

Guidable Local Hamiltonian Problems with Implications to Heuristic Ansatz State Preparation and the Quantum PCP Conjecture

Jordi Weggemans ✉ 🏠 

CWI & QuSoft, Amsterdam, The Netherlands

Fermioniq, Amsterdam, The Netherlands

Marten Folkertsma ✉

CWI & QuSoft, Amsterdam, The Netherlands

Chris Cade ✉

Fermioniq, Amsterdam, The Netherlands

QuSoft & University of Amsterdam (UvA), The Netherlands

Abstract

We study “Merlinized” versions of the recently defined Guided Local Hamiltonian problem, which we call “*Guidable* Local Hamiltonian” problems. Unlike their guided counterparts, these problems do not have a guiding state provided as a part of the input, but merely come with the promise that one *exists*. We consider in particular two classes of guiding states: those that can be prepared efficiently by a quantum circuit; and those belonging to a class of quantum states we call *classically evaluable*, for which it is possible to efficiently compute expectation values of local observables classically. We show that guidable local Hamiltonian problems for both classes of guiding states are QCMA-complete in the inverse-polynomial precision setting, but lie within NP (or NqP) in the constant precision regime when the guiding state is classically evaluable.

Our completeness results show that, from a complexity-theoretic perspective, classical Ansätze selected by classical heuristics are just as powerful as quantum Ansätze prepared by quantum heuristics, as long as one has access to quantum phase estimation. In relation to the quantum PCP conjecture, we (i) define a complexity class capturing quantum-classical probabilistically checkable proof systems and show that it is contained in $BQP^{NP[1]}$ for constant proof queries; (ii) give a no-go result on “dequantizing” the known quantum reduction which maps a QPCP-verification circuit to a local Hamiltonian with constant promise gap; (iii) give several no-go results for the existence of quantum gap amplification procedures that preserve certain ground state properties; and (iv) propose two conjectures that can be viewed as stronger versions of the NLTS theorem. Finally, we show that many of our results can be directly modified to obtain similar results for the class MA.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Quantum complexity theory, local Hamiltonian problem, quantum state ansatzes, QCMA, quantum PCP conjecture

Digital Object Identifier 10.4230/LIPIcs.TQC.2024.10

Related Version *Full Version*: <https://arxiv.org/abs/2302.11578> [48]

Funding CC was supported by Fermioniq B.V., and thanks QuSoft and CWI for their accommodation whilst this work was completed. MF and JW were supported by the Dutch Ministry of Economic Affairs and Climate Policy (EZK), as part of the Quantum Delta NL programme.

Acknowledgements The authors thank Jonas Helsen and Harry Buhrman for useful discussions and Sevag Gharibian and François Le Gall for their comments on an earlier version of the manuscript. We also thank Ronald de Wolf for providing feedback on the introduction. We thank the anonymous reviewers for their helpful comments, and in particular an anonymous STOC reviewer who provided an annotated version of an earlier version of this manuscript with very detailed and helpful comments.



© Jordi Weggemans, Marten Folkertsma, and Chris Cade;
licensed under Creative Commons License CC-BY 4.0

19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024).

Editors: Frédéric Magniez and Alex Bredariol Grilo; Article No. 10; pp. 10:1–10:24

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 **Introduction**

Quantum chemistry and quantum many-body physics are generally regarded as two of the most promising application areas of quantum computing [1, 12]. Whilst perhaps the original vision of the early pioneers of quantum computing was to simulate the *time-dynamics* of quantum systems [13, 26], for many applications one is interested in *stationary* properties. One particularly noteworthy quantity is the *ground state energy* (which corresponds to the smallest eigenvalue) of a local Hamiltonian describing a quantum mechanical system of interest, say a small molecule or segment of material. The precision to which one can estimate the ground state energy plays a crucial role in practice: for instance, in chemistry the relative energies of molecular configurations enter into the exponent of the term computing reaction rates, making the latter exceptionally sensitive to small (non-systematic) errors in energy calculations. The problem of estimating the smallest eigenvalue of a local Hamiltonian up to some additive error relative to the operator norm (the decision variant of which is known as the *local Hamiltonian problem*) is well-known to be QMA-hard when the required accuracy scales inversely with a polynomial. Therefore, it is generally believed that, without any additional help or structure, quantum computers are not able to accurately estimate the smallest eigenvalues of general local Hamiltonians, and there is some evidence that this hardness carries over to those Hamiltonians relevant to chemistry and materials science [40]. A natural question to ask is then the following: how much “extra help” needs to be provided in order to accurately estimate ground state energies using a quantum computer?

In the quantum chemistry community, it is often suggested that this extra help could come from a classical heuristic that first finds some form of *guiding state*: a classical description of a quantum state that can be used as an input to a quantum algorithm to compute the ground state energy accurately [38]. Concretely, this comes down to the following two-step procedure [17]:

- Step 1 (Guiding state preparation): A classical heuristic algorithm is applied to obtain a *guiding state* $|\psi\rangle$, which is hoped to have “good”¹ fidelity with the ground space.
- Step 2: (Ground state energy approximation): The guiding state $|\psi\rangle$ is used as input to Quantum Phase Estimation (QPE) to efficiently and accurately compute the corresponding ground state energy.

Step 2 of the above procedure can be formalised by the *Guided k -local Hamiltonian problem* (k -GLH), which was introduced in [28] and shown to be BQP-complete under certain parameter regimes that were subsequently improved and tightened in [17]. The problem k -GLH is stated informally as follows: given a k -local Hamiltonian H , an appropriate classical “representation” of a guiding state $|u\rangle$ promised to have ζ -fidelity with the ground space of H , and real thresholds $b > a$, decide if the ground state energy of H lies above or below the interval $[a, b]$.

In a series of works [17, 18, 28], it was shown that 2-GLH is BQP-complete for *inverse polynomial* precision and fidelity, i.e. $b - a \geq 1/\text{poly}(n)$ and $\zeta = 1 - 1/\text{poly}(n)$ respectively. In contrast, when $b - a \in \Theta(1)$ and $\zeta = \Omega(1)$, k -GLH can be efficiently solved *classically* by using a dequantized version of the quantum singular value transformation [28].

The GLH problem forms the starting point of this work. We study “*Merlinized*” versions of GLH – in which guiding states are no longer given as part of the input but instead are only promised to exist – and use these as a way to gain some insight into important theoretical questions in quantum chemistry and complexity theory. In the subsequent paragraphs, we introduce some of the motivating questions guiding the study of the complexity of these so-called “guidable” local Hamiltonian problems.

¹ “Good” here means at least inverse polynomial in the number of qubits the Hamiltonian acts on.

Ansätze for state preparation. Step 1 of the aforementioned two-step procedure generally requires one to have access to classical heuristics capable of finding guiding states whose energies can be estimated classically (as a metric to test whether candidate states are expected to be close to the actual ground state or not). Furthermore, these “trial states” should also be preparable as quantum states on a quantum computer, so that they can be used as input to phase estimation in Step 2. In [28], inspired by a line of works that focused on the dequantization of quantum machine learning algorithms [20, 33, 44], a particular notion of “sampling-access” to the guiding state $|u\rangle$ is assumed. Specifically, it is assumed that one can both query the amplitude of arbitrary basis states, and additionally that one can sample basis states according to their l_2 norm with respect to the overall state $|u\rangle$. This can be a somewhat powerful model [22], and it is closely related to the assumption of QRAM access to classical data, and thus in the context of quantum machine learning (where such access is commonly assumed), it makes sense to compare quantum machine learning algorithms to classical algorithms with sampling access to rule out quantum speed-ups that come merely from having access to quantum states that are constructed from exponential-size classical data. However, for quantum chemistry and quantum many-body applications, this type of access to quantum states seems to be somewhat artificial. From a theoretical perspective, one might wonder to what extent this sampling access model “hides” some complexity, allowing classical algorithms to perform well on the problem when they otherwise would not.

Moreover, one may ask whether the fact that the ground state preparation in Step 1 considers only *classical* heuristics might be too restrictive. *Quantum* heuristics for state preparation, such as variational quantum eigensolvers [46] and adiabatic state preparation techniques [8], have received considerable attention as possible quantum approaches within the NISQ era, and one can argue that even in the fault-tolerant setting, such heuristics will likely still be viable approaches to state preparation, in particular when used in conjunction with Quantum Phase Estimation.

The quantum PCP conjecture. Arguably the most fundamental result in classical complexity theory is the Cook-Levin Theorem [21, 36], which states that constraint satisfaction problems (CSPs) are NP-complete. The PCP theorem [10, 11], which originated from a long line of research on the complexity of interactive proof systems, can be viewed as a “strengthening” of the Cook-Levin theorem. In its proof-checking form, it states that all decision problems in NP can be decided, with a constant probability of error, by only checking a constant number of bits of a polynomially long proof string y (selected randomly from the entries of y). There are also alternative equivalent formulations of the PCP theorem. One is in terms of *hardness of approximation*: it states that it remains NP-hard to decide whether an instance of CSP is either completely satisfiable, or whether no more than a constant fraction of its constraints can be satisfied.² Naturally, quantum complexity theorists have proposed proof-checking and hardness of approximation versions of PCP in the quantum setting. Given the close relationship between QMA and the local Hamiltonian problem, the most natural formulation is in terms of hardness of approximation: in this context, the *quantum* PCP conjecture roughly states that energy estimation of a (normalized) local Hamiltonian up to *constant* precision, relative to the operator norm of the Hamiltonian, remains QMA-hard. This conjecture is arguably one of the most important open problems in quantum complexity theory and has remained unsolved for nearly two decades.

² The transformation of a CSP to another one which is hard to approximate is generally referred to as *gap amplification*, and is realised in Dinur’s proof of the PCP theorem [24].

One way to shed light on the validity of the quantum PCP conjecture can be to study PCP-type conjectures for other “Merlinized” complexity classes. Up until this point, PCP-type conjectures have not been considered for other classes besides NP and QMA.³ However, there is the beautiful result of [7], which studies the possibility of a gap amplification procedure for the class MA by considering a particular type of Hamiltonian: uniform stoquastic local Hamiltonians. The authors show that deciding whether the energy of such a Hamiltonian is exactly zero or inverse polynomially bounded away from zero is MA-hard, but that the problem is in NP when this interval is increased to be some constant. Consequently, this implies that there can exist a gap-amplification procedure for uniform stoquastic Local Hamiltonians (in analogy to the gap amplification procedure for constraint satisfaction problems in the original PCP theorem) if and only if $MA = NP$ – i.e. if MA can be derandomized. Since $MA \subseteq QMA$, this result also shows that if a gap amplification procedure for the general local Hamiltonian problem would exist that “preserves stoquasticity”, then it could also be used to derandomize MA.

1.1 Summary of main results

1.1.1 Completeness results for guidable local Hamiltonian problems

Inspired by classical heuristics that work with Ansätze to approximate the ground states of local Hamiltonians, we define a general class of states that we call *classically evaluatable and quantumly preparable*.

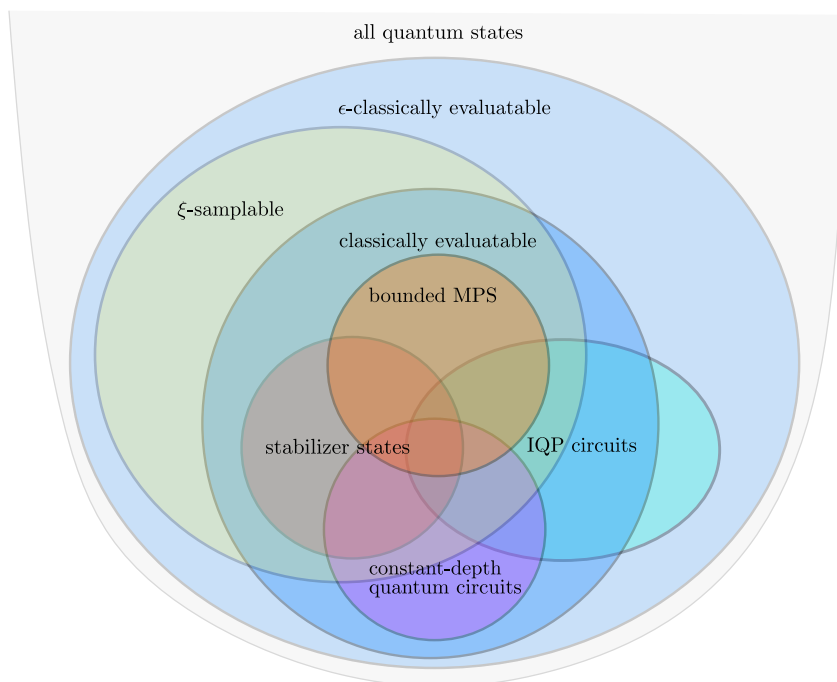
► **Definition 1** (Classically evaluatable and quantumly preparable states). *We say that an n -qubit state $|u\rangle$ is ϵ -classically evaluatable if*

- (i) *it has an efficient classical description which requires at most a polynomial number of bits to write down and*
 - (ii) *one can, given such a description classically efficiently compute expectation values of $\mathcal{O}(\log n)$ -local observables of $|u\rangle$ up to precision ϵ and with probability $\geq 1 - 1/\text{poly}(n)$.*
- In addition, we say that the state is also quantumly preparable if (iii) there exists a quantum circuit that prepares $|u\rangle$ using only a polynomial number of two-qubit gates. Furthermore, if $\epsilon = 0$ the algorithm in (ii) is deterministic instead of probabilistic and we simply say that $|u\rangle$ is classically evaluatable.*

This definition of states is very closely related to the definition of query and sampling access to quantum states given by Gharibian and Le Gall [28], which slightly generalizes the original definition as first proposed by Tang used to dequantize quantum algorithms for recommendation systems [44]. There are three main motivations for introducing this new class of states:

1. It seems rather difficult to find Ansätze that are used in practice for ground state energy estimation that satisfy all conditions of query and sampling access. As one of the main motivations of this work is to investigate the power of quantum versus classical state preparation when one has access to Quantum Phase Estimation, we want to define a class of states that can both be prepared efficiently on a quantum computer and which contains a large class of Ansätze commonly used in practice.

³ Barring a result by Drucker which proves a PCP theorem for the class AM [25]; though there is no direct relationship between QMA and AM and hence it is not clear whether this gives any intuition about the likely validity of the quantum PCP conjecture.



■ **Figure 1** Visualization of the (conjectured) relations between classes of quantum states considered in this work, given a Hilbert space of a fixed dimension. For MPS, we only consider states with polynomially-bounded bond and local dimension. We take $\xi \leq \epsilon/8 \leq 1/3$, such that by Theorem 2 we have that (i) all ξ -samplable states are also ϵ -classically evaluatable and (ii) constant-depth and IQP circuits are not ξ -samplable.

2. Analogous to Dinur's construction, one would expect that determining if a local Hamiltonian has ground state energy (exponentially close to) zero or some constant away from zero is QMA-hard if the quantum PCP conjecture is true. However, there are arguments from physics⁴ as to why one might expect this problem to be in NP [41]. To study the question of containment in NP it is necessary to be able to work with states within a deterministic setting, and therefore it does not make sense to rely on a form of sampling access which inherently relies on a probabilistic model of computation.
3. To add to the previous point, being able to study containment in NP comes with the additional advantage of being able to make statements about whether the problem admits a PCP by the classical PCP theorem. No such theorem is currently known for MA.

To strengthen the first point we find four concrete examples of Ansätze that satisfy all three conditions: matrix product states (MPS), stabilizer states, constant-depth quantum circuits and IQP circuits [15]. Explicit definitions of these classes of states as well as proofs of containment can be found in [48]. The first two examples are in fact also perfectly samplable. However, constant-depth quantum circuits are not even approximately samplable (under the conjecture that $\text{BQP} \not\subseteq \text{AM}$ [45]). We can formalize this in the following theorem which relates ξ -samplable states to ξ -classically evaluatable states

⁴ In this setting the LH problem becomes equivalent to determining whether the free energy of the system becomes negative at a finite temperature. One expects then that at such temperatures, the system loses its quantum characteristics on the large scale, making the effects of long-range entanglement become negligible. Hence, this means that the ground state of such a system should have some classical description, which places the problem in NP [9].

► **Theorem 2.** *For any $\xi > 0$, any ξ -samplable state is also $\mathcal{O}(\xi)$ -classically evaluable. On the other hand, there exist states that are perfectly classically evaluable but not ξ' -samplable for all $0 < \xi' < 1/3$, unless $\text{BQP} \subseteq \text{AM}$.*

The proof of this theorem can be found in the full version of this paper [48]. This theorem gives rise to a (conjectured) hierarchical structure of states as depicted in Figure 1. For the remainder of our work, we will focus on (0-)classically evaluable states, which by Definition 1 means that there exists a deterministic classical algorithm for computing expectation values. A notable advantage of this approach is, as opposed to 0-samplable states, that this allows us to give NP containment results.

Our main focus is on a new family of Hamiltonian problems, in which we are promised that the ground state is close (with respect to fidelity) to some state from a particular class of states, called *guiding states*. We make a distinction between different types of promises one can make with respect to the existence of guiding states: we either assume that the guiding states are of the form of Definition 1 (with or without the promise that the states are also quantumly preparable), or that there exists an efficient quantum circuit that prepares the guiding state.

► **Definition 3** (Guidable Local Hamiltonian problems). *Guidable Local Hamiltonian Problems are problems defined by having the following input, promise, output and some extra promise to be precisely defined below for each of the problems separately:*

Input: *A k -local Hamiltonian H with $\|H\| \leq 1$ acting on n qubits, threshold parameters $a, b \in \mathbb{R}$ such that $b - a \geq \delta > 0$ and a fidelity parameter $\zeta \in (0, 1]$.*

Promise: *We have that either $\lambda_0(H) \leq a$ or $\lambda_0(H) \geq b$ holds, where $\lambda_0(H)$ denotes the ground state energy of H .*

Extra promises: *Let Π_{gs} be the projection on the subspace spanned by the ground states of H . Then for each problem, we have that either one of the following promises holds:*

1. *There exists a classically evaluable state $u \in \mathbb{C}^{2^n}$ for which $\|\Pi_{gs}u\|^2 \geq \zeta$. Then the problem is called the **Classically Guidable Local Hamiltonian Problem**, shortened as $\text{CGaLH}(k, \delta, \zeta)$. If $|u\rangle$ is also quantumly preparable, we call the problem the **Classically Guidable and Quantumly Preparable Local Hamiltonian Problem**, shortened as $\text{CGaLH}^*(k, \delta, \zeta)$.*
2. **Quantumly Guidable k -LH** ($\text{QGaLH}(k, \delta, \zeta)$): *There exists a quantum circuit of polynomially many two-qubit gates that produces the state $|\phi\rangle$ for which $\|\Pi_{gs}|\phi\rangle\|^2 \geq \zeta$.*

Output: *– If $\lambda_0(H) \leq a$, output YES.*

- *If $\lambda_0(H) \geq b$, output NO.*

We note that a guidable local Hamiltonian problem variant for a different class of guiding states was already introduced in Section 5 of [28] without giving any hardness results. Using techniques from Hamiltonian complexity we obtain the following completeness results.⁵

► **Theorem 4** (Complexity of guidable local Hamiltonian problems). *For $k = 2$ and $\delta = 1/\text{poly}(n)$, we have that both $\text{CGaLH}^*(k, \delta, \zeta)$ and $\text{QGaLH}(k, \delta, \zeta)$ are QCMA-complete when $\zeta \in (1/\text{poly}(n), 1 - 1/\text{poly}(n))$.*

A basic version of the hardness proof can be found in Section 3.1, with the remainder written down in the full version [48]. A direct corollary of the above theorem is the following.

⁵ In fact $\text{QGaLH}(k, \delta, \zeta)$ remains QCMA-hard all the way up to $\zeta = 1$.

► **Corollary 5** (Classical versus quantum state preparation). *When one has access to a quantum computer (and in particular quantum phase estimation), then having the ability to prepare any quantum state preparable by a polynomially-sized quantum circuit is no more powerful than the ability to prepare states from the family of classically evaluable and quantumly preparable states, when the task is to decide the local Hamiltonian problem with precision $1/\text{poly}(n)$.*

It should be noted that our result does *not* imply that all Hamiltonians which have efficiently quantumly preparable guiding states also necessarily have guiding states that are classically evaluable. All this result says is that for any instance of the guidable local Hamiltonian problem with the promise that there exist guiding states that can be efficiently prepared by a quantum computer, there exists an (efficient) *mapping* to another instance of the guidable local Hamiltonian problem with the promise that there exist guiding states that are classically evaluable and quantumly preparable. Whilst this reduction is efficient in the complexity-theoretic sense, it might not be for practical purposes, as it would likely remove all the physical structure present in the original Hamiltonian. Hence, the main implication of our result is not that these kinds of reductions are of practical merit, but that at least from a complexity-theoretic point of view the aforementioned classical-quantum hybrid approach of guiding state selection through *classical* heuristics combined with *quantum* energy estimation is at least as powerful as using quantum heuristics for state preparation instead.

We complement our quantum hardness results with classical containment results (of the classically guidable local Hamiltonian problem), obtained through a deterministic dequantized version of Lin and Tong’s ground state energy estimation algorithm [37]. Here CGaLH is just as CGaLH* but without the promise of the guiding state being quantumly preparable.

► **Theorem 6** (Classical containment of the classically guidable local Hamiltonian problem). *Let $k = \mathcal{O}(\log n)$. When δ is constant, we have that $\text{CGaLH}(k, \delta, \zeta)$ is in NP when ζ is constant and is in NqP when $\zeta = 1/\text{poly}(n)$. Here NqP is just as NP but with the Turing machine being allowed to run in quasi-polynomial time.*

Theorem 6 follows directly by applying the spectral amplification technique, as described in Section 3.2.

1.1.2 Quantum-classical probabilistically checkable proofs

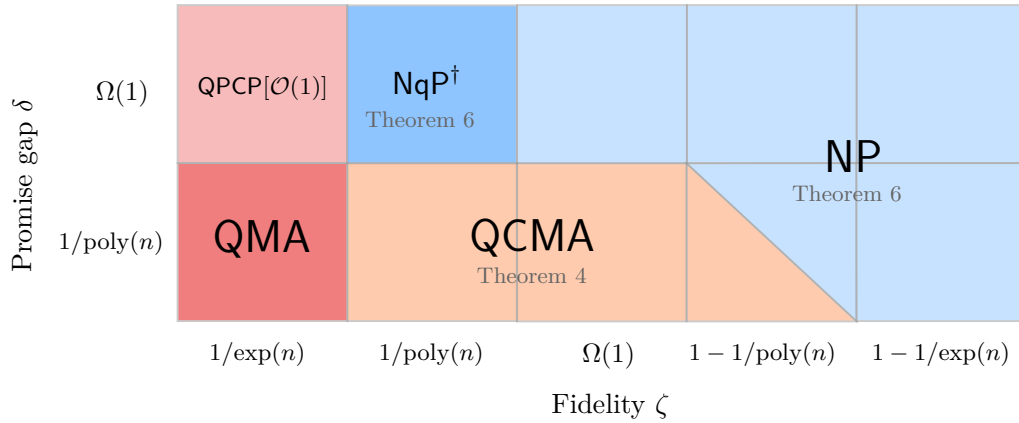
We introduce the notion of a *quantum-classical probabilistically checkable proof system* in the following way.

► **Definition 7** (Quantum-Classical Probabilistically Checkable Proofs (QCPCP)). *Let $n \in \mathbb{N}$ be the input size and $p, q : \mathbb{N} \rightarrow \mathbb{N}$, $c, s : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ with $c - s > 0$. A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ has a $(p(n), q(n), c, s)$ -QCPCP-verifier if there exists a quantum algorithm V which acts on an input $|x\rangle$ and a polynomial number of ancilla qubits, plus an additional bit string $y \in \{0, 1\}^{p(n)}$ from which it is allowed to read at most $q(n)$ bits (non-adaptively), followed by a measurement of the first qubit, after which it accepts only if the outcome is $|1\rangle$, and satisfies:*

Completeness. *If $x \in A_{\text{yes}}$, then there is a $y \in \{0, 1\}^{p(n)}$ such that the verifier accepts with probability at least c ,*

Soundness. *If $x \in A_{\text{no}}$, then for all $y \in \{0, 1\}^{p(n)}$ the verifier accepts with probability at most s .*

A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ belongs to $\text{QCPCP}[p, q, c, s]$ if it has a $(p(n), q(n), c, s)$ -QCPCP verifier. If $p(n) = \mathcal{O}(\text{poly}(n))$, $c = 2/3$, and $s = 1/3$, we simply write $\text{QCPCP}[q]$.



■ **Figure 2** Complexity characterization of $\text{CGaLH}^*(k, \delta, \zeta)$ over parameter regime δ and ζ , for $k = \mathcal{O}(1)$. Any classification indicates completeness for the respective complexity class, except for NqP , for which we only know containment (indicated by the “ \dagger ”). Here completeness for certain parameter combinations means that for all functions of the indicated form, the problem is contained in the complexity class, and for a subset of these functions the problem is also hard. The results for $\text{QPCP}[\mathcal{O}(1)]$ and QMA follow directly from [4] and [35].

We remark that there are likely several ways to characterise a “quantum-classical PCP”, with some being more or less natural than others. With that said, we believe that the above characterisation is well-motivated for the following reasons:

1. It is a natural definition following the structure of a QPCP verifier, now with proofs given as in the standard definition of QCMA . Moreover, one can show that the non-adaptiveness is not restrictive when the number of queries is constant (this is proved in the full version [48]).
2. $\text{QCPCP}[\mathcal{O}(1)]$ captures the power of BQP as well as NP (via the PCP theorem), which are both believed to be strictly different complexity classes. Since techniques used to prove the PCP theorem are difficult (or impossible) to translate to the quantum setting [5], studying $\text{QCPCP}[\mathcal{O}(1)]$ might provide a fruitful direction with which to obtain the first non-trivial lower bound on the complexity of $\text{QPCP}[\mathcal{O}(1)]$. Indeed, the currently best-known lower bound on the complexity of $\text{QPCP}[\mathcal{O}(1)]$ is only NP via the PCP theorem.

Given this definition for QCPCPs , our “quantum-classical” PCP conjecture is naturally formulated as follows.

► **Conjecture 8** (quantum-classical PCP conjecture). *There exists a constant $q \in \mathbb{N}$ such that $\text{QCMA} = \text{QCPCP}[q]$.*

If true, this conjecture would give a “ QCMA lower bound” on the power of quantum PCP systems, showing that a PCP theorem holds for (quantum) classes above NP , taking a step towards proving the quantum PCP conjecture. If it is false, but the quantum PCP conjecture is true, then this suggests that QPCP systems must take advantage of the “quantumness” of their proofs to obtain a probabilistically checkable proof system. In particular, since $\text{QCMA} \subseteq \text{QMA}$, this would imply the existence of a quantum PCP system for every problem in QCMA , but *not* a quantum-classical one, even though the problem admits a classical proof that can be efficiently verified when we are allowed to look at all of its bits.

Our main result regarding $\text{QCPCP}[\mathcal{O}(1)]$ is that we can provide a non-trivial upper bound on the complexity of the class.

► **Theorem 9** (Upper bound on QCPCP, from Theorem 25). $\text{QCPCP}[\mathcal{O}(1)] \subseteq \text{BQP}^{\text{NP}^{[1]}}$.

Here $\text{BQP}^{\text{NP}^{[1]}}$ is the class of all problems that can be solved by a BQP-verifier that makes a single query to an NP-oracle. The key idea behind the proof is that a quantum reduction can be used to transform a QCPCP verification circuit to a local Hamiltonian that is *diagonal* in the computational basis, and thus can be solved with a single query to an NP oracle.

An implication of Theorem 9 is that it can be used to show that under the assumption $\text{NP} \subseteq \text{BQP}$ and the quantum-classical PCP conjecture being true, we have that $\text{PH} \subseteq \text{BQP}$. This follows from the fact that $\text{NP}^{\text{BQP}} \subseteq \text{QCMA}$ and

$$\text{NP}^{\text{NP}} \subseteq \text{NP}^{\text{BQP}} \subseteq \text{BQP}^{\text{NP}} \subseteq \text{BQP}^{\text{BQP}} = \text{BQP},$$

where the first and the third “ \subseteq ” are by assumption, the second is by the assumption of Conjecture 8 to be true and the last equality follows from the fact that BQP is self-low. We then have that $\text{PH} \subseteq \text{BQP}$ follows by induction, just as is the case for BPP [50].⁶ Moreover, this would also imply that under these assumptions $\text{QCMA} \subseteq \text{BQP}$, since

$$\text{QCMA} \subseteq \text{QCPCP}[\mathcal{O}(1)] \subseteq \text{BQP}^{\text{NP}} \subseteq \text{BQP}^{\text{BQP}} \subseteq \text{BQP}.$$

Both of these implications would provide further evidence that it is unlikely that $\text{NP} \subseteq \text{BQP}$. However, it is known that there exists an oracle relative to which $\text{NP}^{\text{BQP}^A} \not\subseteq \text{BQP}^{\text{NP}^A}$ [3]. Nevertheless, this does not necessarily mean the premise (i.e. the quantum-classical PCP conjecture) is false: one can also easily construct an oracle separation between PCP and NP, and both classes are now known to be equal [27]. However, this suggests that, if Conjecture 8 is true, showing so requires non-relativizing techniques, just as was the case for the PCP theorem.

1.1.3 Three implications for the quantum PCP conjecture

We use our obtained results on QCPCP and CGaLH to obtain two new results and a new conjecture with respect to the quantum PCP conjecture. First, we give evidence that it is unlikely that there exists a *classical* reduction from a QPCP-system (see [4] or [16] for a formal definition) to a local Hamiltonian problem with a constant promise gap having the same properties as the known *quantum* reduction (see for example [16, 31]), unless $\text{BQP} \subseteq \text{QCPCP}[\mathcal{O}(1)] \subseteq \text{NP}$, something that is not expected to hold [2, 42].

► **Theorem 10** (No-go for classical polynomial-time reductions). *For any $\epsilon < 1/6$ there cannot exist a classical polynomial-time reduction from a QPCP $[\mathcal{O}(1)]$ verification circuit V to a local Hamiltonian H such that, given a proof $|\psi\rangle$,*

$$|\mathbb{P}[V \text{ accepts } |\psi\rangle] - (1 - \langle \psi | H | \psi \rangle)| \leq \epsilon,$$

unless $\text{QCPCP}[\mathcal{O}(1)] \subseteq \text{NP}$ (which would imply $\text{BQP} \subseteq \text{NP}$).

The proof is given in the full version [48]. This provides strong evidence that allowing for reductions to be quantum is indeed necessary to show equivalence between the gap amplification and proof verification formulations of the quantum PCP conjecture [5].

Second, our classical containment results of CGaLH with constant promise gap can be viewed as no-go theorems for a gap amplification procedure for QPCP having certain properties, as illustrated by the following result.

⁶ See also <https://blog.computationalcomplexity.org/2005/12/pulling-out-quantumness.html>.

► **Theorem 11** (No-go results for Hamiltonian gap amplification). *There cannot exist a polynomial time classical gap amplification procedure for the local Hamiltonian problem that preserves the fidelity between the ground space of the Hamiltonian and any classically evaluable state up to a*

- *multiplicative constant, unless $\text{QCMA} = \text{NP}$, or*
- *multiplicative inverse polynomial, unless $\text{QCMA} \subseteq \text{NqP}$.*

The theorem follows directly from Theorem 6. This result is analogous to the result of [7], which rules out a gap amplification procedure that preserves stoquasticity under the assumption that $\text{MA} \neq \text{NP}$ (or taking a different view, proving the existence of such gap amplifications would allow one to simultaneously prove that MA can be derandomized). Moreover, we point out that many Hamiltonian gadget constructions *do* satisfy such fidelity-preserving conditions, and indeed are precisely those that were used in [17] to improve the hardness results for the guided local Hamiltonian problem. We obtain similar results for the class MA by considering a variant of CGaLH that restricts the Hamiltonian to be stoquastic (See Appendix C in the full version [48]).

Third, we can use our results to formulate a stronger version of the NLTS theorem (and an alternative to the NLSS conjecture [28]), which we will call the *No Low-energy Classically evaluable States conjecture*. This conjecture can hopefully provide a new stepping stone towards proving the quantum PCP conjecture.

► **Conjecture 12** (NLCES conjecture). *There exists a family of local Hamiltonians $\{H_n\}_{n \in \mathbb{N}}$ on n qubits, and a constant $\beta > 0$, such that for sufficiently large n for every classically evaluable state $u \in \mathbb{C}^{2^n}$ as per Definition 1, we have that $\langle u | H_n | u \rangle \geq \lambda_0(H_n) + \beta$.*

Just as is the case for the NLSS conjecture and the NLTS theorem, the NLCES conjecture would, if proven to be true, not necessarily imply the quantum PCP conjecture. For example, it might be that there exist states that can be efficiently described classically but for which computing expectation values is hard (just as, for example, tensor network contraction is $\#\text{P}$ -hard in the worst case [14, 43]). Furthermore, as we have shown in this work, states with high energy but also a large fidelity with the ground state suffice as witnesses to decision problems on Hamiltonian energies, and these would not be excluded by a proof of the NLCES conjecture above. To make this more concrete, in the full version [48] we also formulate an even stronger version of the NLCES conjecture, which states that there must be a family of Hamiltonians for which no classically evaluable state has good fidelity with the low energy spectrum.

2 Preliminaries

2.1 Notation

We write $\lambda_i(A)$ to denote the i th eigenvalue of a Hermitian matrix A , ordered in non-decreasing order, with $\lambda_0(A)$ denoting the smallest eigenvalue (ground state energy). When we write $\|\cdot\|$ we refer to the operator norm when its input is a matrix and Euclidean norm for a vector.

2.2 Complexity theory

All complexity classes will be defined with respect to promise problems. To this end, we take a (promise) problem $A = (A_{\text{yes}}, A_{\text{no}})$ to consist of two non-intersecting sets $A_{\text{yes}}, A_{\text{no}} \subseteq \{0, 1\}^*$ (the YES and NO instances, respectively). We have that $A_{\text{inv}} = \{0, 1\}^* \setminus A_{\text{yes}} \cup A_{\text{no}}$ is the

set of all invalid instances, and we do not care how a verifier behaves on problem instances $x \in A_{\text{inv}}$ (i.e. it can accept or reject arbitrarily). We assume that the reader is familiar with the complexity classes used in this work, and else suggest reading the formal definitions in [48] or the complexity theory zoo (https://complexityzoo.net/Complexity_Zoo). However, since it is crucial to our construction, we will explicitly state the class UQCMA, which is just as QCMA but with a unique accepting witness in the YES-case.

► **Definition 13 (UQCMA).** *A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in UQCMA $[c, s]$ if and only if there exists a polynomial-time uniform family of quantum circuits $\{V_n\}$ and a polynomial p , where V_n takes as input a string $x \in \{0, 1\}^*$ with $|x| = n$, and a $p(n)$ -qubit witness quantum state $|\psi\rangle$ and decides on acceptance or rejection of x such that*

- *if $x \in A_{\text{yes}}$ then there exists a unique $y^* \in \{0, 1\}^{p(n)}$ such that V_n accepts $(x, |y^*\rangle)$ with probability $\geq c$, and for all $y \neq y^*$ we have that V_n accepts $(x, |y\rangle)$ with probability $\leq s$;*
- *if $x \in A_{\text{no}}$ then for every witness state $y \in \{0, 1\}^{p(n)}$, V_n accepts $(x, |y\rangle)$ with probability $\leq s$,*

where $c - s = 1/\text{poly}(n)$. If $c = 2/3$ and $s = 1/3$, we abbreviate to UQCMA.

In [6] it was shown that there exists a randomized reduction from QCMA to UQCMA, analogous to the Valiant-Vazirani theorem for NP [47].

Oracle access. For a (promise) class \mathcal{C} with complete (promise) problem A , the class $\mathcal{P}^{\mathcal{C}} = \mathcal{P}^A$ is the class of all (promise) problems that can be decided by a polynomial-time verifier circuit V with the ability to query an oracle for A . If V makes invalid queries (i.e. $x \in A_{\text{inv}}$), the oracle may respond arbitrarily. However, since V is deterministic, it is required to output the same final answer regardless of how such invalid queries are answered [29, 30]. Hence, the answer to any query outside of the promise set should not influence the final output bit. For a function f , we define $\mathcal{P}^{\mathcal{C}[f]}$ to be just as $\mathcal{P}^{\mathcal{C}}$ but with the additional restriction that V may ask at most $f(n)$ queries on an input of length n . One defines $\text{NP}^{\mathcal{C}}$ or $\text{NP}^{\mathcal{C}[f]}$ in the same way but replacing the polynomial-time deterministic verifier V by a nondeterministic polynomial-time verifier V' , taking an additional input $y \in \{0, 1\}^{p(n)}$ for some polynomial $p(n)$.

3 (Partial) proofs of a selection of results

In this section we will give some of the key lemmas and theorems which are behind the results presented in Section 1. The full proofs as well as in-depth discussions can be found in the full version [48].

3.1 QCMA-completeness of guidable local Hamiltonian problems

We prove a basic version of the reduction that shows that Guidable Local Hamiltonian problems are QCMA-hard in the inverse polynomial precision regime. Our construction is based on a combination of the ideas needed to show BQP-hardness for the Guided Local Hamiltonian problem [17, 18, 28] and the small penalty clock construction of [23].

The first obstruction one encounters in adopting the ideas from the BQP-hardness proofs of the Guided Local Hamiltonian problem to the guidable setting is the fact that QCMA verifiers, unlike BQP, have a proof register. In QCMA the promises of completeness and soundness are always with respect to computational basis state witnesses. Hence, these might no longer hold when *any* quantum state can be considered as witness: for example, in the NO-case there might be highly entangled states which are accepted with probability

10:12 Guidable Local Hamiltonian Problems with Implications to HASP and QPCP

$\geq 2/3$. When considering a circuit problem, the verifier can easily work around this by simply measuring the witness and then proceeding to verify with the resulting computational basis state. However, there is also another trick, which retains the unitarity of the verification circuit – and which we will denote as the “CNOT-trick” from now on – to force the witness to be classical, first used in proving QCMA-completeness of the *Low complexity low energy states* problem in [49].

► **Lemma 14** (The “CNOT-trick”). *Let $p(n) : \mathbb{N} \rightarrow \mathbb{R}_{>0}, q(n) : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be polynomials. Let U_n be a quantum polynomial-time verifier circuit that acts on an n -qubit input register A , a $p(n)$ -qubit witness register B and a $q(n)$ -qubit workspace register C , initialized to $|0\rangle^{\otimes q(n)}$. Denote Π_0 for the projection on the first qubit being zero. Let Q be the Marriott-Watrous operator of the circuit, defined as*

$$Q = \left(|x\rangle \otimes I_w \otimes |0\rangle^{\otimes q(n)} \right) U_n^\dagger \Pi_0 U_n \left(|x\rangle \otimes I_w \otimes |0\rangle^{\otimes q(n)} \right). \quad (1)$$

Consider yet another additional $p(n)$ -qubit workspace D initialized to $|0\rangle^{\otimes p(n)}$, on which U_n does not act. Then by prepending U_n with $p(n)$ CNOT-operations, each of which is controlled by a single qubit in register B and targeting the corresponding qubit in register D , the corresponding Marriott-Watrous operator becomes diagonal in the computational basis.

The corresponding lemma and proof can be found in the full version [48]. The next obstruction one faces is that in the QCMA setting there might be multiple proofs which all have exponentially close, or even identical, acceptance probabilities. The analysis of the BQP-hardness proof fails to translate directly to this setting, and another technique is needed. For this, we resort to (i) using the fact that QCMA is equal to UQCMA under randomized reductions and (ii) use a *small-penalty clock construction* of [23]. The key idea is to use a Feynman-Kiteav circuit-to-Hamiltonian mapping modified with a tunable parameter ϵ , which maps a quantum verification circuit U_n , consisting of T gates from a universal gate set of at most 2-local gates, taking input x and a quantum proof $|\psi\rangle \in (\mathbb{C}^2)^{\otimes \text{poly}(n)}$ to a k -local Hamiltonian of the form

$$H_{FK}^x = H_{\text{in}} + H_{\text{clock}} + H_{\text{prop}} + \epsilon H_{\text{out}}. \quad (2)$$

The value of k depends on the used construction. Intuitively, the first three terms check that the Hamiltonian is faithful to the computation and the last term shifts the energy level depending on the acceptance probability of the circuit. Just as in [23], we will use Kempe and Regev’s 3-local construction. A precise description of the individual terms in (2) can be found in [34], and will not be relevant for our work, except for the fact that H_{FK}^x has a polynomially bounded operator norm. The ground state of the first three terms $H_0 = H_{\text{in}} + H_{\text{clock}} + H_{\text{prop}}$ is given by the so-called *history state*, which is given in [34] by

$$|\eta(\psi)\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T U_t \dots U_1 |\psi\rangle |0\rangle |\hat{t}\rangle, \quad (3)$$

where $|\psi\rangle$ is the quantum proof and \hat{t} the unary representation of the time step of the computation given by

$$\hat{t} = |\underbrace{1 \dots 1}_t \underbrace{0 \dots 0}_{T-t}\rangle.$$

From the construction in [34], it is easily verified that if U_n accepts $(x, |\psi\rangle)$ with probability p then we have that the corresponding history state has energy

$$\langle \eta(\psi) | H_{FK}^x | \eta(\psi) \rangle = \epsilon \frac{1-p}{T+1}. \quad (4)$$

Though the core idea behind the small-penalty clock construction is identical to the one used in the BQP-hardness proof – rescaling the weight of the H_{out} term as compared to the other terms in a Feynman-Kiteav circuit-to-Hamiltonian mapping – the analysis differs: using tools from the Schrieffer-Wolff transformation one can find precise bounds on intervals in which the energies in the low-energy sector must lie, gaining fine control over the relation between the acceptance probabilities of the circuit and the low-energy sector of the Hamiltonian. The main lemma we use from [23] is adopted from the proof of Lemma 26 in their work.

► **Lemma 15** (Small-penalty clock construction, adopted from Lemma 26 in [23]). *Let U_n be a quantum verification circuit for inputs x , $|x| = n$, where U_n consists of $T = \text{poly}(n)$ gates from some universal gate-set using at most 2-local gates. Denote $P(\psi)$ for the probability that U_n accepts $(x, |\psi\rangle)$, and let H_{FK}^x be the corresponding 3-local Hamiltonian from the circuit-to-Hamiltonian mapping in [34] with a ϵ -factor in front of H_{out} , as in Eq. (2). Then for all $\epsilon \leq c/T^3$ for some constant $c > 0$, we have that low-energy subspace \mathcal{S}_ϵ of H , i.e.*

$$\mathcal{S}_\epsilon = \text{span}\{|\Phi\rangle : \langle \Phi | H | \Phi \rangle \leq \epsilon\}$$

has that its eigenvalues λ_i satisfy

$$\lambda_i \in \left[\epsilon \frac{1-P(\psi_i)}{T+1} - \mathcal{O}(T^3 \epsilon^2), \epsilon \frac{1-P(\psi_i)}{T+1} + \mathcal{O}(T^3 \epsilon^2) \right], \quad (5)$$

where $\{|\psi_i\rangle\}$ are the eigenstates of the Marriott-Watrous operator of the circuit U_n given by Eq. (1).

Having a QCMA-verifier with the CNOT-trick of Lemma 14 ensures that in Lemma 15 all $|\psi_i\rangle$ are computational basis states, as the CNOT-trick diagonalizes the Marriott-Watrous operator. The small-penalty clock construction, in combination with the CNOT-trick and some properties of the class QCMA, allows us to show QCMA-hardness of guidable local Hamiltonian problems in a wide range of parameter settings.

► **Theorem 16.** *CGaLH(k, δ, ζ) is QCMA-hard under randomized reductions for $k \geq 2$, $\zeta \in (1/\text{poly}(n), 1 - 1/\text{poly}(n))$ and $\delta = 1/\text{poly}(n)$.*

Proof. We will only state a “basic version” reduction, which uses basis states as guiding states which trivially satisfy the conditions of Definition 1, for which we prove completeness and soundness. One can improve its parameters in terms of the achievable fidelity and locality domains, which is done in the full manuscript [48].

The basic reduction. Let $\langle U_n, p_1, p_2 \rangle$ be a QCMA promise problem. By the result of [6], there exists randomized reduction to a UQCMA (which is QCMA but with a unique accepting witness in the YES-case) promise problem $\langle \hat{U}_n, \hat{p}_1, \hat{p}_2 \rangle$, $\hat{p}_1 - \hat{p}_2 \geq 1/q(n)$ for some polynomial q , which uses witnesses $y \in \{0, 1\}^{p(n)}$ for some polynomial $p(n)$ and uses at most $T = \text{poly}(n)$ gates. We will now apply the following modifications to the UQCMA instance:

1. First, we force the witness to be classical by adding another register to which we “copy” all bits of y (through CNOT operations), before running the actual verification protocol – i.e. we use the CNOT trick of Lemma 14, which diagonalizes the corresponding Marriott-Watrous operator in the computational basis.

2. We apply *error reduction* to the circuit. This is done by applying the so-called “Marriot and Watrous trick” for error reduction, described in [39], which allows one to repeat the verification circuit several times whilst re-using the same witness. It is shown in [39], Theorem 3.3, that for any quantum circuit V_n using $T = \text{poly}(n)$ 2-qubit gates which decides on acceptance or rejectance of an input x , $|x| = n$, using a $p(n)$ -qubit witness $|\psi\rangle$ for some polynomial p , satisfying completeness and soundness probabilities c, s such that $c - s \geq 1/q(n)$ there is another circuit \tilde{V}_n that again uses a $p(n)$ -qubit witness $|\psi\rangle$ but has completeness and soundness $1 - 2^{-r}$ and 2^{-r} , respectively, at the cost of using $\tilde{T} = \mathcal{O}(q^2 r T)$ gates.

Let the resulting protocol be denoted by $\langle \tilde{U}_n, \tilde{c}, \tilde{s} \rangle$, where \tilde{U}_n has an input register A , a witness register W and ancilla register B , uses $\tilde{T} = \mathcal{O}(q^2 r T)$ gates and has completeness and soundness $C = 1 - 2^{-r}$ and $\hat{s} = 2^{-r}$. We denote y^* for the (unique) witness with acceptance probability $\geq C$ in the YES-case. We keep r as a parameter to be tuned later in our construction. We will also write $P(y) := \Pr[\hat{U} \text{ accepts } (y)]$. Now consider the 4-local Hamiltonian

$$H^x = H_{yes} \otimes |0\rangle\langle 0|_D + H_{yes} \otimes |1\rangle\langle 1|_D, \quad (6)$$

where $H_{yes} = H_{\text{FK}}^x$ is the Hamiltonian given by Eq. (2) using the circuit \tilde{U}_n and parameter ϵ and H_{no} is given by

$$H_{no} = \sum_{i=0}^{R-1} |1\rangle\langle 1|_i + bI, \quad (7)$$

where R is the total size of the registers A, W, B and the clock register C , and $b > 0$ is yet another tunable parameter. Note that H_{no} has a *unique* ground state with energy b given by the all zeros state, and the spectrum after that increases in steps of 1 (and so it in particular has a *spectral gap* of 1). We also have that $\|H_{no}\| = R + b = \text{poly}(n)$. As a guiding state in the YES-case we will use the following basis state

$$|u_{yes}\rangle = |x\rangle_A |y^*\rangle_W |0 \dots 0\rangle_B |0\rangle_C |0\rangle_D, \quad (8)$$

which satisfies $(\langle \eta(y^*) | \langle 0|_D) |u_{yes}\rangle = 1/\sqrt{T+1} = \mathcal{O}(1/\text{poly}(N))$, with $|\eta(y^*)\rangle$ being the history state of witness y^* for Hamiltonian H_{yes} . In the NO-case, we will show that the state

$$|u_{no}\rangle = |0 \dots 0\rangle_{AWBC} |1\rangle_D, \quad (9)$$

will be in fact the ground state. We will now show that setting $b := \mathcal{O}(1/\tilde{T}^7)$ and $\epsilon := \mathcal{O}(1/\tilde{T}^5)$, our reduction achieves the desired result.

Completeness. Let us first analyse the YES-case. By Lemma 15, we have that the eigenvalue $\lambda(y)$ corresponding to the witness y^* is upper bounded by

$$\lambda(y^*) \leq \epsilon \frac{2^{-r}}{\tilde{T} + 1} + \mathcal{O}(\tilde{T}^3 \epsilon^2).$$

On the other hand, we have that for any $y \neq y^*$

$$\lambda(y) \geq \epsilon \frac{1 - 2^{-r}}{\tilde{T} + 1} - \mathcal{O}(\tilde{T}^3 \epsilon^2) = \Omega\left(\frac{1}{\tilde{T}^6}\right)$$

for our choice of ϵ and $r \geq 1$. Hence, for our choice of ϵ we must have that the ground state $|\psi\rangle$ of H_{yes} is unique and has a spectral gap that can be bounded as

$$\gamma(H_{\text{yes}}) \geq \epsilon \frac{1 - 2^{-r+1}}{\tilde{T} + 1} - \mathcal{O}(\tilde{T}^3 \epsilon^2) = \Omega\left(\frac{1}{\tilde{T}^6}\right), \quad (10)$$

for some $r \geq \Omega(1)$ (we will pick r to be much larger later). Let us consider the fidelity of the history state $|\eta(y^*)\rangle$ with the actual ground state. First, we have that the energy of $|\eta(y^*)\rangle$ is upper bounded by

$$\langle \eta(y^*) | H_{\text{yes}} | \eta(y^*) \rangle \leq \epsilon \frac{2^{-r}}{\tilde{T} + 1} = \mathcal{O}\left(\frac{2^{-r}}{\tilde{T}^6}\right),$$

which follows directly from Eq. 4 and the fact that $P(y^*) \geq 1 - 2^{-r}$. We can write $|\eta(y^*)\rangle$ in the eigenbasis of H_{yes} as $|\eta(y^*)\rangle = \alpha |\psi\rangle + \sqrt{1 - \alpha^2} |\psi^\perp\rangle$, for some real number $\alpha \in [0, 1]$, where $|\psi\rangle$ is the actual ground state of H_{yes} and $|\psi^\perp\rangle$ another state orthogonal to $|\psi\rangle$. We have that the energy of $|\eta(y^*)\rangle$ is upper bounded by

$$\langle \eta(y^*) | H_{\text{yes}} | \eta(y^*) \rangle \leq \epsilon \frac{2^{-r}}{\tilde{T} + 1} = \mathcal{O}\left(\frac{2^{-r}}{\tilde{T}^6}\right).$$

On the other hand, the energy of $|\eta(y^*)\rangle$ is lower bounded by

$$\langle \eta(y^*) | H_{\text{yes}} | \eta(y^*) \rangle = \alpha^2 \langle \psi | H_{\text{yes}} | \psi \rangle + (1 - \alpha^2) \langle \psi^\perp | H_{\text{yes}} | \psi^\perp \rangle \geq \Omega\left(\frac{1 - \alpha^2}{\tilde{T}^6}\right),$$

using the fact that H_{yes} is PSD. Combining the upper and lower bounds, we find

$$\alpha^2 = |\langle \eta(y^*) | \psi \rangle|^2 \geq 1 - \mathcal{O}(2^{-r}), \quad (11)$$

which can be made $\geq 1 - 2^{-c\tilde{T}}$ for some $r = c\tilde{T} + \mathcal{O}(1)$. Hence, we have that the fidelity of $|\eta(y^*)\rangle$ with the unique ground state of H can be lower bounded as

$$\begin{aligned} |\langle u_{\text{yes}} | \psi \rangle|^2 &\geq 1 - \left(\sqrt{1 - |\langle u_{\text{yes}} | (\eta(y^*)) | 0 \rangle|^2} + \sqrt{1 - |\langle (\eta(y^*)) | \langle 0 | \psi \rangle|^2} \right)^2 \\ &\geq 1 - \left(\sqrt{1 - \frac{1}{\tilde{T} + 1}} + 2^{-c\tilde{T}/2} \right)^2 \\ &\geq \Omega\left(\frac{1}{\tilde{T}}\right), \end{aligned}$$

as desired.

Soundness. We have that all witnesses y get accepted by \hat{U} with at most an exponentially small probability, and hence have that $H_{\text{yes}} \succeq \Omega(1/\tilde{T}^6)$. By our choice b we have therefore ensured that the ground state in the NO-case must be the state given by Eq. (9), which has energy $b = \Omega(1/\tilde{T}^7)$. Hence, the promise gap between YES and NO cases is $\delta = \Omega(1/\tilde{T}^7) = \Omega(1/q^2 T^8) = 1/\text{poly}(n)$.

In the full version [48] the rest of the proof can be found, which uses similar tricks as in [17, 18] to improve the basic construction in terms of the fidelity range and locality. ◀

Now that we have established QCMA-completeness for CGaLH*, we get QCMA-completeness for QGaLH *for free* for the same range of parameter settings, as the latter is a generalization of the former (containing CGaLH* as a special case), and containment holds by the same argument as used in the proof of Theorem 2 in [18]. However, with just a little bit more work we can see that QCMA-hardness for QGaLH actually persists even when the overlap is *exponentially* close to one. A proof of this is given in the full version [48].

3.2 Spectral amplification

In this subsection we will discuss spectral amplification, which is the key technique behind showing the containment results of Theorem 6. Let $H = \sum_{i=0}^{m-1} H_i$ be a Hamiltonian on n qubits which is a sum of k -local terms H_i , which satisfies $\|H\| \leq 1$. Since H is Hermitian, we can write H as

$$H = \sum_{i=0}^{2^n-1} \lambda_i |\psi_i\rangle \langle \psi_i|,$$

where $\lambda_i \in [-1, 1]$ (by assumption on the operator norm) denotes the i 'th eigenvalue of H with corresponding eigenvector $|\psi_i\rangle$. Consider a polynomial $P \in \mathbb{R}[x]$ of degree d , and write

$$P(x) = a_0 + a_1x + \cdots + a_dx^d.$$

The *polynomial spectral amplification* of H for P is then defined as

$$\begin{aligned} P(H) &= a_0I + a_1H + \cdots + a_dH^d \\ &= a_0I + a_1 \sum_{i=0}^{2^n-1} \lambda_i |\psi_i\rangle \langle \psi_i| + \cdots + a_d \sum_{i=0}^{2^n-1} \lambda_i^d |\psi_i\rangle \langle \psi_i| \\ &= \sum_{i=0}^{2^n-1} P(\lambda_i) |\psi_i\rangle \langle \psi_i|. \end{aligned}$$

Now for $\alpha \in [-1, 1]$, denote

$$\Pi_\alpha = \sum_{\{i: \lambda_i \leq \alpha\}} |\psi_i\rangle \langle \psi_i| \tag{12}$$

for the projection on all eigenstates of H which have eigenvalues at most α , which we will call a *low-energy projector* of H . Note that for any $\alpha \geq \lambda_0$, we must have that $\Pi_{\text{gs}}\Pi_\alpha = \Pi_\alpha\Pi_{\text{gs}} = \Pi_{\text{gs}}$. We can utilize such a projector to solve $\text{CGaLH}(k, \delta, \zeta)$, simply by computing $\|\Pi_\alpha |u\rangle\|$ for $\alpha = a$ given a classically evaluable state $|u\rangle$. To see why this works, note that in the YES-case, for the witness $\text{desc}(u)$ we have that $\|\Pi_a |u\rangle\| \geq \|\Pi_{\text{gs}} |u\rangle\| \geq \sqrt{\zeta}$ and in the NO-case we have that $\|\Pi_a |v\rangle\| = 0$ for all states, which means that the two cases are separated by $\sqrt{\zeta}$. However, it is unlikely that an efficient description exists of Π_a , and even if it did, it would not be k -local and therefore $\|\Pi_a |u\rangle\|$ would not even be necessarily efficiently computable.

The idea is now to approximate this low-energy projector Π_α by a polynomial in H . To see this, note that Π_α can be written exactly as

$$\Pi_\alpha = \frac{1}{2} (1 - \text{sgn}(H - \alpha I)),$$

where $\text{sgn}(x)$ is the sign function, which for our purposes is defined on $\mathbb{R} \rightarrow \mathbb{R}$ as

$$\text{sgn}(x) = \begin{cases} 1 & \text{if } x > 0, \\ -1 & \text{if } x \leq 0. \end{cases}$$

From [32] we can then use the polynomial approximation of the sign function, which can subsequently be shifted to obtain the desired approximate low-energy projector $\tilde{\Pi}_\alpha$.

► **Lemma 17** (Polynomial approximation to the sign function, from [32]). *For all $\delta' > 0, \epsilon' \in (0, 1/2)$ there exists an efficiently computable odd polynomial $P \in \mathbb{R}[x]$ of degree $d = \mathcal{O}\left(\frac{\log(1/\epsilon')}{\delta'}\right)$, such that*

- for all $x \in [-2, 2] : |P(x)| \leq 1$, and
- for all $x \in [-2, 2] \setminus (-\delta', \delta') : |P(x) - \text{sgn}(x)| \leq \epsilon'$.

Since Lemma 17 holds on the entire interval $[-2, 2]$, choosing any $\alpha \in [-1, 1]$ and scaling the $\text{sgn}(x)$ function with the factor $1/2$ will ensure that the error, as in the lemma, will be $\leq \epsilon/2$. Let $q_\alpha(x) : \mathbb{R} \rightarrow [0, 1]$ defined as $q_\alpha(x) = \frac{1}{2}(1 - \text{sgn}(x - \alpha))$ be this function, with polynomial approximation $Q_\alpha \in \mathbb{R}[x]$ of degree d . Note that Q_α can be written as a function of P as $Q_\alpha(x) = \frac{1}{2}(1 - P(x - \alpha))$. We will write $\tilde{\Pi}_\alpha = Q_\alpha(H)$ for the corresponding polynomial approximation of the approximate low-energy ground state “projector”. Note that $\tilde{\Pi}_\alpha$ is Hermitian (since H is Hermitian), but that $\tilde{\Pi}_\alpha$ is no longer necessarily a projector and therefore $\tilde{\Pi}_\alpha^2 \neq \tilde{\Pi}_\alpha$. If we now replace Π_α in $\|\Pi_\alpha |u\rangle\|$ by $\tilde{\Pi}_\alpha$, we get $\|\tilde{\Pi}_\alpha |u\rangle\| = \sqrt{\langle u | \tilde{\Pi}_\alpha^\dagger \tilde{\Pi}_\alpha |u\rangle} = \sqrt{\langle u | \tilde{\Pi}_\alpha^2 |u\rangle} = \sqrt{\langle u | (Q_\alpha(H))^2 |u\rangle}$, which means that we have to evaluate up to degree $2d$ powers of H . The next lemma (proof in full version [48]) will give an upper bound on the number of expectation values that have to be computed when evaluating a polynomial of H of degree d .

► **Lemma 18.** *Given access to a classically evaluable state $|u\rangle$, a Hamiltonian $H = \sum_{i=0}^{m-1} H_i$, where each H_i acts on at most k qubits non-trivially, and a polynomial $P[x]$ of degree d , there exists a classical algorithm that computes $\langle u | P(H) |u\rangle$ in $\mathcal{O}(m^d)$ computations of $\langle u | O_i |u\rangle$, where the observables $\{O_i\}$ are at most kd -local.*

All that remains to show is that for constant promise gap δ , using a good enough approximation $\tilde{\Pi}_\alpha$ with a suitable choice of α , will ensure that we can still distinguish the two cases in the $\text{CGaLH}(k, \delta, \zeta)$ problem in a polynomial (resp. quasi-polynomial) number of computations in m when $\zeta = \Omega(1)$ (resp. $\zeta = 1/\text{poly}(n)$).

► **Theorem 19.** *Let $H = \sum_{i=0}^{m-1} H_i$ be some Hamiltonian, and $\text{desc}(u)$ be a description of a classically evaluable state $u \in \mathbb{C}^{2^n}$. Let $a, b \in [-1, 1]$ such that $b - a \geq \delta$, where $\delta > 0$ and let $\zeta \in (0, 1]$. Consider the following two cases of H , with the promise that either one holds:*

- (i) H has an eigenvalue $\leq a$, and $\|\Pi_{gs} |u\rangle\|^2 \geq \zeta$ holds, or
- (ii) all eigenvalues of H are $\geq b$.

Then there exists a classical algorithm that distinguishes between cases (i) and (ii) using

$$\mathcal{O}\left(m^{c(\log(1/\sqrt{\zeta}))/\delta}\right)$$

computations of local expectation values, for some constant $c > 0$.

Proof. Let $\tilde{\Pi}_\alpha := Q_\alpha(H)$, where Q is a polynomial of degree d , be the approximate low-energy projector that approximates $\Pi_\alpha = \frac{1}{2}(1 - \text{sgn}(H - (\alpha I)))$. We set $\alpha := \frac{a+b}{2}$, $\delta' := \delta/2$ and $\epsilon' = 1/10$. We propose the following algorithm:

1. Compute $\|\tilde{\Pi}_\alpha |u\rangle\|$ using a polynomial of degree $2d$ where $d = \mathcal{O}(\log(1/\epsilon')/\delta')$, for $\epsilon' := \frac{1}{10}\sqrt{\zeta}$ and $\delta' = \delta/2$.
2. If $\|\tilde{\Pi}_\alpha |u\rangle\| \geq \frac{9}{10}\sqrt{\zeta}$ output (i), and otherwise output (ii).

Clearly, by Lemma 18, we have that this can be done in at most $\mathcal{O}\left(m^{c(\log(1/\sqrt{\zeta}))/\delta}\right)$ computations of expectation values of local observables, for some constant c . Let us now prove the correctness of the algorithm. Note that we can write $\tilde{\Pi}_\alpha$ as

$$\tilde{\Pi}_\alpha = \sum_{i=0}^{2^n-1} Q(\lambda_i) |\psi_i\rangle \langle \psi_i|,$$

where we have that

$$\begin{cases} 1 - \sqrt{\zeta}/2 \leq Q(\lambda_i) \leq 1 & \text{if } \lambda_i \leq a, \\ 0 \leq Q(\lambda_i) \leq \zeta/2 & \text{if } \lambda_i \geq b, \\ 0 \leq Q(\lambda_i) \leq 1 & \text{else,} \end{cases}$$

by Lemma 17. Let us analyse both cases (i) and (ii) separately.

(i) H has an eigenvalue $\leq a$, and $\|\Pi_{\text{gs}} |u\rangle\|^2 \geq \zeta$ holds.

$$\begin{aligned} \|\tilde{\Pi}_\alpha |u\rangle\| &\geq \|\tilde{\Pi}_\alpha \Pi_{\text{gs}} |u\rangle\| \\ &= \|\Pi_\alpha \Pi_{\text{gs}} |u\rangle - (\Pi_\alpha - \tilde{\Pi}_\alpha) \Pi_{\text{gs}} |u\rangle\| \\ &= \left\| \Pi_{\text{gs}} |u\rangle - \left(\sum_{i:\lambda_i \leq \alpha} (1 - Q(\lambda_i)) |\psi_i\rangle \langle \psi_i| - \sum_{i:\lambda_i > \alpha} Q(\lambda_i) |\psi_i\rangle \langle \psi_i| \right) \Pi_{\text{gs}} |u\rangle \right\| \\ &\geq \left\| \Pi_{\text{gs}} |u\rangle - \left(\sum_{i:\lambda_i \leq \alpha} \frac{1}{10} |\psi_i\rangle \langle \psi_i| \right) \Pi_{\text{gs}} |u\rangle \right\| \\ &= \left\| \Pi_{\text{gs}} |u\rangle - \frac{1}{10} \Pi_\alpha \Pi_{\text{gs}} |u\rangle \right\| \\ &= \left(1 - \frac{1}{10}\right) \|\Pi_{\text{gs}} |u\rangle\| \\ &\geq \frac{9}{10} \sqrt{\zeta}. \end{aligned}$$

(ii) All eigenvalues of H are $\geq b$. We must have that $\|\tilde{\Pi}_\alpha |u\rangle\| \leq \frac{1}{2} \sqrt{\zeta}$, since $\lambda_i \geq b$ for all $i \in \{0, \dots, 2^n - 1\}$.

Hence, we have that the promise gap between both cases is lower bounded by $\frac{9}{10} \sqrt{\zeta} - \frac{1}{2} \sqrt{\zeta} = \frac{2}{5} \sqrt{\zeta}$, which is $1/\text{poly}(n)$ when $\zeta \geq 1/\text{poly}(n)$. \blacktriangleleft

► **Remark 20.** It should be straightforward to adopt the same derivation as above to a more general setting by considering sparse matrices, a promise with respect to the fidelity with the low-energy subspace (i.e. all states with energy $\leq \lambda_0 + \gamma$ for some small γ), as well as $\epsilon > 0$ for ϵ -classically evaluable states.

3.3 Upper bound on QCPCP with constant proof queries

Here we show that QCPCP with a constant number of proof queries is contained in $\text{BQP}^{\text{NP}[1]}$, i.e. in BQP with only a single query to an NP-oracle. The full proof is rather long, but the idea is simple: just as is the case for QPCP, a *quantum* reduction can be used to transform a QCPCP system into a local Hamiltonian problem. However, since the proof is now classical, one can directly learn a diagonal (i.e. classical) Hamiltonian that captures the input/output behaviour of the QCPCP-circuit on basis state inputs. The main technical work required is to derive sufficient parameters in the reduction, thereby ensuring that the reduction succeeds with the desired success probability.

We will use the following two lemmas, whose proofs can be bound in the full version [48].

► **Lemma 21.** Let $H = \sum_{i \in [m]} w_i H_i$ be a k -local Hamiltonian consisting of weights $w_i \in [0, 1]$ such that $\sum_{i \in [m]} w_i = W$, and k -local terms H_i for which $\|H_i\| \leq 1$ for all $i \in [m]$. Let $\Omega_{\geq \gamma} = \{i \mid w_i \geq \gamma\}$ and $\Omega_{< \gamma} = [m] \setminus \Omega_{\geq \gamma}$, for some parameter $\gamma \in [0, 1]$. Suppose $\tilde{H} = \sum_{i \in \Omega_{\geq \gamma}} \tilde{w}_i \tilde{H}_i$ is another Hamiltonian such that, for all $i \in \Omega_{\geq \gamma}$, we have $|\tilde{w}_i - w_i| \leq \epsilon_0$ and $\|H_i - \tilde{H}_i\| \leq \epsilon_1$. Then

$$\|H - \tilde{H}\| \leq m(\gamma + \epsilon_0) + (W + m\epsilon_0)\epsilon_1$$

► **Lemma 22** (Upper bound on the non-uniform double dixie cup problem). Given samples from the set $N = [n]$, according to a distribution \mathcal{P} , consider the subset $M_\gamma \subseteq N$ such that $M_\gamma = \{i \in N : \mathcal{P}(i) \geq \gamma\}$, for some $\gamma \in [0, 1]$. Let $T_m^{\mathcal{P}}(M)$ be the random variable indicating the first time that all elements in M_γ have been sampled at least m times when sampling from N over the distribution \mathcal{P} . Write $T_m(S)$ when the distribution over some set S is uniform. Then we have that

$$\mathbb{E}[T_m^{\mathcal{P}}(M_\gamma)] \leq \mathbb{E}[T_m(\lceil 1/\gamma \rceil)],$$

where $\mathbb{E}[T_m(\lceil 1/\gamma \rceil)] = \lceil 1/\gamma \rceil \ln \lceil 1/\gamma \rceil + (m-1)\lceil 1/\gamma \rceil \ln \ln \lceil 1/\gamma \rceil + \mathcal{O}(\lceil 1/\gamma \rceil)$.

Let us now consider the quantum algorithm used to learn the diagonal Hamiltonian whose spectrum encodes the acceptance probabilities of the QCPCP-verifier. Let V_x be the QCPCP-verifier circuit with the input x hardcoded into it. The idea of the algorithm is that it runs V_x many times, simultaneously gathering statistics on which indices are most likely to be queried by V_x (which is independent of the proof when the verifier is non-adaptive) as well as the probability of acceptance given that the proof locally looks like a string $z \in \{0, 1\}^q$. For every run, indexed by $t \in [T]$ for some $T \in \mathbb{N}$, this generates a tuple $O^{t,z} = ((i_1^{t,z}, \dots, i_q^{t,z}), o^{t,z})$, in which the proof y was supposed to be queried at indices i_1, \dots, i_q , and in which those bits were assigned the values $y_{i_1} = z_1, \dots, y_{i_q} = z_q$, and where o is the accept/reject measurement outcome. It repeats this process T times for every z . The resulting algorithm can be specified as follows:

1. For $z \in \{0, 1\}^q$:
 - a. Run V_x for a total of T times to obtain samples $\{O^{t,z}\}_{t \in [T]}$.
 - b. For all observed $(i_1^{t,z}, \dots, i_q^{t,z})$, set

$$\tilde{\lambda}_{x, (i_1, \dots, i_q)}(z) := \frac{\# \text{ times } o^{t,z} = 1 \text{ and } i_1, \dots, i_q \text{ observed}}{\# \text{ times } i_1, \dots, i_q \text{ observed}}.$$

2. Set

$$\tilde{\mathcal{P}}_x(i_1, \dots, i_q) = \sum_{z \in \{0, 1\}^q} \frac{\# \text{ times } (i_1^{t,z}, \dots, i_q^{t,z}) \text{ observed}}{2^{qT}},$$

3. For any estimated $\tilde{\mathcal{P}}_x(i_1, \dots, i_q) \leq \gamma$ remove both $\tilde{\mathcal{P}}_x(i_1, \dots, i_q)$ and associated $\tilde{\lambda}_{x, (i_1, \dots, i_q)}(z)$ for all z , and output all of the remaining ones.

The resulting diagonal Hamiltonian will then be constructed as

$$\tilde{H}_x = \sum_{(i_1, \dots, i_q) \in \Omega_{\geq \gamma}} \tilde{\mathcal{P}}_x(i_1, \dots, i_q) \tilde{H}_{x, (i_1, \dots, i_q)},$$

where

$$\tilde{H}_{x, (i_1, \dots, i_q)} = \sum_{z \in \{0, 1\}^q} (1 - \tilde{\lambda}_{x, (i_1, \dots, i_q)}(z)) |z\rangle \langle z|_{i_1, \dots, i_q}.$$

10:20 Guidable Local Hamiltonian Problems with Implications to HASP and QPCP

By Lemma 21 we can upper bound the difference between the Hamiltonian and the learned Hamiltonian. The next lemma shows that all relevant parameters can be learned up to inverse polynomial precision in polynomial time.

► **Lemma 23.** *Let $q \in \mathbb{N}$ be some constant and x an input with $|x| = n$. Consider a QCPCP[q] protocol with verification circuit V_x (which is V but with the input x hardcoded into the circuit), and proof $y \in \{0, 1\}^{p(n)}$, and let*

$$\mathcal{P}_x(i_1, \dots, i_q) = \mathbb{P}[V_x \text{ queries the proof at indices } (i_1, \dots, i_q)]$$

and

$$\lambda_{x, (i_1, \dots, i_q)}(z) = \mathbb{P} \left[\begin{array}{l} V_x \text{ accepts given proof bits } i_1, \dots, i_q \text{ are queried} \\ \text{and are given by } y_{i_1} = z_1, \dots, y_{i_q} = z_q \end{array} \right].$$

Let $\Omega = \{(i_1, \dots, i_q) : i_j \in [p(n)], \forall j \in [q]\}$, $\Omega_{\geq \gamma} = \{(i_1, \dots, i_q) \in \Omega \mid \mathcal{P}_x(i_1, \dots, i_q) \geq \gamma\}$ and $\Omega_{< \gamma} = \Omega \setminus \Omega_{\geq \gamma}$, for some parameter $\gamma \in [0, 1]$. Then there exists a quantum algorithm that, for all $(i_1, \dots, i_q) \in \Omega_{\geq \gamma}$ and all $z \in \{0, 1\}^q$, provides estimates $\tilde{\mathcal{P}}_x(i_1, \dots, i_q)$ and $\tilde{\lambda}_{x, (i_1, \dots, i_q)}(z)$ such that

$$|\tilde{\mathcal{P}}_x(i_1, \dots, i_q) - \mathcal{P}_x(i_1, \dots, i_q)| \leq \epsilon_0,$$

and

$$|\tilde{\lambda}_{x, (i_1, \dots, i_q)}(z) - \lambda_{x, (i_1, \dots, i_q)}(z)| \leq \epsilon_1,$$

with probability $1 - \delta$, and runs in time $\text{poly}(n, 1/\gamma, 1/\epsilon_0, 1/\epsilon_1, 1/\delta)$.

Proof. Let us now show that there exists a T not too large such that the criteria of the theorem are satisfied. Since the $\mathcal{P}_x(i_1, \dots, i_q)$ form a discrete distribution over the set Ω , where $|\Omega| = \binom{n}{q} \leq \left(\frac{en}{q}\right)^q$ which for constant q is polynomial in n , we know by a standard result in learning theory (see for example [19]) that a total of

$$\Theta \left(\frac{|\Omega| + \log(1/\delta_0)}{\epsilon_0^2} \right)$$

samples of $O^{t,z}$ (the “ z ”-value is in fact irrelevant here) suffices to get, with probability at least $1 - \delta_0$, estimates $\tilde{\mathcal{P}}_x(i_1, \dots, i_q)$ which satisfy

$$|\tilde{\mathcal{P}}_x(i_1, \dots, i_q) - \mathcal{P}_x(i_1, \dots, i_q)| \leq \epsilon_0.$$

To learn estimates $\tilde{\lambda}_{x, (i_1, \dots, i_q)}(z)$ for a single index configuration (i_1, \dots, i_q) and proof configuration z , Hoeffding’s inequality tells us that we only need

$$\mathcal{O} \left(\frac{\log(1/\delta_1)}{\epsilon_1^2} \right)$$

samples of $O^{t,z}$ to have that $|\tilde{\lambda}_{x, (i_1, \dots, i_q)}(z) - \lambda_{x, (i_1, \dots, i_q)}(z)| \leq \epsilon_1$, with probability $1 - \delta_1$. This means that any index configuration (i_1, \dots, i_q) such that $\mathcal{P}_x(i_1, \dots, i_q) \geq \gamma$ needs to appear $\mathcal{O} \left(\frac{\log(1/\delta_1)}{\epsilon_1^2} \right)$ many times, to get a good estimate of $\tilde{\lambda}_{x, (i_1, \dots, i_q)}(z)$. Lemma 22 shows that the expected number of samples needed such that this condition is met is upper bounded by

$$\left\lceil \frac{1}{\gamma} \right\rceil \ln \left\lceil \frac{1}{\gamma} \right\rceil + \left(\mathcal{O} \left(\frac{\log(1/\delta_1)}{\epsilon_1^2} \right) - 1 \right) \left\lceil \frac{1}{\gamma} \right\rceil \ln \ln \left\lceil \frac{1}{\gamma} \right\rceil + \mathcal{O} \left(\left\lceil \frac{1}{\gamma} \right\rceil \right),$$

which by Markov's inequality means that

$$\frac{1}{\delta_\lambda} \left(\left\lceil \frac{1}{\gamma} \right\rceil \ln \left\lceil \frac{1}{\gamma} \right\rceil + \left(\mathcal{O} \left(\frac{\log(1/\delta_1)}{\epsilon_1^2} \right) - 1 \right) \left\lceil \frac{1}{\gamma} \right\rceil \ln \ln \left\lceil \frac{1}{\gamma} \right\rceil + \mathcal{O} \left(\left\lceil \frac{1}{\gamma} \right\rceil \right) \right)$$

samples of $O^{t,z}$ suffice to turn this into an algorithm that achieves success probability $\geq 1 - \delta_\lambda$. To ensure that our entire algorithm succeeds with probability $1 - \delta$, we require that

$$(1 - \delta_\lambda)^{2^q} (1 - \delta_0) (1 - \delta_1)^{2^q \lceil \frac{1}{\gamma} \rceil} \geq 1 - \delta,$$

which can be achieved by setting $\delta_\lambda = \delta/(2^{q+2})$, $\delta_0 = \delta/4$ and $\delta_1 = \delta/(\lceil 1/\gamma \rceil 2^{q+2})$. Both the statistics for probabilities over the set of indices, as well as the output probabilities, are gathered at the same time. This means that the requirements on the number of samples needed for both estimations can be met at the same time, therefore the total number of samples T that we must take satisfies

$$T \geq \max \left\{ \Theta \left(\frac{\lceil \frac{1}{\gamma} \rceil + \log \left(\frac{1}{\delta} \right)}{\epsilon_0^2} \right), \frac{2^{2(q+1)}}{\delta} \left(\left\lceil \frac{1}{\gamma} \right\rceil \ln \left\lceil \frac{1}{\gamma} \right\rceil + \mathcal{O} \left(\frac{q \log \left(\lceil \frac{1}{\gamma} \rceil / \delta \right)}{\epsilon_1^2} \right) \left\lceil \frac{1}{\gamma} \right\rceil \ln \ln \left\lceil \frac{1}{\gamma} \right\rceil \right) \right\},$$

which yields a total runtime of $\mathcal{O}(\text{poly}(n, \lceil 1/\gamma \rceil, 1/\delta, 1/\epsilon_1, 1/\epsilon_0))$ when $q = \mathcal{O}(1)$. ◀

Lemma 23 can then be combined with Lemma 21 to show that a diagonal Hamiltonian whose spectrum encodes the acceptance probabilities of V_x can be learned in polynomial time with high probability.

► **Lemma 24.** *Let $q \in \mathbb{N}$ be some constant, then there exists a quantum algorithm that can reduce any problem solvable by a QCPCP[q] protocol, without access to the proof y , to a diagonal Hamiltonian \tilde{H}_x with the following properties:*

- $x \in P_{yes} \Rightarrow \exists y \in \{0, 1\}^{p(n)} : \langle y | \tilde{H}_x | y \rangle \leq \frac{1}{3} + \epsilon$
- $x \in P_{no} \Rightarrow \forall y \in \{0, 1\}^{p(n)} : \langle y | \tilde{H}_x | y \rangle \geq \frac{2}{3} - \epsilon$.

This reduction succeeds with probability $1 - \delta$ and runs in time $\text{poly}(n, 1/\epsilon, 1/\delta)$.

Finally, from Lemma 24 the main theorem follows, as a BQP verifier can perform the quantum reduction and, conditioned on succeeding, solve the resulting diagonal local Hamiltonian problem making only a single query to the NP-oracle.

► **Theorem 25.** *For all constant $q \in \mathbb{N}$, we have that $\text{QCPCP}[q] \subseteq \text{BQP}^{\text{NP}[1]}$.*

The full proofs of Theorem 25 and Lemma 24 are given in the full version [48].

References

- 1 Scott Aaronson. Computational complexity: Why quantum chemistry is hard. *Nature Physics*, 5:707–708, October 2009. doi:10.1038/nphys1415.
- 2 Scott Aaronson. Bqp and the polynomial hierarchy. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, STOC '10, pages 141–150, New York, NY, USA, 2010. Association for Computing Machinery. arXiv:0910.4698. doi:10.1145/1806689.1806711.
- 3 Scott Aaronson, DeVon Ingram, and William Kretschmer. The Acrobatics of BQP. In Shachar Lovett, editor, *37th Computational Complexity Conference (CCC 2022)*, volume 234 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:17, 2022. arXiv:2111.10409. doi:10.4230/LIPIcs.CCC.2022.20.
- 4 Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh Vazirani. The detectability lemma and quantum gap amplification. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, pages 417–426, New York, NY, USA, 2009. arXiv:0811.3412. doi:10.1145/1536414.1536472.

- 5 Dorit Aharonov, Itai Arad, and Thomas Vidick. Guest column: the quantum pcp conjecture. *SIGACT News*, 44(2):47–79, June 2013. [arXiv:1309.7495](#). [doi:10.1145/2491533.2491549](#).
- 6 Dorit Aharonov, Michael Ben-Or, Fernando G.S.L. Brandão, and Or Sattath. The Pursuit of Uniqueness: Extending Valiant-Vazirani Theorem to the Probabilistic and Quantum Settings. *Quantum*, 6:668, March 2022. [arXiv:0810.4840](#). [doi:10.22331/q-2022-03-17-668](#).
- 7 Dorit Aharonov and Alex Bredariol Grilo. Stoquastic pcp vs. randomness. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1000–1023, 2019. [arXiv:1901.05270](#). [doi:10.1109/FOCS.2019.00065](#).
- 8 Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Rev. Mod. Phys.*, 90:015002, January 2018. [arXiv:1611.04471](#). [doi:10.1103/RevModPhys.90.015002](#).
- 9 Itai Arad. A note about a partial no-go theorem for quantum pcp. *Quantum Info. Comput.*, 11(11–12):1019–1027, November 2011. [arXiv:1012.3319](#).
- 10 Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, May 1998. [doi:10.1145/278298.278306](#).
- 11 Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: a new characterization of np. *J. ACM*, 45(1):70–122, January 1998. [doi:10.1145/273865.273901](#).
- 12 Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet K. Chan. Quantum algorithms for quantum chemistry and quantum materials science. *Chemical Reviews*, 120(22):12685–12717, October 2020. [arXiv:2001.03685](#). [doi:10.1021/acs.chemrev.9b00829](#).
- 13 Paul Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics*, 22(5):563–591, May 1980. [doi:10.1007/BF01011339](#).
- 14 Jacob D. Biamonte, Jason Morton, and Jacob W. Turner. Tensor network contractions for #SAT. *Journal of Statistical Physics*, 160:1389–1404, June 2015. [arXiv:1405.7375](#). [doi:10.1007/s10955-015-1276-z](#).
- 15 Michael J Bremner, Richard Jozsa, and Dan J Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2011. [arXiv:1005.1407](#). [doi:10.1098/rspa.2010.0301](#).
- 16 Harry Buhrman, Jonas Helsen, and Jordi Weggemans. Quantum PCPs: on Adaptivity, Multiple Provers and Reductions to Local Hamiltonians. *CoRR*, March 2024. [arXiv:2403.04841](#).
- 17 Chris Cade, Marten Folkertsma, Sevag Gharibian, Ryu Hayakawa, François Le Gall, Tomoyuki Morimae, and Jordi Weggemans. Improved Hardness Results for the Guided Local Hamiltonian Problem. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261, pages 32:1–32:19, 2023. [arXiv:2207.10250](#). [doi:10.4230/LIPIcs.ICALP.2023.32](#).
- 18 Chris Cade, Marten Folkertsma, and Jordi Weggemans. Complexity of the guided local Hamiltonian problem: Improved parameters and extension to excited states. *CoRR*, July 2022. [arXiv:2207.10097](#).
- 19 Clément L Canonne. A short note on learning discrete distributions. *CoRR*, February 2020. [arXiv:2002.11457](#).
- 20 Nai-Hui Chia, András Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pages 387–400, New York, NY, USA, 2020. Association for Computing Machinery. [arXiv:190.06151](#). [doi:10.1145/3357713.3384314](#).
- 21 Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, pages 151–158, New York, NY, USA, 1971. Association for Computing Machinery. [doi:10.1145/800157.805047](#).
- 22 Jordan Cotler, Hsin-Yuan Huang, and Jarrod R. McClean. Revisiting dequantization and quantum advantage in learning tasks. *CoRR*, December 2021. [arXiv:2112.00811](#).

- 23 Abhinav Deshpande, Alexey V. Gorshkov, and Bill Fefferman. Importance of the spectral gap in estimating ground-state energies. *PRX Quantum*, 3:040327, December 2022. arXiv:2007.11582. doi:10.1103/PRXQuantum.3.040327.
- 24 Irit Dinur. The pcp theorem by gap amplification. *J. ACM*, 54(3):12–es, June 2007. doi:10.1145/1236457.1236459.
- 25 Andrew Drucker. A PCP characterization of AM. *ICALP*, pages 581–592, July 2011. arXiv:1002.3664. doi:10.1007/978-3-642-22006-7_49.
- 26 Richard P Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6/7), 1982. doi:10.1007/BF02650179.
- 27 Lance Fortnow. The role of relativization in complexity theory. *Bulletin of the EATCS*, 52:229–243, 1994. URL: <https://bibbase.org/network/publication/fortnow-theroleofrelativizationincomplexitytheory-1994>.
- 28 Sevag Gharibian and François Le Gall. Dequantizing the quantum singular value transformation: hardness and applications to quantum chemistry and the quantum pcp conjecture. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022*, pages 19–32, New York, NY, USA, 2022. Association for Computing Machinery. arXiv:2111.09079. doi:10.1145/3519935.3519991.
- 29 Sevag Gharibian and Justin Yirka. The complexity of simulating local measurements on quantum systems. *Quantum*, 3:189, September 2019. arXiv:1606.05626. doi:10.22331/q-2019-09-30-189.
- 30 Oded Goldreich. On promise problems: A survey. In *Theoretical Computer Science: Essays in Memory of Shimon Even*, pages 254–290. Springer, 2006. doi:10.1007/11685654_12.
- 31 Alex B. Grilo. *Quantum proofs, the local Hamiltonian problem and applications*. PhD thesis, Université Sorbonne Paris Cité, April 2018. URL: <https://www.irif.fr/~abgrilo/thesis.pdf>.
- 32 Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by uniform spectral amplification. *CoRR*, July 2017. arXiv:1707.05391.
- 33 Dhawal Jethwani, François Le Gall, and Sanjay K. Singh. Quantum-Inspired Classical Algorithms for Singular Value Transformation. In Javier Esparza and Daniel Král', editors, *45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*, volume 170 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 53:1–53:14, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. arXiv:1910.05699. doi:10.4230/LIPIcs.MFCS.2020.53.
- 34 Julia Kempe and Oded Regev. 3-local hamiltonian is qma-complete. *Quantum Info. Comput.*, 3(3):258–264, May 2003. arXiv:0302079.
- 35 Alexei Y. Kitaev, Alexander Shen, and Mikhail N. Vyalyi. *Classical and quantum computation*. Number 47 in Graduate Studies in Mathematics. American Mathematical Society, 2002.
- 36 Leonid A. Levin. Universal sequential search problems. *Problemy peredachi informatsii*, 9(3):115–116, 1973.
- 37 Lin Lin and Yu Tong. Near-optimal ground state preparation. *Quantum*, 4:372, December 2020. arXiv:2002.12508. doi:10.22331/q-2020-12-14-372.
- 38 Hongbin Liu, Guang Hao Low, Damian S. Steiger, Thomas Häner, Markus Reiher, and Matthias Troyer. Prospects of quantum computing for molecular sciences. *Materials Theory*, 6(1):1–17, March 2022. arXiv:2102.10081. doi:10.1186/s41313-021-00039-z.
- 39 Chris Marriott and John Watrous. Quantum Arthur-Merlin games. *CCC*, pages 275–285, June 2004. arXiv:cs/0506068. doi:10.1007/s00037-005-0194-x.
- 40 Bryan O’Gorman, Sandy Irani, James Whitfield, and Bill Fefferman. Intractability of electronic structure in a fixed basis. *PRX Quantum*, 3:020322, May 2022. arXiv:2103.08215. doi:10.1103/PRXQuantum.3.020322.
- 41 David Poulin and Matthew B. Hastings. Markov entropy decomposition: A variational dual for quantum belief propagation. *Phys. Rev. Lett.*, 106:080403, February 2011. arXiv:1012.2050. doi:10.1103/PhysRevLett.106.080403.

- 42 Ran Raz and Avishay Tal. Oracle separation of bqp and ph. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 13–23, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3313276.3316315.
- 43 Norbert Schuch, Michael M. Wolf, Frank Verstraete, and J. Ignacio Cirac. Computational complexity of projected entangled pair states. *Phys. Rev. Lett.*, 98:140506, April 2007. arXiv:0611050. doi:10.1103/PhysRevLett.98.140506.
- 44 Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 217–228, New York, NY, USA, 2019. Association for Computing Machinery. arXiv:1807.04271. doi:10.1145/3313276.3316310.
- 45 Barbara M. Terhal and David P. DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games. *Quantum Information & Computation*, 4(2):134–145, March 2004. arXiv:quant-ph/0205133. doi:10.26421/QIC4.2-5.
- 46 Jules Tilly, Hongxiang Chen, Shuxiang Cao, Dario Picozzi, Kanav Setia, Ying Li, Edward Grant, Leonard Wossnig, Ivan Rungger, George H. Booth, et al. The variational quantum eigensolver: a review of methods and best practices. *Physics Reports*, 986:1–128, November 2022. arXiv:2111.05176. doi:10.1016/j.physrep.2022.08.003.
- 47 L G Valiant and V V Vazirani. Np is as easy as detecting unique solutions. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC '85, pages 458–463, New York, NY, USA, 1985. Association for Computing Machinery. doi:10.1145/22145.22196.
- 48 Jordi Weggemans, Marten Folkertsma, and Chris Cade. Guidable local hamiltonian problems with implications to heuristic ansatz state preparation and the quantum pcg conjecture. *CoRR*, February 2023. arXiv:302.11578.
- 49 Pawel Wocjan, Dominik Janzing, and Thomas Beth. Two QCMA-complete problems. *Quantum Information & Computation*, 3(6):635–643, November 2003. arXiv:quant-ph/0305090.
- 50 Stathis Zachos. Probabilistic quantifiers and games. *Journal of Computer and System Sciences*, 36(3):433–451, 1988. doi:10.1016/0022-0000(88)90037-2.