Hamiltonian Complexity and QMA-Completeness

Sam Pallister

Contents

I. Introduction

One of the cornerstone results in computational complexity theory is the Cook-Levin Theorem [1]: the theorem that establishes that Boolean satisfiability problems are NP-complete. A Boolean satisfiability problem (SAT) asks the following question: given a set of clauses built from literals (variables that take values "TRUE" and "FALSE"), Boolean operators and parentheses, does there exist an assignment of values for the literals such that all clauses evaluate to TRUE? The Cook-Levin theorem demonstrates, in some sense, that satisfiability problems are the stereotypical set of NP problems. Specifically, SAT is a problem in NP and there is at most a polynomial overhead in reduction of other languages in NP to a SAT problem. It is also worth pursuing what restrictions can be made on the clauses such that this property of NP-completeness remains; for example, if each clause contains at most k literals ("k-SAT"). It was shown by Karp [2] that 3-SAT is NP-complete, but 2-SAT was shown by Krom [3] to be in P. This sharp jump in problem complexity between 2-SAT and 3-SAT was generalised in a pleasing result known as Schaefer's dichotomy theorem [4], which demonstrates this step-change in complexity when the constraints imposed are drawn from a particular set.

This narrative plays out in a surprisingly similar way when the complexity classes under consideration are related to the power of quantum computers rather than classical ones. The quantum equivalent of NP is known as QMA (an initialism of "Quantum Merlin-Arthur", a name inspired by the "call and response" type protocol that problems in this class are solved by). In this class, a sender with unbounded computational resources (Merlin) gives a poly-sized proof of an affirmative answer to a decision problem to a receiver (Arthur). Arthur's verification protocol must then lie in BQP; that is, he must be able to verify or reject Merlin's proof with a quantum computer with polynomial resources in polynomial time. One may then ask whether there exists any problem that is QMA-complete, and whether these are natural quantum equivalents of problems that are known to be NP-complete. Interestingly, in the same sense that a computational problem in NP can be encoded as a satisfiability problem for a set of constraints, a computational problem in QMA can be encoded as the problem of finding a ground-state of some local Hamiltonian given a set of constraints. This problem, k-LocalHamiltonian, is the quantum analogue of the Boolean satisfiability problems introduced above (in particular, the analogue of k-MaxCSP, the problem of identifying the maximum number of contraints satisified when clauses have at most k literals in a general constraint satisfaction problem). The k-LocalHamiltonian problem was initially introduced by Kitaev [5]. In direct analogy with k-MaxCSP, which is NP-complete for $k \geq 2$, it can be shown that k-LocalHamiltonian is QMA-complete for $k \geq 2$ [6].

In addition to the constraint of locality, it may be of interest to consider constraints on the Hamiltonian based on the physical quantum system it describes. This is of great importance for adiabatic quantum computers, as implementation of any QMA-complete Hamiltonian is sufficient for universal adiabatic quantum computing [7]. In [7], it was demonstrated that the Ising model can be rendered QMA-complete with appropriate additional couplings, and in [8] it was shown that the Heisenberg and Hubbard models are QMA-complete in the case of physically local couplings. In a manner analogous to Schaefer's dichotomy theorem, these constraints have been recently generalised in [9] to give a dichotomy in computational complexity for physically realisable local Hamiltonians.

The structure of this essay is to expand on this narrative; first, by giving a concise definition of both QMA and the k-LocalHamiltonian problem and then by considering Schaefer's dichotomy theorem and its quantum equivalent. Finally, the relevance of this work to adiabatic quantum computing will be outlined.

II. QMA-completeness

We begin with a formal definition of the complexity class QMA. This class is phrased in a "call and response" manner - Merlin provides a proof of a problem with a promise that it is correct; Arthur must use a quantum circuit and polynomial quantum resources in order to accept correct proofs above a certain threshold probability and to reject incorrect proofs below a certain threshold probability.

Definition (QMA)**.** *A promise problem $L = (L_{yes} \cup L_{no})$ is in the complexity class* QMA *iff, given a string $x \in \Sigma^*$ of size $|x| = n$ and a quantum proof $|y\rangle \in (\mathbb{C}^2)^{\otimes \mathsf{poly}(n)}$ from Merlin, Arthur checks given access to quantum resources: the family of all quantum circuits that run in*

$\mathsf{poly}(n)$ *time,* $\{Q_n\}$, *and* $\mathsf{poly}(n)$ *ancillas. Then, Arthur must:*

- *Verify correct proofs: if* $x \in L_{yes}$, *then there exists a proof* $|y\rangle$ *such that Arthur accepts the pair* $(x, |y\rangle)$ *with probability at least 2/3.*
- *Reject incorrect proofs: if* $x \in L_{no}$, *then there exists no proof* $|y\rangle$ *such that Arthur accepts the pair* $(x, |y\rangle)$ *with probability more than 1/3.*

A problem is QMA-complete if it is both QMA-hard (that is, any problem in QMA can be reduced to it with polynomial overhead) and it is itself in QMA. Any problem that is QMA-complete is likely to require superpolynomial time to evaluate on both classical and quantum hardware.

There are a few notable relationships between QMA and other complexity classes. Firstly, it is clear that NP $\subseteq$ MA $\subseteq$ QMA, as the only difference between the three classes is the computational power of Arthur (deterministic classical computation, probabilistic classical computation and quantum computation, respectively). Secondly, it was shown by Kitaev and Watrous that QMA is in PP [10].

III. K-LocalHamiltonian

The first QMA-complete problem to be discovered took its cue from the Cook-Levin theorem, as it is in some sense the most natural quantum equivalent of a constraint satisfaction problem. The problem concerns 'emphlocal Hamiltonians; that is, ones that can be written

$$H = \sum_{i=1}^{m} H_i, \tag{1}$$

where each H_i acts on at most k qubits for the Hamiltonian to be considered k-local. The k-LocalHamiltonian problem is then defined as follows:

Definition (k-LocalHamiltonian)**.** *Consider a k-local Hamiltonian acting on a set of m qubits; i.e.* $H = \sum_{i=0}^{k} H_i$ *for* $k = \mathsf{poly}(m)$, *each* H_i *acts non-trivially on no more than k qubits, and each* H_i *has bounded norm* $\|H_i\| \leq \mathsf{poly}(m)$. *Then an instance of the problem k-*LocalHamiltonian *is a promise problem; either:*

- *Verify that H has an eigenvalue less than a, or*
- *Verify that H has no eigenvalue less than b,*

under the promise that one of these is true and with condition that $a - b \leq \mathsf{poly}(m)$.

The connection to classical constraint satisfaction can be made explicit; in fact, it is straightforward to show that CSP problems can be encoded in a local Hamiltonian. For example, consider an instance Ψ of, say, k-CSP with clauses c_i that evaluate to a constant Boolean value. Then construct a Hamiltonian with terms that give an energy penalty to constraint-violating assignments, i.e. add terms like

$$H_i = \sum_{\substack{x \in \{0,1\}^k \\ \text{s.t.} c_i(x)=0}} |x\rangle \langle x| \,. \tag{2}$$

Then, an assignment that satisfies all constraints encoded in a state $|y\rangle$ will be a eigenstate of energy with eigenvalue 0 (as it is a computational basis state orthogonal to every $|x\rangle$ in the formula above), whereas any state that encodes a non-satisfying assignment will have non-zero energy. Therefore finding the ground state of the k-local Hamiltonian written here will be a state that satisfies Ψ.

The problem k-LocalHamiltonian was first shown to be QMA-complete by Kitaev for $k \geq 5$ [5]; this was later refined by Kempe, Kitaev and Regev to show QMA-completeness for $k \geq 2$. A physicist may then ask: if the Hamiltonian is restricted further by now introducing a set of constraints given instead by physics, does the problem of finding ground states of that Hamiltonian remain QMA-complete? There are case-by-case answers to this question for a handful of physical Hamiltonians. For example, the problem is still QMA-complete for qubits arranged on a line with only nearest neighbour interactions [11] and is still QMA-complete if the qubits are arranged on a square lattice [12]. One can also restrict the physicality of the qubits themselves; the problem is still QMA-complete if the qubits are bosons [13] or fermions [14].

IV. Schaefer's dichotomy theorem

The generalisation of the handful of results on QMA-completeness for physically restricted local Hamiltonians is introduced in the next section as the "quantum dichotomy theorem". However, the quantum dichotomy theorem is heavily inspired by a classical counterpart concerned with the characterisation of the hardness of constraint satisfaction problems known as "Schaefer's dichotomy theorem". Given the significance of this result, it is presented here first.

The form of the dichotomy theorem presented by Schaefer [4] is as follows:

Theorem 1 (Schaefer's Dichotomy Theorem)**.** *Given a set* $\mathcal{S}$ *of constraints over the Boolean domain, define an instance of the class* $\mathsf{CSP}(\mathcal{S})$ *as a conjunction of constraints drawn from* $\mathcal{S}$ *on a set of propositional variables. Then, provided that* $\mathcal{S}$ *satisfies any of the following conditions, the problem of checking satisfiablity of the instance is in* P*:*

- *All constraints in* $\mathcal{S}$ *which aren't constantly false are true when all its arguments are true.*
- *All constraints in* $\mathcal{S}$ *which aren't constantly true are false when all its arguments are false.*
- *All constraints in* $\mathcal{S}$ *are equivalent to a conjunction of binary clauses.*
- *All constraints in* $\mathcal{S}$ *are equivalent to a conjunction of Horn clauses.**
- *All constraints in* $\mathcal{S}$ *are equivalent to a conjunction of dual-Horn clauses.*†
- *All constraints in* $\mathcal{S}$ *are equivalent to a conjunction of affine clauses.*‡

If none of these conditions hold for $\mathcal{S}$, *then the class* $\mathsf{CSP}(\mathcal{S})$ *is* NP*-complete.*

The dichotomy theorem is particularly surprising given the ubiquity of constraint satisfaction problems and a separate result by Ladner:

*A Horn clause is a clause with at most one positive literal.

†A dual Horn clause is a clause with at most one negated literal.

‡An affine clause is defined by Schaefer as a clause composed exclusively of literals and $\oplus$, the XOR operation, such that it evaluates to a constant Boolean value.

Theorem 2 (Ladner's Theorem)**.** *Provided that* $\mathsf{P} \neq \mathsf{NP}$*, then there exist problems in* NP *that are neither in* P*, nor are* NP*-complete.*

At first glance, this appears contradictory - Schaefer's theorem appears to state that there are no constraint satisfaction problems that lie in the NP-intermediate space defined by Ladner. The resolution of the contradiction is that the instances generated by Schaefer, i.e. instances generated by conjunction and substitution of propositional variables, form a somewhat special set of problems (and cannot be used to encode a general SAT problem). However, this caveat should not be taken to mean that Schaefer's theorem is only concerned with a peculiar class of problems. It is a very strong result; for example, it implies the NP-completeness of 3-SAT and that $\mathsf{HORN}-\mathsf{SAT}$ (satisfiability with Horn clauses) is in P, for example. That this dichotomy exists at all for constraint satisfaction problems is a remarkable result.

V. The quantum dichotomy theorem

Inspired by the result of Schaefer and the connection between satisfaction problems and k-LocalHamiltonian, Cubitt and Montanaro [9] recently developed a *quantum* dichotomy theorem that establishes step-changes in the complexity of physically constrained, local Hamiltonian problems. The form of the quantum dichotomy theorem as presented in [9] is the following:

Theorem 3 (The quantum dichotomy theorem)**.** *Define the set* $\mathcal{S}$ *to be an arbitrary set of Hermitian matrices acting on no more than 2 qubits. Given this set* $\mathcal{S}$*, define the problem* $\mathcal{S}$-Hamiltonian *in the following way:* $\mathcal{S}$-Hamiltonian *is a special case of the problem k-*LocalHamiltonian *where the Hamiltonian is now given as a summation of terms* $\alpha_i H_i \in \mathcal{S}$*, where* $\alpha_i \in \mathbf{R}\ \forall i$*.*

Then the following dichotomy holds:

- *If every matrix in* $\mathcal{S}$ *is 1-local, then* $\mathcal{S}$-Hamiltonian *is in* P.
- *Otherwise, if* $\mathcal{S}$ *is locally diagonalised by a matrix* $U \in SU(2)$*, then* $\mathcal{S}$-Hamiltonian *is* NP*-complete.*
- *Otherwise, if* $\exists$ *a matrix* $U \in SU(2)$ *such that for all 2-qubit gates* $H_i \in \mathcal{S}$*,* $U^{\otimes 2} H_i U^{\otimes 2\dagger} = \alpha_i Z^{\otimes 2} + A_i \mathbf{1} + \mathbf{1} B_i$ *(where* $\alpha_i \in \mathbf{R}$ *and* A_i, B_i *are arbitrary single-qubit gates), then the complexity of* $\mathcal{S}$-Hamiltonian *is* TIM*-complete, where* TIM *is the complexity class of the Ising model with transverse external magnetic field.*
- *If none of the above hold, then* $\mathcal{S}$-Hamiltonian *is* QMA*-complete.*

This dichotomy theorem subsumes many of the individual cases stated previously. In particular, the QMA-completeness of the Heisenberg model as shown by Schuch and Verstraete [8] is a specific case of this result.[§] It also subsumes the result by Biamonte and Love [7] that the XY Hamiltonian is QMA-complete. It also settles the complexity of a large list of Hamiltonians with couplings previously untouched in the literature.

[§]In fact, the Hamiltonian considered in [8] was the Heisenberg Hamiltonian with the allowance of arbitrary 1-local terms. The dichotomy theorem demonstrates the QMA-completeness of the Heisenberg model both with and without this additional allowance.

Much like Schaefer's construction of $\mathsf{CSP}(\mathcal{S})$ had limited scope compared to general satisfaction problems, the quantum dichotomy theorem does not cover all cases of Hamiltonian problems with . For example, Bravyi's work concerning the complexity of stoquastic Hamiltonians [15] does not fit exactly into the $\mathcal{S}-\mathsf{Hamiltonian}$ scheme, and the results concerning the QMA-completeness of bosonic and fermionic systems [13], [12] remain independent.

VI. Adiabatic quantum computing and universality

The final section of this paper is to remark on the relevance of this discussion for adiabatic quantum computing.

Adiabatic quantum computation [16] is the "slow" evolution of a Hamiltonian from some easily preparable form, H_0, to a form H_s that encodes the answer to some computationally difficult problem, whilst maintaining the state in the lowest energy configuration. For example, consider the time-dependent Hamiltonian

$$H(t) = (1 - \frac{t}{T})H_0 + \frac{t}{T}H_s, \quad (3)$$

where H_0 is some easily manufactured Hamiltonian and H_s is the example Hamiltonian given in equation 2 for solving SAT instances. As the Hamiltonian evolves from $t = 0$ to $t = T$, provided that the evolution is slow enough then the system remains equilibrated in its ground state (a result known as the adiabatic theorem). The speed at which the Hamiltonian can be evolved depends on the minimum energy gap between the ground state and the first excited state (here written as g_{min})). If the instantaneous eigenvalues and eigenstates of $H(t)$ are given by $E_l(t)$ and $|l, t\rangle$, respectively, then the condition for adiabaticity is given by

$$T \gg \frac{\max_t |\langle 1, t/T| \frac{dH}{dt} |0, t/T\rangle|}{\min_t(E_1(t/T) - E_0(t/T))} := \frac{\mathcal{E}}{g_{min}^2}, \quad (4)$$

a proof of which is in [16]. An optimistic reader might hope that this scheme is sufficient for finding answers to NP-hard problems, where the evolution moves slowly enough to retain adiabaticity but fast enough that solutions are found efficiently. However, in [17] it was shown that there is a polynomial equivalence between adiabatic quantum computing and the standard, gate-based model of quantum computing.[¶] The proof of this statement is predicated on two things: firstly, that the spectral gap g_{min} only increases inverse polynomially with the size of the problem; and secondly, that the Hamiltonian used to encode solutions is local.

This is, ultimately, where the interest from implementers of adiabatic quantum computing in QMA-complete Hamiltonians stems. Any adiabatic quantum computing platform that can encode a QMA-complete Hamiltonian problem can encode any other Hamiltonian problem in QMA, including the problem k-LocalHamiltonian. By the result in [17], this is then enough to be universal for quantum computation. This is also part of the motivation for considering physically constrained Hamiltonians like those in [11], [12], [13],

[¶]The term "polynomial equivalence" is used to mean that adiabatic and gate-based quantum computing can simulate each other with only a polynomial overhead.

[14], and ultimately those that are the motivation for the quantum dichotomy theorem.

The recipe the quantum dichotomy theorem gives from the perspective of adiabatic quantum computing is as follows: from any physically implemented Hamiltonian, reconstruct the corresponding set of Hermitian matrices $\mathcal{S}$. Then provided that $\mathcal{S}$-$\mathsf{LocalHamiltonian}$ is QMA-complete, implementing this Hamiltonian as H_s is enough to be universal for quantum computation.

Interestingly, it is unclear from the literature what the computational power is of an adiabatic quantum computer that can only encode an "easier" Hamiltonian in H_s. Specifically, the interest is in the class TIM, as Ising models with transverse magnetic fields are readily implementable in practice and form the basis of D-Wave's prototype quantum computing devices [18].

An indication of the answer to this question is provided by Hastings in [19]. Here, the systems under consideration are adiabatic quantum computers with H_s given by a local Hamiltonian without a "sign problem" (also known as "stoquastic"); i.e. Hamiltonians whose off-diagonal elements are all negative. In can be shown that $\mathsf{TIM} \subseteq \mathsf{StoqMA}$, the set of local Hamiltonian problems where the Hamiltonian is assumed to be stoquastic. It does not appear to be an unreasonable conjecture that restricting H_s to be stoquastic is enough to hamstring the computational power of an adiabatic quantum computer; say, to conjecture that its computational power is restricted to solving problems in BPP. While [19] gives no definite answer to this question, it does give a strong indication that this conjecture is untrue. Specifically, the claim is that the most straightforward way to show that the power of the adiabatic quantum computer is reduced is to take a classical probabilistic algorithm for evaluating ground-state energies of some Hamiltonian (here, the algorithm considered is a path integral quantum Monte Carlo method), and to show that it is just as efficient as the adiabatic algorithm when the Hamiltonian is stoquastic. Hastings explicitly shows that this is not the case, by constructing counterexamples where the quantum Monte Carlo method fails to equilibrate efficiently.

On the other hand, a recent result by Bravyi [15] indicates some restriction on the computational power of an adiabatic quantum computer given a stoquastic Hamiltonian. This work approaches the problem from a similar tack as that by Hastings, with the modification that the classical probabilistic algorithm under consideration is now a diffusion Monte Carlo algorithm rather than a path integral approach. Whilst the setting is very similar, Bravyi also makes another large assumption not made by Hastings. The added assumption is the inclusion of a "guiding state"; that is, that there is a readily preparable state $|\phi_g\rangle$ with non-negligible overlap with the ground-state of the stoquastic Hamiltonian; if this ground-state is $|\psi_0\rangle$ then

$$\langle x|\phi_g\rangle \geq \frac{\langle x|\psi_0\rangle}{\mathsf{poly}(n)} \text{for all n-bit binary strings x.} \tag{5}$$

Whilst this may seem like a strong restriction to place on the problem, an assumption of this flavour is made at the beginning of the quantum phase estimation algorithm (although there, the assumption is only that there exists a guiding state with non-negligible overlap with the ground-state, and not the stronger, pointwise correlation that is given above). Whether an arbitrary stoquastic Hamiltonian admits a guiding state is not obvious, although the claim by Bravyi is that it appears unlikely. The main result of this work is that, with the inclusion of the guiding state, the complexity class of the k-$\mathsf{LocalHamiltonian}$ problems that are both *guided* and *stoquastic* is in MA and is MA-complete for $k \geq 6$, as opposed to QMA-complete without the added restrictions of guiding and stoquasticity. A corollary of the mathematics used to derive this result is that a *ferromagnetic* TIM Hamiltonian, i.e. a TIM Hamiltonian where the spin-coupling coefficients are all non-negative, admits an approximation scheme with polynomial runtime on a probabilistic computer.

These results clearly give no definitive answers on the complexity of TIM Hamiltonians, and therefore give no definitive answer on the universality of current adiabatic quantum computing implementations. The sentiment for builders of adiabatic quantum computers is still the same - encode a provably QMA-complete Hamiltonian in H_s; for example, the 2-local Hamiltonians discussed by Biamonte and Love [7] or the examples discussed in the context of the quantum dichotomy theorem, in order to ensure universality.

VII. Open questions

There are at least two immediate avenues for further work. As mentioned previously, the construction of the $\mathcal{S}-\mathsf{Hamiltonian}$ problem is not of infinite scope and fails to capture a handful of physically relevant Hamiltonians. Since the quantum dichotomy theorem takes its inspiration from Schaefer, it might be worth pursuing generalisations of Schaefer's result that already exist in the literature and then teasing out their implications for the quantum case. For example, in [20] Schaefer's dichotomy theorem is modified such that the constraints are built from propositional logic on graphs rather than Boolean logic.

Secondly, the settling of the debate on the complexity of the class TIM is of relevance both from the point of view of the dichotomy theorem and, as discussed above, for implementations of adiabatic quantum computing. It seems unlikely that TIM will be either "easy enough" or "hard enough" for the dichotomy to collapse, but it may be the case that $\mathsf{TIM} \supseteq \mathsf{MA}$, for example.

References

[1] S. A. Cook, "The complexity of theorem-proving procedures," in *Proceedings of the third annual ACM symposium on Theory of computing - STOC '71*, (New York, New York, USA), pp. 151–158, ACM Press, May 1971.

[2] R. Karp, *Reducibility among Combinatorial Problems*. Boston, MA: Springer US, 1972.

[3] 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*, vol. 13, no. 1-2, pp. 15–20, 1967.

[4] T. J. Schaefer, "The complexity of satisfiability problems," in *Proceedings of the tenth annual ACM symposium on Theory of computing - STOC '78*, (New York, New York, USA), pp. 216–226, ACM Press, May 1978.

[5] A. Y. Kitaev, A. Shen, and M. N. Vyalyi, *Classical and Quantum Computation*. American Mathematical Society, 2002.

[6] J. Kempe, A. Kitaev, and O. Regev, "The Complexity of the Local Hamiltonian Problem," *SIAM Journal on Computing*, vol. 35, pp. 1070–1097, Jan. 2006.

[7] J. D. Biamonte and P. J. Love, "Realizable Hamiltonians for universal adiabatic quantum computers," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 78, p. 012352, July 2008.

[8] N. Schuch and F. Verstraete, "Computational complexity of interacting electrons and fundamental limitations of density functional theory," *Nature Physics*, vol. 5, pp. 732–735, Aug. 2009.

[9] T. Cubitt and A. Montanaro, "Complexity classification of local Hamiltonian problems," Nov. 2013.

[10] C. Marriott and J. Watrous, "Quantum Arthur-Merlin games," in *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.*, pp. 275–285, IEEE, 2004.

[11] D. Nagaj, *Local Hamiltonians in Quantum Computation*. PhD thesis, Aug. 2008.

[12] R. Oliveira and B. M. Terhal, "The complexity of quantum spin systems on a two-dimensional square lattice," *Quantum Information & Computation*, vol. 8, pp. 900–924, Nov. 2008.

[13] T.-C. Wei, M. Mosca, and A. Nayak, "Interacting Boson Problems Can Be QMA Hard," *Physical Review Letters*, vol. 104, p. 040501, Jan. 2010.

[14] J. D. Whitfield, P. J. Love, and A. Aspuru-Guzik, "Computational complexity in electronic structure.," *Physical chemistry chemical physics : PCCP*, vol. 15, pp. 397–411, Jan. 2013.

[15] S. Bravyi, "Monte Carlo simulation of stoquastic Hamiltonians," Feb. 2014.

[16] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, "Quantum Computation by Adiabatic Evolution," Jan. 2000.

[17] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, "Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation," in *45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 42–51, IEEE, 2004.

[18] Z. Bian, F. Chudak, R. Israel, B. Lackey, W. G. Macready, and A. Roy, "Discrete optimization using quantum annealing on sparse Ising models," *Frontiers in Physics*, vol. 2, Sept. 2014.

[19] M. B. Hastings, "Obstructions to classically simulating the quantum adiabatic algorithm," *Quantum Information & Computation*, vol. 13, pp. 1038–1076, Nov. 2013.

[20] M. Bodirsky and M. Pinsker, "Schaefer's theorem for graphs," Nov. 2010.