension element $\sqrt{D'}$, the denominators $D'\delta$, and the numerators $A = M$ and $B = 2^n + MN$ (or $A = -M$ and $B = -2^n - MN$, which yields an equivalent vector in $\mathbb{Q}[\sqrt{D'}]$). From such a representation, one could compute MN simply as $B - 2^n$ (or $-B - 2^n$ respectively), which requires time $O(n)$.

The instance we construct is on a chain of $2n + 2$ qubits, labelled $v \in \{0, 1, 2, \dots, 2n+1\}$, as follows. For $1 \leq i \leq n$, we define matrices

$$\mathsf{T}_{i-1,i} = \begin{pmatrix} 1 & 0 \\ M_{n-i} & 2 \end{pmatrix}, \quad \mathsf{T}_{n+i,n+1+i} = \begin{pmatrix} 1 & 0 \\ N_{n-i} & 2 \end{pmatrix}; \quad (5)$$

and also two matrices $\mathsf{T}_{n,n+1}$ and $\mathsf{T}'_{0,1}$:

$$\mathsf{T}_{n,n+1} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \mathsf{T}'_{0,1} = \begin{pmatrix} 0 & 1 \\ 0 & M_{n-1} \end{pmatrix}. \quad (6)$$

For each $i \in \{0, 1, 2, \dots, 2n\}$, we include a constraint $\Pi_{i,i+1}$ between qubits i and $i + 1$, with transfer matrix $\mathsf{T}_{i,i+1}$; and we also include a second constraint $\Pi'_{0,1}$ between 0 and 1, with transfer matrix $\mathsf{T}'_{0,1}$. The resulting instance of 2-QSAT has two rank-1 constraints between qubits 0 and 1, and one rank-1 constraint between all other consecutive pairs of qubits. By Chen *et al.* [4], this instance is then satisfiable by a product state if it is satisfiable at all. It is easy to show that all of the projectors have rational coefficients in this case, so we take the field of the representation to be $\mathbb{Q}$ itself.

We show that there is a unique product state which satisfies the above instance of 2-QSAT. It is easy to show that the opposite transfer operator to $\mathsf{T}'_{0,1}$ is

$$\mathsf{T}'_{1,0} \propto \begin{pmatrix} -M_{n-1} & 1 \\ 0 & 0 \end{pmatrix} \quad (7)$$

so that $T'_{1,0}T_{0,1} \propto |0\rangle\langle 1|$. The only eigenvector of this operator is $|0\rangle$, which is therefore the only single-qubit state on qubit 0 which is consistent with a satisfying solution. As all other transfer operators are non-singular, this determines a unique assignment for all other qubits i in the chain, determined by the first column of the walk operator $T_{[0,i]} := T_{i-1,i} \cdots T_{1,2}T_{0,1}$. It is easy to show for $1 \leq i \leq n$ that

$$T_{[0,i]} = \begin{pmatrix} 1 & 0 \\ \sum_{1 \leq t \leq i} M_{n-t} 2^{i-t} & 2^i \end{pmatrix}, \quad (8)$$

and that in particular

$$T_{[0,n]} = \begin{pmatrix} 1 & 0 \\ M & 2^n \end{pmatrix}; \quad (9)$$

from this we easily obtain

$$T_{[0,n+1]} = \begin{pmatrix} M & 2^n \\ 1 & 0 \end{pmatrix}; \quad (10)$$

from which point we may prove by induction for $1 \leq i \leq n$ that

$$T_{[0,n+1+i]} = \begin{pmatrix} M & 2^n \\ 2^i + M \sum_{1 \leq t \leq i} N_{n-t} 2^{i-t} & 2^n \sum_{1 \leq t \leq i} N_{n-t} 2^{i-t} \end{pmatrix}; \quad (11)$$

so that

$$T_{[0,2n+1]} = \begin{pmatrix} M & 2^n \\ 2^n + MN & 2^n N \end{pmatrix}. \quad (12)$$

Let q be qubit $2n+1$. The only assignment to this qubit which is consistent with a satisfying assignment is then the state given by the first column of $T_{[0,2n+1]}$, which is $M|0\rangle + (2^n + MN)|1\rangle$; the vector given by Eqn. (3) is the normalised version of this vector.

Using the techniques of Laumann *et al.* [21], we may show that the space of satisfying assignments of this instance has dimension 2, spanned by the product solution above, and an entangled solution on all of the qubits. Considering all projectors except for $\Pi'_{0,1}$, there is an invertible (non-unitary) local transformation mapping all projectors $\Pi_{i-1,i}$ to $|\Psi^-\rangle\langle\Psi^-|$, the two-qubit antisymmetric projector. Thus the satisfying states for these projectors are the symmetric subspace on $S = 2n+2$ qubits, which is spanned by any collection of states of the form $|\alpha_i\rangle^{\otimes 2n+2}$, for $S+1 = 2n+3$ distinct states $|\alpha_i\rangle$. Any state in this space which is not a product state, is entangled across the entire chain of qubits. Undoing this change of local co-ordinates, it follows that any state which satisfies the above instance of 2-QSAT which is not a product state, is also entangled across the entire chain of qubits (*i.e.*, it cannot be factorized across any cut). Since we require each factor to be explicitly written in the standard basis, such a solution would then require explicitly writing out the standard basis elements of a vector of dimension 2^{2n+2} ; such solutions would require vectors of dimension 2^{2n+2} to represent. Any algorithm which in polynomial time exhibits one of the tensor factors of the solution, must therefore exhibit factors of the product solution. In particular, it must compute $|\psi_q\rangle$ if this is the required tensor factor. As we have already shown an $O(n)$ reduction from computing the product MN to computing the minimal representation of $|\psi_q\rangle$, it follows that there is an $\Omega(M(n))$ lower bound for such an algorithm in the worst case. ◀

► **Corollary 6.2.** *If there does not exist a $\Theta(n)$ -time algorithm for multiplying two n -bit integers, then there does not exist an $O(m+n)$ -time algorithm to present single-qubit marginals of satisfying solutions to instances of 2-QSAT.*

We would also like to show lower bounds for algorithms such as SOLVE_Q , which do not necessarily compute its output in reduced form, but which does compute an *explicit* output, in the sense of presenting a complete description of a satisfying solution via tensor factors. We may obtain such lower bounds even for algorithms which produce non-normalised outputs, as follows.

► **Lemma 6.3.** *There exist instances of 2-QSAT on n vertices and $m \in O(n)$ clauses, such that an explicit rationalised (but not necessarily normalised) assignment for a satisfying state requires $\Omega(n^2)$ bits.*

Proof. We may simplify the construction of Lemma 6.1 by omitting the qubits $n+1, \dots, 2n+1$ and the projectors which act on them. This yields an instance in which there is a unique product solution (with all other solutions requiring a vector of dimension 2^{n+1} to represent). In this product state, the qubit n is in a state $|\psi_n\rangle \propto |0\rangle + M|1\rangle$. More generally, each qubit $1 \leq i \leq n$ is in a state

$$|\psi_i\rangle \propto |0\rangle + M^{(i)}|1\rangle \quad (13)$$

where $M^{(i)} = \sum_{t=1}^i M_{n-t} 2^{i-t}$. As $M_{n-1}M_{n-2} \cdots M_2M_1 \in \{0,1\}^{n-1}$ may be an arbitrary $n-1$ bit string, and as we require the tensor factors on the qubits i to be presented independently of one another, the integers $M^{(i)}$ cannot be represented any more succinctly in the worst case; at best, by applying arbitrary scalar factors, we may consider representations $|\psi_i\rangle = \frac{1}{\alpha_i}|0\rangle + \frac{M^{(i)}}{\alpha_i}|1\rangle$, in which the representation of the $|1\rangle$ component of $|\psi_i\rangle$ may be reduced if α_i divides $M^{(i)}$, but at the cost of increasing the size of the representation of the $|0\rangle$ component. (More formally, if the pair $(1/\alpha_i, M^{(i)}/\alpha_i)$ has asymptotically smaller Kolmogorov complexity than the pair $(1, M^{(i)})$, we would have a contradiction, since the former allows us to extract $M^{(i)}$ – thus, we would have a shorter description of $M^{(i)}$ than its Kolmogorov complexity allows.) Thus, for any constant $0 < \alpha < 1$, the qubits $\lfloor \alpha n \rfloor < i < n$ all require $\Omega(n)$ bits to represent, yielding a total lower bound of $\Omega(n^2)$. ◀

► **Corollary 6.4.** *Up to $\Omega(\log(n)^{1+o(1)})$ factors, SOLVE_Q is optimal among algorithms which present explicit expressions for satisfying assignments.*

Acknowledgements. NdB thanks Rick, Linda, Colin, and Michelle de Beaudrap for their hospitality during an academic absence, during which the bulk of this work was performed. This project was initiated while NdB was affiliated with the Centrum Wiskunde & Informatica, with support from the European Commission FET-Proactive project Quantum Algorithms (QALGO) 600700, and completed with support from the UK Quantum Technology Hub project NQIT. SG thanks Ronald de Wolf and Centrum Wiskunde & Informatica, where this work was initiated, for their hospitality. SG acknowledges support from a Government of Canada NSERC Banting Postdoctoral Fellowship, the Simons Institute for the Theory of Computing at UC Berkeley, and NSF grant CCF-1526189.

References

- 1 I. Arad, M. Santha, A. Sundaram, and S. Zhang. Linear time algorithm for quantum 2SAT. arXiv:1508.06340, 2015.
- 2 B. Aspvall, M. F. Plass, and R. E. Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Information Processing Letters*, 8(3):121–123, 1979.
- 3 S. Bravyi. Efficient algorithm for a quantum analogue of 2-SAT. Available at arXiv.org e-Print quant-ph/0602108v1, 2006.
- 4 J. Chen, X. Chen, R. Duan, Z. Ji, and B. Zeng. No-go theorem for one-way quantum computing on naturally occurring two-level systems. *Physical Review A*, 83:050301(R), 2011.
- 5 H. Cohen. *A Course in Computational Algebraic Number Theory*. Graduate Texts in Mathematics. Springer, 1993.
- 6 S. Cook. The complexity of theorem proving procedures. In *Proceedings of the 3rd ACM Symposium on Theory of Computing (STOC 1972)*, pages 151–158, 1972.
- 7 M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201, 1960.
- 8 N. de Beaudrap. Difficult instances of the counting problem for 2-quantum-sat are very atypical. In *Proceedings of TQC'14*, pages 118–140, 2014. arXiv:1403.1588.
- 9 N. de Beaudrap, T. J. Osborne, and J. Eisert. Ground states of unfrustrated spin hamiltonians satisfy an area law. *New Journal of Physics*, 12, 2010. arXiv:1009.3051.
- 10 S. Even, A. Itai, and A. Shamir. On the complexity of the time table and multi-commodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976.
- 11 M. Fürer. Faster integer multiplication. In *Proceedings of the 39th ACM Symposium on the Theory of Computing (STOC 2007)*, pages 55–67, 2007.
- 12 Oded Goldreich and David Zuckerman. Another proof that $BPP \subseteq PH$ (and more). In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, volume 6650 of *Lecture Notes in Computer Science*, pages 40–53. Springer Berlin Heidelberg, 2011. doi:10.1007/978-3-642-22670-0_6.
- 13 D. Gosset and D. Nagaj. Quantum 3-SAT is QMA1-complete. In *Proceedings of the 54th IEEE Symposium on Foundations of Computer Science (FOCS 2013)*, pages 756–765, 2013.
- 14 Z. Ji, Z. Wei, , and B. Zeng. Complete characterization of the ground space structure of two-body frustration-free hamiltonians for qubits. *Physical Review A*, 84, 2011.
- 15 S. P. Jordan, H. Kobayashi, D. Nagaj, and H. Nishimura. Achieving perfect completeness in classical-witness quantum Merlin-Arthur proof systems. *Quantum Information & Computation*, 12(5 & 6):461–471, 2012.
- 16 R. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. New York: Plenum, 1972.
- 17 J. Kempe, A. Kitaev, and O. Regev. The complexity of the local Hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006.
- 18 A. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003.
- 19 A. Kitaev, A. Shen, and M. Vyalı. *Classical and Quantum Computation*. American Mathematical Society, 2002.
- 20 M. R. Krom. The decision problem for a class of first-order formulas in which all disjunctions are binary. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 13:15–20, 1967.
- 21 C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi. Phase transitions and random quantum satisfiability. *Quantum Information & Computation*, 10:1–15, 2010.
- 22 L. Levin. Universal search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.

- 23 C. Papadimitriou. On selecting a satisfying truth assignment. In *Proceedings of the 32nd Annual IEEE Symposium on Foundations of Computing (FOCS 1991)*, pages 163–169, 1991.
- 24 W. V. Quine. On cores and prime implicants of truth functions. *The American Mathematical Monthly*, 66(5):755–760, 1959.
- 25 R. E. Tarjan. Depth first search and linear graph algorithms. *SIAM Journal on Computing*, pages 146–160, 1972.
- 26 J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 2003.
- 27 S. Zachos and M. Furer. Probabilistic quantifiers vs. distrustful adversaries. In *Foundations of Software Technology and Theoretical Computer Science, 7th Conference*, pages 443–455, 1987. Volume 287 of *Lecture Notes in Computer Science*.