The Complexity of Simulating Local Measurements on Quantum Systems*†

Sevag Gharibian[‡1] **and Justin Yirka**[2]

1 **Virginia Commonwealth University, Richmond, USA**
sgharibian@vcu.edu
2 **Virginia Commonwealth University, Richmond, USA**
yirkajk@vcu.edu

Abstract

An important task in quantum physics is the estimation of local quantities for ground states of local Hamiltonians. Recently, [Ambainis, CCC 2014] defined the complexity class $\mathrm{P}^{\mathrm{QMA[log]}}$, and motivated its study by showing that the physical task of estimating the expectation value of a local observable against the ground state of a local Hamiltonian is $\mathrm{P}^{\mathrm{QMA[log]}}$-complete. In this paper, we continue the study of $\mathrm{P}^{\mathrm{QMA[log]}}$, obtaining the following results.

- The $\mathrm{P}^{\mathrm{QMA[log]}}$-completeness result of [Ambainis, CCC 2014] requires $O(\log n)$-local observables and Hamiltonians. We show that simulating even a *single qubit* measurement on ground states of 5-local Hamiltonians is $\mathrm{P}^{\mathrm{QMA[log]}}$-complete, resolving an open question of Ambainis.
- We formalize the complexity theoretic study of estimating two-point correlation functions against ground states, and show that this task is similarly $\mathrm{P}^{\mathrm{QMA[log]}}$-complete.
- $\mathrm{P}^{\mathrm{QMA[log]}}$ is thought of as "slightly harder" than QMA. We justify this formally by exploiting the hierarchical voting technique of [Beigel, Hemachandra, Wechsung, SCT 1989] to show $\mathrm{P}^{\mathrm{QMA[log]}} \subseteq \mathrm{PP}$. This improves the containment $\mathrm{QMA} \subseteq \mathrm{PP}$ [Kitaev, Watrous, STOC 2000].
- A central theme of this work is the subtlety involved in the study of oracle classes in which the oracle solves a *promise* problem. In this vein, we identify a flaw in [Ambainis, CCC 2014] regarding a $\mathrm{P}^{\mathrm{UQMA[log]}}$-hardness proof for estimating spectral gaps of local Hamiltonians. By introducing a "query validation" technique, we build on [Ambainis, CCC 2014] to obtain $\mathrm{P}^{\mathrm{UQMA[log]}}$-hardness for estimating spectral gaps under polynomial-time Turing reductions.

1998 ACM Subject Classification F.1.3 Complexity Measures and Classes

Keywords and phrases Complexity theory, Quantum Merlin Arthur (QMA), local Hamiltonian, local measurement, spectral gap

Digital Object Identifier 10.4230/LIPIcs.TQC.2017.2

1 Introduction

The use of computational complexity theory to study the inherent difficulty of computational problems has proven remarkably fruitful over the last decades. For example, the theory of NP-completeness [8, 21, 17] has helped classify the worst-case complexity of hundreds of computational problems which elude efficient classical algorithms. In the quantum setting,

* A full version of the paper is available at https://arxiv.org/abs/1606.05626.
† Part of this work was completed while SG was supported by a Government of Canada NSERC Banting Postdoctoral Fellowship and the Simons Institute for the Theory of Computing at UC Berkeley. SG acknowledges support from NSF grants CCF-1526189 and CCF-1617710.
‡ SG acknowledges support from NSF grants CCF-1526189 and CCF-1617710, an NSERC Banting Postdoctoral Fellowship and a Simons Postdoctoral Fellow at the University of California, Berkeley.

the study of a quantum analogue of NP, known as Quantum Merlin Arthur[1] (QMA), was started in 1999 by the seminal "quantum Cook-Levin theorem" of Kitaev [19], which showed that estimating the ground state energy of a given k-local Hamiltonian is QMA-complete for $k \geq 5$. Here, a k-local Hamiltonian[2] H can be thought of as a quantum constraint satisfaction system in which each quantum clause acts non-trivially on k qubits. The "largest total weight of satisfiable clauses" is given by the *ground state energy* of H, i.e. the smallest eigenvalue of H. Physically, the ground state energy and its corresponding eigenvector, the *ground state*, are motivated in that they represent the energy level and state of a given quantum system at low temperature, respectively. For this reason, since Kitaev's work [19], a number of physically motivated problems have been shown complete for QMA (see, e.g., [5] and [14] for surveys), a number of which focus on estimating ground state energies of local Hamiltonians.

In recent years, however, new directions in quantum complexity theory involving other physical properties of local Hamiltonians have appeared. For example, Brown, Flammia and Schuch [6] (also Shi and Zhang [25]) introduced a quantum analogue of #P, denoted #BQP, and showed that computing the ground state degeneracy or density of states of local Hamiltonians is #BQP-complete. Gharibian and Kempe [12] introduced cq-Σ_2, a quantum generalization of Σ_2^p, and showed that determining the smallest subset of interaction terms of a given local Hamiltonian which yields a frustrated ground space is cq-Σ_2-complete (and additionally, cq-Σ_2-hard to approximate). Gharibian and Sikora [13] showed that determining whether the ground space of a local Hamiltonian has an "energy barrier" is QCMA-complete, where QCMA [2] is Merlin-Arthur (MA) with a classical proof and quantum prover. Finally, and most relevant to this work, Ambainis [3] introduced $\mathrm{P}^{\mathrm{QMA[log]}}$, which is the class of decision problems decidable by a polynomial time Turing machine with logarithmically many queries to a QMA oracle (i.e. a quantum analogue of $\mathrm{P}^{\mathrm{NP[log]}}$). He showed that $\mathrm{P}^{\mathrm{QMA[log]}}$ captures the complexity of a very natural physical problem: "Simulating" a local measurement against the ground state of a local Hamiltonian (more formally, computing the expectation value of a given local observable against the ground state).

It is worth noting here that, given a local Hamiltonian, often one is not necessarily interested in a description of the *entire* ground state [14]. Rather, one may be interested in local quantities such as the evaluation of a local observable or of a correlation function. This makes $\mathrm{P}^{\mathrm{QMA[log]}}$ a well-motivated complexity class, whose study we continue here.

Our results (summarized under three headings)

1. $\mathrm{P}^{\mathrm{QMA[log]}}$-completeness of estimating local quantities. We begin with the study of two physically motivated problems. The first, APX-SIM, was formalized by Ambainis [3] (formal definitions in Section 2): *Given a k-local Hamiltonian H and an l-local observable A, estimate the expectation value of the measurement A against the ground state of H, i.e. estimate $\langle A \rangle := \langle \psi | A | \psi \rangle$ for $|\psi\rangle$ a ground state of H.* The second problem, which we introduce here and denote APX-2-CORR, is defined similarly to APX-SIM, except one is given observables A and B, and asked to estimate the *two-point correlation function* $\langle A \otimes B \rangle - \langle A \rangle \langle B \rangle$.

Previously, Ambainis [3] showed that APX-SIM is $\mathrm{P}^{\mathrm{QMA[log]}}$-complete for $O(\log n)$-local Hamiltonians and $O(\log n)$-local observables. From a physical standpoint, however, it is

[1] More accurately, QMA is Merlin-Arthur (MA) with a quantum proof and quantum verifier.

[2] $H \in \mathbb{C}^{2^n \times 2^n}$ is a Hermitian matrix with a succinct description $H = \sum_i H_i$, where each local clause $H_i \in \mathbb{C}^{2^k \times 2^k}$ acts non-trivially on k qubits. Implicitly, if H_i acts on a subset $S_i \subseteq [n]$ of qubits non-trivially, then more accurately one writes $H_i \otimes I_{[n] \setminus S_i}$. We write $H = \sum_i H_i$ for simplicity.

typically desirable to have $O(1)$-local Hamiltonians and observables, and whether $\mathrm{P}^{\mathrm{QMA[log]}}$-hardness holds in this regime was left as an open question. We thus first ask: *Is* APX-SIM *still hard for an* $O(1)$*-local Hamiltonian and* 1*-local observables?*

A priori, one might guess that simulating 1-local measurements might not be difficult — for example, the ground state energy of a 1-local Hamiltonian can be estimated efficiently. Yet, this intuition is incorrect: By embedding a 3-SAT instance ϕ into a 3-local Hamiltonian, and using the ability to repeatedly locally measure observable Z against single qubits of the ground state, we can extract a solution to ϕ! Thus, the 1-local observable case is at least NP-hard. Indeed, we show it is much harder, resolving Ambainis's open question.

▶ **Theorem 1.1.** *Given a* 5*-local Hamiltonian* H *on* n *qubits and a* 1*-local observable* A*, estimating* $\langle A \rangle$ *(i.e.* APX-SIM*) is* $\mathrm{P}^{\mathrm{QMA[log]}}$*-complete.*

Thus, measuring just a *single* qubit of a local Hamiltonian H's ground state with a fixed single-qubit observable A (in our construction, A is independent of H) is harder than QMA (assuming $\mathrm{QMA} \neq \mathrm{P}^{\mathrm{QMA[log]}}$, which is likely as otherwise co-QMA $\subseteq$ QMA).

Using similar techniques, we also show APX-2-CORR is $\mathrm{P}^{\mathrm{QMA[log]}}$-complete.

▶ **Theorem 1.2.** *Given a* 5*-local Hamiltonian* H *on* n *qubits and a pair of* 1*-local observables* A *and* B*, estimating* $\langle A \otimes B \rangle - \langle A \rangle \langle B \rangle$ *(i.e.* APX-2-CORR*) is* $\mathrm{P}^{\mathrm{QMA[log]}}$*-complete.*

2. An upper bound on the power of $\mathrm{P}^{\mathrm{QMA[log]}}$. Since $\mathrm{P}^{\mathrm{QMA[log]}}$ is thought of as "slightly harder" than QMA (note $\mathrm{QMA} \subseteq \mathrm{P}^{\mathrm{QMA[log]}}$), we next ask: *How much harder than QMA is* $\mathrm{P}^{\mathrm{QMA[log]}}$*?* Recall that QMA $\subseteq$ PP [20, 26, 23] (note [26] actually shows the stronger containment QMA $\subseteq \mathrm{A_0PP}$). Here, PP is the set of promise problems solvable in probabilistic polynomial time with *unbounded* error. Our next result shows that $\mathrm{P}^{\mathrm{QMA[log]}}$ is "not too much harder" than QMA in the following rigorous sense.

▶ **Theorem 1.3.** $\mathrm{P}^{\mathrm{QMA[log]}} \subseteq \mathrm{PP}$.

3. Estimating spectral gaps and oracles for promise problems. A central theme in this work is the subtlety involved in the study of oracle classes in which the oracle solves a *promise* problem (such as $\mathrm{P}^{\mathrm{QMA[log]}}$), as opposed to a decision problem (such as $\mathrm{P}^{\mathrm{NP[log]}}$, where $\mathrm{P}^{\mathrm{NP[log]}}$ is $\mathrm{P}^{\mathrm{QMA[log]}}$ except with an NP oracle). As discussed in "Proof techniques and discussions" below, the issue is that a P machine cannot in general determine if the query it makes to a QMA oracle satisfies the promise gap of the oracle. For queries which violate this promise, the oracle is allowed to give an arbitrary answer. We observe that this point appears to have been missed in [3], rendering a claimed proof that determining the spectral gap of a given $O(\log n)$-local Hamiltonian H is $\mathrm{P}^{\mathrm{UQMA[log]}}$-hard incorrect. ($\mathrm{P}^{\mathrm{UQMA[log]}}$ is $\mathrm{P}^{\mathrm{QMA[log]}}$ except with a Unique QMA oracle. Unique QMA is roughly QMA with a unique accepting quantum witness in the YES case.) Our last result both shows how to overcome this difficulty (at the expense of obtaining a "slightly weaker" hardness claim involving a Turing reduction, whereas [3] claimed hardness under a mapping reduction), and improves the locality of H to $O(1)$.

▶ **Theorem 1.4.** *Given a* 4*-local Hamiltonian* H*, estimating its spectral gap (i.e. the problem* SPECTRAL-GAP*) is* $\mathrm{P}^{\mathrm{UQMA[log]}}$*-hard under polynomial time Turing reductions.*

Proof techniques and discussion

1. $\mathrm{P}^{\mathrm{QMA[log]}}$-completeness for estimating local quantities. The proofs of our first two $\mathrm{P}^{\mathrm{QMA[log]}}$-hardness results (Theorem 1.1 and Theorem 1.2) are similar, so we focus on APX-SIM here. Intuitively, our aim is simple: To design our local Hamiltonian H so that its

ground state encodes a so-called history state[3] [19] $|\psi\rangle$ for a given $\mathrm{P}^{\mathrm{QMA[log]}}$ instance, such that measuring observable Z on the designated "output qubit" of $|\psi\rangle$ reveals the answer of the computation. At a high level, this is achieved by combining a variant of Kitaev's circuit-to-Hamiltonian construction [19] (which forces the ground state to follow the P circuit) with Ambainis's "query Hamiltonian" [3] (which forces the ground state to encode correctly answered queries to the QMA oracle). Making this rigorous requires developing a few ideas, including: A careful analysis of Ambainis's query Hamiltonian's ground space when queries violating the promise gap of the oracle are allowed (Lemma 3.1), a simple but useful corollary (Cor. 2.3) of Kempe, Kitaev, and Regev's Projection Lemma [18] (Corollary 2.3, showing that any low energy state of H must be close to a valid history state), and application of Kitaev's unary encoding trick [19] to bring the locality of the Hamiltonian H down to $O(1)$ (Lemma 3.2).

Next, to show containment of APX-2-CORR in $\mathrm{P}^{\mathrm{QMA[log]}}$ (Theorem 1.2), a natural approach would be to run Ambainis's $\mathrm{P}^{\mathrm{QMA[log]}}$ protocol for APX-SIM independently for each term $\langle A \otimes B \rangle$, $\langle A \rangle$, and $\langle B \rangle$. However, if a cheating prover does not send the *same* ground state $|\psi\rangle$ for each of these measurements, soundness of the protocol can be violated. To circumvent this, we exploit a trick of Chailloux and Sattath [7] from the setting of QMA(2): we observe that the correlation function requires only knowledge of the two-body reduced density matrices $\{\, \rho_{ij} \,\}$ of $|\psi\rangle$. Thus, a prover can send classical descriptions of the $\{\, \rho_{ij} \,\}$ along with a "consistency proof" for the QMA-complete Consistency problem [22].

2. An upper bound on the power of $\mathrm{P}^{\mathrm{QMA[log]}}$. We now move to our third result, which is perhaps the most technically involved. To show $\mathrm{P}^{\mathrm{QMA[log]}} \subseteq \mathrm{PP}$ (Theorem 1.3), we exploit the technique of *hierarchical voting* (used by Beigel, Hemachandra, and Wechsung [4] to show $\mathrm{P}^{\mathrm{NP[log]}} \subseteq \mathrm{PP}$), in conjunction with the QMA strong amplification results of Marriott and Watrous [23]. The intuition is best understood in the context of $\mathrm{P}^{\mathrm{NP[log]}}$ [4]. There, the PP machine first attempts to *guess* the answers to each NP query by picking random assignments to the SAT formula ϕ_i representing query i, in the hope of guessing a satisfying assignment for ϕ_i. Since such a guess can succeed only if ϕ_i is satisfiable, it can be seen that the lexicographically *largest* string y^* attainable by this process must be the correct query string (i.e. string of query answers). The scheme then uses several rounds of "hierarchical voting," in which lexicographically smaller query strings reduce their probability of being output to the point where y^* is guaranteed to be the "most likely" query string output. While the quantum variant of this scheme we develop is quite natural, its analysis is markedly more involved than the classical setting due to both the bounded-error nature of QMA and the possibility of "invalid queries" violating the QMA promise gap. (For example, it is no longer necessarily true that the lexicographically largest obtainable y^* is a "correct" query string.)

3. Estimating spectral gaps and oracles for promise problems. Finally, we discuss our fourth result and the theme of "invalid queries". Assume that all calls by the $\mathrm{P}^{\mathrm{QMA[log]}}$ machine to the QMA oracle Q are for an instance (H, a, b) of the Local Hamiltonian Problem (LH): *Is the ground state energy of H at most a (YES case), or at least b (NO case), for $b - a \geq 1/\mathrm{poly}(n)$?* Unfortunately, a P machine cannot in general tell whether the instance (H, a, b) it feeds to Q satisfies the promise conditions of LH (i.e. the ground state energy

[3] A *history state* can be seen as a quantum analogue of the "tableaus" which appear in the proof of the Cook-Levin theorem, i.e. a history state encodes the history of a quantum computation. In contrast to tableaus, however, the history encodes information in quantum superposition.

may lie in the interval (a, b)). If the promise is violated, we call such a query *invalid*, and in this case Q is allowed to either accept or reject. This raises the issue of how to ensure a YES instance (or NO instance) of a $\mathrm{P}^{\mathrm{QMA[log]}}$ problem is well-defined. To do so, we stipulate (see, e.g., Definition 3 of Goldreich [16]) that the P machine must output the *same* answer regardless of how any invalid queries are answered by the oracle. As mentioned earlier, this point appears to have been missed in [3], where all queries were assumed to satisfy the LH promise. This results in the proofs of two key claims of [3] being incorrect. The first claim was used in the proof of $\mathrm{P}^{\mathrm{QMA[log]}}$-completeness for APX-SIM (Claim 1 in [3]); we provide a corrected statement and proof in Lemma 3.1 (which suffices for the $\mathrm{P}^{\mathrm{QMA[log]}}$-hardness results in [3] regarding APX-SIM to hold).

The error in the second claim (Claim 2 of [3]), wherein $\mathrm{P}^{\mathrm{UQMA[log]}}$-hardness of determining the spectral gap of a local Hamiltonian is shown, appears arguably more serious. The construction of [3] requires a certain "query Hamiltonian" to have a spectral gap, which indeed holds if the $\mathrm{P}^{\mathrm{QMA[log]}}$ machine makes no invalid queries. However, if the machine makes invalid queries, this gap can close, and it is not clear how one can recover $\mathrm{P}^{\mathrm{QMA[log]}}$-hardness under mapping reductions. To overcome this, we introduce a technique of "query validation": Given a query to the QMA oracle, we would like to determine if the query is valid or "far" from valid. While it is not clear how a P machine alone can perform such "query validation", we show how to use a SPECTRAL GAP oracle to do so, allowing us to eliminate "sufficiently invalid" queries. Combining this idea with Ambainis's original construction [3], we show Theorem 1.4, i.e. $\mathrm{P}^{\mathrm{UQMA[log]}}$-hardness for SPECTRAL-GAP for $O(1)$-local Hamiltonians. Since our "query validation" requires a polynomial number of calls to the SPECTRAL-GAP oracle, this result requires a polynomial-time *Turing* reduction. Whether this can be improved to a mapping reduction is left as an open question.

Significance. The problems studied here explore the line of research recently initiated by Ambainis [3] on $\mathrm{P}^{\mathrm{QMA[log]}}$, and focus on central problems for local Hamiltonian systems. The complexity theoretic study of such problems is appealing in that it addresses the original motivation of celebrated physicist Richard Feynman in proposing quantum computers [10], who was interested in avenues for simulating quantum systems. Indeed, hardness results, such as Kitaev's Cook-Levin theorem, rigorously justify Feynman's intuition that such simulation problems are "hard". Our work (e.g. Theorem 1.1), in particular, strongly supports this view by demonstrating that even some of the "simplest" and most natural simulation tasks, such as measuring a *single qubit (!)* of a ground state, can be harder than QMA.

Our work on the complexity of estimating spectral gaps (Theorem 1.4) further highlights another theme: The subtleties which must be carefully treated when studying oracle classes for promise problems (such as $\mathrm{P}^{\mathrm{QMA[log]}}$). As quantum complexity theory commonly focuses on such promise problems, we believe this theme would potentially be of interest to a broader computer science audience.

Open questions. Although we resolve one of the open questions from [3], there are others we leave open, along with some new ones. Do our results for APX-SIM and APX-2-CORR hold for more restricted classes of Hamiltonians, such as 2-local Hamiltonians, local Hamiltonians on a 2D lattice, or specific Hamiltonian models of interest (see e.g. [9, 24] for QMA-completeness results for estimating ground state energies of the spin-1/2 Heisenberg anti-ferromagnet)? Is SPECTRAL-GAP $\mathrm{P}^{\mathrm{UQMA[log]}}$-complete or $\mathrm{P}^{\mathrm{QMA[log]}}$-complete (recall SPECTRAL-GAP $\in \mathrm{P}^{\mathrm{QMA[log]}}$, and [3] and our work together show $\mathrm{P}^{\mathrm{UQMA[log]}}$-hardness)? What is the relationship between $\mathrm{P}^{\mathrm{QMA[log]}}$ and $\mathrm{P}^{\mathrm{UQMA[log]}}$? Finally, what is the complexity of other physical tasks "beyond" estimating ground state energies?

Organization. Section 2 gives notation, formal definitions, and a corollary of the Projection Lemma. Section 3 shows various lemmas regarding Ambainis's query Hamiltonian. Section 4 proves Theorem 1.1. As the proof of Theorem 1.2 uses techniques similar to Theorem 1.1, we defer its proof to the full version of this article. Section 5 shows Theorem 1.3. Theorem 1.4 is given in Section 6. Full proofs of selected claims are deferred to the full version.

2 Preliminaries

Notation. For $x \in \{0,1\}^n$, $|x\rangle \in (\mathbb{C}^2)^{\otimes n}$ denotes the computational basis state labeled by x. Let $\mathcal{X}$ be a complex Euclidean space. Then, $\mathrm{L}(\mathcal{X})$ and $\mathrm{D}(\mathcal{X})$ denote the sets of linear and density operators acting on $\mathcal{X}$, respectively. For subspace $\mathcal{S} \subseteq \mathcal{X}$, $\mathcal{S}^\perp$ denotes the orthogonal complement of $\mathcal{S}$. For Hermitian operator H, $\lambda(H)$ and $\lambda(H|_{\mathcal{S}})$ denote the smallest eigenvalue of H and the smallest eigenvalue of H restricted to space $\mathcal{S}$, respectively. The spectral and trace norms are defined $\|A\|_\infty := \max\{\|A|v\rangle\|_2 : \||v\rangle\|_2 = 1\}$ and $\|A\|_{\mathrm{tr}} := \mathrm{Tr}\sqrt{A^\dagger A}$, respectively, where $:=$ denotes a definition. We set $[m] := \{1,\ldots,m\}$.

Definitions and lemmas. PP [15] is the set of decision problems for which there exists a polynomial-time probabilistic Turing machine which accepts any YES instance with probability $> 1/2$, and accepts any NO instance with probability $\leq 1/2$.

$\mathrm{P}^{\mathrm{QMA[log]}}$, defined by Ambainis [3], is the set of decision problems decidable by a polynomial-time deterministic Turing machine with the ability to query an oracle for a QMA-complete problem (e.g. the 2-local Hamiltonian problem (2-LH) [18]) $O(\log n)$ times, where n is the size of the input. 2-LH is defined as: Given a 2-local Hamiltonian H and inverse polynomially separated thresholds $a, b \in \mathbb{R}$, decide whether $\lambda(H) \leq a$ (YES-instance) or $\lambda(H) \geq b$ (NO-instance). Note that the P machine is allowed to make queries which violate the promise gap of 2-LH, i.e. with $\lambda(H) \in (a,b)$; in this case, the oracle can output either YES or NO. The P machine is nevertheless required to output the same final answer (i.e. accept or reject) regardless of how such "invalid" queries are answered [16].

For any P machine M making m queries to a QMA oracle, we use the following terminology throughout this article. A *valid* (*invalid*) query satisfies (violates) the promise gap of the QMA oracle. A *correct* query string $y \in \{0,1\}^m$ encodes a sequence of correct answers to all of the m queries. Note that for any invalid query of M, any answer is considered "correct", yielding the possible existence of multiple correct query strings. An *incorrect* query string is one which contains at least one incorrect query answer.

We now recall the definition of APX-SIM.

▶ **Definition 2.1** (APX-SIM$(H, A, k, l, a, b, \delta)$ (Ambainis [3])). Given a k-local Hamiltonian H, an l-local observable A, and real numbers a, b, and δ such that $a - b \geq n^{-c}$ and $\delta \geq n^{-c'}$, for n the number of qubits H acts on and $c, c' > 0$ some constants, decide:

- If H has a ground state $|\psi\rangle$ satisfying $\langle\psi|A|\psi\rangle \leq a$, output YES.
- If for any $|\psi\rangle$ satisfying $\langle\psi|H|\psi\rangle \leq \lambda(H) + \delta$, it holds that $\langle\psi|A|\psi\rangle \geq b$, output NO.

Next, we briefly review Kitaev's circuit-to-Hamiltonian construction from the "quantum Cook-Levin theorem" [19]. Given a quantum circuit $U = U_L \cdots U_1$ consisting of 1- and 2-qubit gates U_i and acting on registers Q (proof register) and W (workspace register), this construction maps U to a 5-local Hamiltonian $H = H_{\mathrm{in}} + H_{\mathrm{out}} + H_{\mathrm{prop}} + H_{\mathrm{stab}}$. Here, we use two key properties of $H_{\mathrm{in}} + H_{\mathrm{prop}} + H_{\mathrm{stab}}$. First, the null space of $H_{\mathrm{in}} + H_{\mathrm{prop}} + H_{\mathrm{stab}}$

is spanned by *history states*, which for any $|\psi\rangle$ have form

$$|\psi_{\text{hist}}\rangle = \sum_{t=0}^{L} U_t \cdots U_1 |\psi\rangle_Q |0\cdots 0\rangle_W |t\rangle_C \,, \tag{1}$$

where C is a clock register keeping track of time [19]. Second, we use the following lower bound[4] on the smallest non-zero eigenvalue of $H_{\text{in}} + H_{\text{prop}} + H_{\text{stab}}$:

▶ **Lemma 2.2** (Lemma 3 (Gharibian, Kempe [12])). *The smallest non-zero eigenvalue of* $\Delta(H_{\text{in}} + H_{\text{prop}} + H_{\text{stab}})$ *is at least* $\pi^2\Delta/(64L^3) \in \Omega(\Delta/L^3)$*, for* $\Delta \in \mathbb{R}^+$ *and* $L \geq 1$*. construction.*

A useful fact for complex unit vectors $|v\rangle$ and $|w\rangle$ is (see, e.g., Equation 1.33 of [11]):

$$\| |v\rangle\langle v| - |w\rangle\langle w| \|_{\text{tr}} = 2\sqrt{1 - |\langle v|w\rangle|^2} \leq 2 \| |v\rangle - |w\rangle \|_2 \,. \tag{2}$$

Next, let V denote a QMA verification circuit acting on M proof qubits with completeness c and soundness s. If one runs V on "proof" $\rho = I/2^M$, then for a YES instance, V accepts with probability $\geq c/2^M$ (since $I/2^M$ can be viewed as "guessing" a correct proof with probability $\geq 1/2^M$), and in a NO instance, V accepts with probability $\leq s$ (see, e.g., [23, 27]). The class PQP is defined analogously to BQP, except in the YES case, the verifier accepts with probability $> 1/2$, and in the NO case, the verifier accepts with probability $\leq 1/2$.

A corollary of the Projection Lemma. Finally, we give a simple but useful corollary of the Projection Lemma of Kempe, Kitaev, Regev [18]. The Projection Lemma, along with the proof of Corollary 2.3, are given in the full version.

▶ **Corollary 2.3.** *Let* $H = H_1 + H_2$ *be the sum of two Hamiltonians operating on some Hilbert space* $\mathcal{H} = \mathcal{S} + \mathcal{S}^\perp$*. The Hamiltonian* H_1 *is such that* $\mathcal{S}$ *is a zero eigenspace and the eigenvectors in* $\mathcal{S}^\perp$ *have eigenvalue at least* $J > 2\|H_2\|_\infty$*. Let* $K := \|H_2\|_\infty$*. Then, for any* $\delta \geq 0$ *and vector* $|\psi\rangle$ *satisfying* $\langle\psi| H |\psi\rangle \leq \lambda(H) + \delta$*, there exists a* $|\psi'\rangle \in \mathcal{S}$ *such that* $|\langle\psi|\psi'\rangle|^2 \geq 1 - \left(\frac{K+\sqrt{K^2+\delta(J-2K)}}{J-2K}\right)^2$.

3 Ambainis's Query Hamiltonian

We now show various results regarding Ambainis's "query Hamiltonian" [3], which intuitively aims to have its ground space contain correct answers to a sequence of QMA queries. Let U be a $\text{P}^{\text{QMA[log]}}$ computation, and let $H_{\mathcal{Y}_i}^{i,y_1\cdots y_{i-1}}$ be the 2-local Hamiltonian corresponding to the ith query made by U given that the answers to the previous $i-1$ queries are given by $y_1 \cdots y_{i-1}$. (Without loss of generality, we may assume $H_{\mathcal{Y}_i}^{i,y_1\cdots y_{i-1}} \succeq 0$ by adding multiples of the identity and rescaling.) The oracle query made at step i corresponds to an input $(H_{\mathcal{Y}_i}^{i,y_1\cdots y_{i-1}}, \epsilon, 3\epsilon)$ to 2-LH, for $\epsilon > 0$ a fixed inverse polynomial. Then, Ambainis's [3] $O(\log(n))$-local query Hamiltonian H acts on $\mathcal{X} \otimes \mathcal{Y}$, where $\mathcal{X} = (\mathcal{X}_i)^{\otimes m} = (\mathbb{C}^2)^{\otimes m}$ and $\mathcal{Y} = \otimes_{i=1}^m \mathcal{Y}_i$, such that $\mathcal{X}_i$ is intended to encode the answer to query i with $\mathcal{Y}_i$ encoding the

[4] This bound is stated as $\Omega(\Delta/L^3)$ in [12]; the constant $\pi^2/64$ can be derived from the analysis therein.

ground state of the corresponding query Hamiltonian $H_{\mathcal{Y}_i}^{i,y_1\cdots y_{i-1}}$. Specifically,

$$\begin{aligned} H &= \sum_{i=1}^{m} \frac{1}{4^{i-1}} \sum_{y_1,\ldots,y_{i-1}} \bigotimes_{j=1}^{i-1} |y_j\rangle\langle y_j|_{\mathcal{X}_j} \otimes \left(2\epsilon\,|0\rangle\langle 0|_{\mathcal{X}_i} \otimes I_{\mathcal{Y}_i} + |1\rangle\langle 1|_{\mathcal{X}_i} \otimes H_{\mathcal{Y}_i}^{i,y_1\cdots y_{i-1}}\right) \\ &=: \sum_{i=1}^{m} \frac{1}{4^{i-1}} \sum_{y_1,\ldots,y_{i-1}} M_{y_1\cdots y_{i-1}}. \end{aligned} \tag{3}$$

Recall from Section 2 that a sequence of query answers $y = y_1 \cdots y_m \in \{0,1\}^m$ is *correct* if it corresponds to a possible execution of U. Since U can make queries to its QMA oracle which violate the QMA promise gap, the set of correct y is generally not a singleton. However, we henceforth assume without loss of generality that U makes at least one valid query (i.e. which satisfies the QMA promise gap). For if not, then a P machine can solve such an instance by simulating the $\mathrm{P}^{\mathrm{QMA[log]}}$ machine on all possible (polynomially many) query strings $y \in \{0,1\}^m$. If U corresponds to a YES (NO) instance, then *all* query strings lead to accept (reject), which the P machine can verify. We now prove the following about H.

▶ **Lemma 3.1.** *Define for any $x \in \{0,1\}^m$ the space $\mathcal{H}_{x_1\cdots x_m} := \bigotimes_{i=1}^{m} |x_i\rangle\langle x_i| \otimes \mathcal{Y}_i$. Then, there exists a correct query string $x \in \{0,1\}^m$ such that the ground state of H lies in $\mathcal{H}_{x_1\cdots x_m}$. Moreover, suppose this space has minimum eigenvalue λ. Then, for any incorrect query string $y_1 \cdots y_m$, any state in $\mathcal{H}_{y_1\cdots y_m}$ has energy at least $\lambda + \frac{\epsilon}{4^m}$.*

As discussed in Section 1, Claim 1 of [3] proved a similar statement under the assumption that the correct query string x is unique. In that setting, [3] showed the ground state of H is in $\mathcal{H}_x$, and that for *all* query strings $y \neq x$, the space $\mathcal{H}_y$ has energy at least $\lambda + \frac{\epsilon}{4^{m-1}}$. However, in general invalid queries must be allowed, and in this setting this claim no longer holds — two distinct correct query strings can have eigenvalues which are arbitrarily close if they contain queries violating the promise gap. The key observation we make here is that even in the setting of non-unique x, a spectral gap between the ground space and all *incorrect* query strings can be shown. The proof is deferred to the full version of this article.

The next lemma converts H from an $O(\log n)$-local Hamiltonian to an $O(1)$-local one. Its proof uses Kitaev's unary encoding trick [19], and is given in the full version.

▶ **Lemma 3.2.** *For any $x \in \{0,1\}^m$, let $\hat{x}$ denote its unary encoding. Then, for any $\mathrm{P}^{\mathrm{QMA[log]}}$ circuit U acting on n bits and making $m \geq 1$ queries to a QMA oracle, there exists a mapping to a 4-local Hamiltonian H' acting on space $(\mathbb{C}^2)^{\otimes 2^m-1} \otimes \mathcal{Y}$ such that there exists a correct query string $x = x_1 \cdots x_m$ satisfying:*

1. *The ground state of H' lies in subspace $|\hat{x}\rangle\langle\hat{x}| \otimes \mathcal{Y}$.*
2. *For any state $|\psi\rangle$ in subspace $|\hat{x}'\rangle\langle\hat{x}'| \otimes \mathcal{Y}$ where either $\hat{x}'$ is not a unary encoding of a binary string x' or x' is an incorrect query string, one has $\langle\psi| H' |\psi\rangle \geq \lambda(H') + \epsilon/4^m$, for inverse polynomial ϵ.*
3. *For all strings $x' \in \{0,1\}^m$, H' acts invariantly on subspace $|\hat{x}'\rangle\langle\hat{x}'| \otimes \mathcal{Y}$.*
4. *The mapping can be computed in time polynomial in n (recall $m \in O(\log n)$).*

4 Measuring 1-local observables

Proof of Theorem 1.1. Containment in $\mathrm{P}^{\mathrm{QMA[log]}}$ was shown for $k, l \in O(\log n)$ in [3]; we show $\mathrm{P}^{\mathrm{QMA[log]}}$-hardness. Let U' be an arbitrary $\mathrm{P}^{\mathrm{QMA[log]}}$ circuit for instance Π, such that U' acts on workspace register W and query result register Q. Suppose U' consists of L' gates and makes $m = c\log(n)$ queries, for $c \in O(1)$ and n the input size. Without loss of generality,

U' can be simulated with a similar unitary U which treats Q as a *proof* register which it does not alter at any point: Namely, U does not have access to a QMA oracle, but rather reads bit Q_i whenever it desires the answer to the ith query. Thus, if a correct query string $y_1 \cdots y_m$ corresponding to an execution of U' on input x is provided in Q as a "proof", then the output statistics of U' and U are identical. We can also assume that Q is encoded not in binary, but in unary. Thus, Q consists of $2^m - 1 \in \text{poly}(n)$ bits. For simplicity, however, in our discussion we will speak of m-bit query strings $y = y_1 \cdots y_m$ in register Q.

Next, we map U to a 5-local Hamiltonian H_1 via a modification of the circuit-to-Hamiltonian construction of Kitaev [19], such that H_1 acts on registers W (workspace register), Q (proof register), and C (clock register). Recall (Section 2) that Kitaev's construction outputs Hamiltonian terms $H_{\text{in}} + H_{\text{prop}} + H_{\text{stab}} + H_{\text{out}}$. Set $H_1 = \Delta(H_{\text{in}} + H_{\text{prop}} + H_{\text{stab}})$ for Δ to be set as needed. It is crucial that H_{out} be omitted from H_1, as we require our final Hamiltonian H to enforce a certain structure on the ground space *regardless* of whether the computation should accept or reject. The job of "checking the output" is instead delegated to the observable A. Formally, H_1 has a non-trivial null space, which is its ground space, consisting of history states $|\psi_{\text{hist}}\rangle$ (Equation (1)) which simulate U on registers W and Q. These history states correctly simulate U' *assuming that* Q is initialized to a correct proof.

To thus enforce that Q is initialized to a correct proof, let H_2 be our variant of Ambainis's query Hamiltonian from Lemma 3.2, such that H_2 acts on registers Q and Q' (where for clarity $Q = (\mathbb{C}^2)^{\otimes 2^m - 1}$ (recall $m \in O(\log n)$) and $Q' = \mathcal{Y}$ from Lemma 3.2). Hence, our final Hamiltonian is $H = H_1 + H_2$, which is 5-local since H_1 is 5-local. Suppose without loss of generality that U's output qubit is W_1, which is set to $|0\rangle$ until the final time step, in which the correct output is copied to it. Then, set observable $A = (I + Z)/2$ such that A acts on qubit W_1. Set $a = 1 - 1/(L+1)$, and $b = 1 - 1/2L$ for L the number of gates in U. Fix $\eta \geq \max(\|H_2\|_\infty, 1)$ (such an η can be efficiently computed by applying the triangle inequality and summing the spectral norms of each term of H_2 individually). Set $\Delta = L^3\eta\gamma$ for γ a monotonically increasing polynomial function of L to be set as needed. Finally, set $\delta = 1/\Delta$. This completes the construction.

Correctness. Suppose Π is a YES instance. Then, by Lemma 3.2, the ground space of H_2 is the span of states of the form $|\hat{x}\rangle_Q \otimes |\phi\rangle_{Q'}$ where $\hat{x}$ is a correct query string encoded in unary. Fix an arbitrary such ground state $|\hat{x}\rangle_Q \otimes |\phi\rangle_{Q'}$. Note that setting Q to $\hat{x}$ in this manner causes U to accept with certainty. Consider the history state $|\psi_{\text{hist}}\rangle$ on registers W, C, Q, and Q' (Q and Q' together are the "proof register", and the contents of Q' are not accessed by U), which lies in the ground space of H_1. Since U can read but does not alter the contents of Q, the history state has the tensor product form $|\psi'_{\text{hist}}(x)\rangle_{W,C} \otimes |\hat{x}\rangle_Q \otimes |\phi\rangle_{Q'}$ for some $|\psi'_{\text{hist}}(x)\rangle_{W,C}$, i.e. the action of H_2 on the history state is unaffected. We conclude that $|\psi'_{\text{hist}}(x)\rangle_{W,C} \otimes |\hat{x}\rangle_Q \otimes |\phi\rangle_{Q'}$ is in the ground space of H. Moreover, since U accepts $\hat{x}$, the expectation of this state against A is $1 - 1/(L+1)$.

Conversely, suppose we have a NO instance Π, and consider any $|\psi\rangle$ satisfying $\langle\psi| H |\psi\rangle \leq \lambda(H) + \delta$. By Lemma 2.2, the smallest non-zero eigenvalue of ΔH_1 is at least $J = \pi^2\Delta/(64L^3) = \pi^2\eta\gamma/64$. Recalling that $\delta = 1/\Delta$, apply Corollary 2.3 to obtain that there exists a valid history state $|\psi'\rangle$ on W, C, Q, and Q' such that $|\langle\psi|\psi'\rangle|^2 \geq 1 - O(\gamma^{-2}L^{-6})$, which by Equation (2) implies

$$\| |\psi\rangle\langle\psi| - |\psi'\rangle\langle\psi'| \|_{\text{tr}} \leq \frac{c}{\gamma L^3} \tag{4}$$

for some constant $c > 0$. By definition, such a history state $|\psi'\rangle$ simulates U given "quantum proof" $|\phi\rangle_{Q,Q'}$ in registers Q and Q', i.e. $|\psi'\rangle = \sum_t U_t \cdots U_1 |0 \cdots 0\rangle_W |t\rangle_C |\phi\rangle_{Q,Q'}$. By

Equation (4) and the Hölder inequality, $|\mathrm{Tr}(H|\psi\rangle\langle\psi|) - \mathrm{Tr}(H|\psi'\rangle\langle\psi'|)| \leq \frac{c}{\gamma L^3}\|H\|_\infty =: \gamma'$. Thus, $\langle\psi'|H|\psi'\rangle \leq \lambda(H) + (\delta+\gamma')$.

We now analyze the structure of $|\phi\rangle_{Q,Q'}$. By Lemma 3.2, the ground space $\mathcal{G}$ of H_2 is contained in the span of states of the form $|\hat{x}\rangle_Q \otimes |\phi'\rangle_{Q'}$ where $\hat{x}$ is a correct query string encoded in unary. Since the ground spaces of H_1 and H_2 have non-empty intersection, i.e. history states acting on "quantum proofs" from $\mathcal{G}$ (which lie in the null space of H_1 and obtain energy $\lambda(H_2)$ against H_2), we know $\lambda(H) = \lambda(H_2)$. Thus, since $H_1 \succeq 0$,

$$\langle\psi'|H_2|\psi'\rangle \leq \langle\psi'|H|\psi'\rangle \leq \lambda(H_2) + (\delta+\gamma'). \tag{5}$$

Write $|\phi\rangle = \alpha|\phi_1\rangle + \beta|\phi_2\rangle$ for $|\phi_1\rangle \in \mathrm{Span}\{\ |\hat{x}\rangle_Q \otimes |\phi'\rangle_{Q'} \mid \text{correct query string } x\}$ and $|\phi_2\rangle \in \mathrm{Span}\{\ |\hat{x}\rangle_Q \otimes |\phi'\rangle_{Q'} \mid \text{incorrect query string } x\}$ ($|\phi_1\rangle$, $|\phi_2\rangle$ normalized), $\alpha, \beta \in \mathbb{C}, |\alpha|^2 + |\beta|^2 = 1$. Since any history state $|\psi'\rangle$, for any amplitudes α_x and unit vectors $|\phi'_x\rangle$, has form $\sum_{t,x}\alpha_x U_t\cdots U_1|0\cdots0\rangle_W|t\rangle_C|\hat{x}\rangle_Q|\phi'_x\rangle_{Q'} = \sum_x \alpha_x|\psi'_{\mathrm{hist}}(x)\rangle_{W,C}|\hat{x}\rangle_Q|\phi'_x\rangle_{Q'}$ (i.e. for any fixed x, $|\hat{x}\rangle_Q$ is not altered), and since H_2 is block-diagonal with respect to strings in Q, by Equation (5) and Lemma 3.2 we have

$$\begin{aligned}\lambda(H_2) + (\delta+\gamma') &\geq \langle\psi'|H_2|\psi'\rangle = |\alpha|^2\langle\phi_1|H_2|\phi_1\rangle + |\beta|^2\langle\phi_2|H_2|\phi_2\rangle \\ &\geq |\alpha|^2\lambda(H_2) + |\beta|^2\left(\lambda(H_2) + \frac{\epsilon}{4^m}\right),\end{aligned}$$

which implies $|\beta|^2 \leq 4^m(\delta+\gamma')/\epsilon$. Thus, defining $|\psi''\rangle$ as the history state for "proof" $|\phi_1\rangle_{Q,Q'}$, we have that $\||\psi\rangle\langle\psi| - |\psi''\rangle\langle\psi''|\|_{\mathrm{tr}}$ is at most

$$\||\psi\rangle\langle\psi| - |\psi'\rangle\langle\psi'|\|_{\mathrm{tr}} + \||\phi\rangle\langle\phi| - |\phi_1\rangle\langle\phi_1|\|_{\mathrm{tr}} \leq \frac{c}{\gamma L^3} + 2\sqrt{\frac{4^m(\delta+\gamma')}{\epsilon}}, \tag{6}$$

which follows from the triangle inequality and the structure of the history state. Observe now that increasing γ by a polynomial factor decreases $\delta+\gamma'$ by a polynomial factor. Thus, set γ as a large enough polynomial in L such that

$$\frac{c}{\gamma L^3} + 2\sqrt{\frac{4^m(\delta+\gamma')}{\epsilon}} \leq \frac{1}{2L}. \tag{7}$$

Since U rejects any correct query string (with certainty) in the NO case, and since $|\psi''\rangle$ is a valid history state whose Q register is a superposition over correct query strings (all of which must lead to reject), we conclude that $\langle\psi''|A|\psi''\rangle = 1$. Moreover, we have that $|\mathrm{Tr}(A|\psi\rangle\langle\psi|) - \mathrm{Tr}(A|\psi''\rangle\langle\psi''|)| \leq \|A\|_\infty\||\psi\rangle\langle\psi| - |\psi''\rangle\langle\psi''|\|_{\mathrm{tr}} \leq \frac{1}{2L}$, where the first inequality follows from Hölder's inequality, and the second by Equations (6) and (7). We conclude that $\langle\psi|A|\psi\rangle \geq 1 - 1/(2L)$, completing the proof. ◀

5 $\mathsf{P}^{\mathsf{QMA[log]}}$ is in PP

We now prove Theorem 1.3. Our approach is to develop a variant of the hierarchical voting scheme used in the proof of $\mathrm{P}^{\mathrm{NP[log]}} \subseteq \mathrm{PP}$ [4] which uses the strong error reduction technique of Marriott and Watrous [23]. We also require a more involved analysis than present in [4], since QMA is a class of promise problems, not decision problems.

Proof of Theorem 1.3. Let Π be a P machine which makes $m = c\log n$ queries to an oracle for 2-LH, for $c \in O(1)$ and n the input size. . Without loss of generality, we assume all queries involve Hamiltonians on M qubits (M some fixed polynomial in n). Define $q := (M+2)m$. We give a PQP computation simulating Π; since PQP = PP [27], this yields the claim. Let V denote the verification circuit for 2-LH. The PQP computation is (intuition to follow):

1. For i from 1 to m:
   a. Prepare $\rho = I/2^M \in \mathrm{D}\left((\mathbb{C}^2)^{\otimes M}\right)$.
   b. Run V on the ith query Hamiltonian $H_{y_i}^{i,y_1\cdots y_{i-1}}$ (see Equation (3)) and proof ρ, and measure the output qubit in the standard basis. Set bit y_i to the result.
2. Let $y = y_1 \cdots y_m$ be the concatenation of bits set in Step 1(b).
3. For i from 1 to $n^c - 1$:
   a. If $|y| < i$, then with probability $1 - 2^{-q}$, set $y = \#$, and with probability 2^{-q}, leave y unchanged.
4. If $y = \#$, output a bit in $\{0,1\}$ uniformly at random. Else, run Π on query string y and output Π's answer.

Intuition. In Step 1, one tries to determine the correct answer to query i by guessing a satisfying quantum proof for verifier V. Suppose for the moment that V has zero error, i.e. has completeness 1 and soundness 0, and that Π only makes valid queries. Then, if Step 1(b) returns $y_i = 1$, one knows with certainty that the query answer should be 1. And, if the correct answer to query i is 0, then Step 1(b) returns $y_i = 0$ with certainty. Thus, analogous to the classical case of an NP oracle (as done in [4]), it follows that the lexicographically *largest* query string y^* obtainable by this procedure must be the (unique) correct query string (note that $y^* \neq 1^m$ necessarily[5]). Thus, ideally one wishes to obtain y^*, simulate Π on y^*, and output the result. To this end, Step 3 ensures that among all values of $y \neq \#$, y^* is more likely to occur than all other $y \neq y^*$ combined. We now make this intuition rigorous (including in particular the general case where V is not zero-error and Π makes invalid queries).

Correctness. To analyze correctness of our PQP computation, it will be helpful to refine our partition of the set of query strings $\{0,1\}^m$ into three sets:

- **(Correct query strings)** Let $A \subseteq \{0,1\}^m$ denote the set of query strings which correspond to correctly answering each of the m queries. Note we may have $|A| > 1$ if invalid queries are made.
- **(Incorrect query strings)** Let $B \subseteq \{0,1\}^m \setminus A$ denote the set of query strings such that for any $y \in B$, all bits of y which encode an incorrect query answer are set to 0 (whereas the correct query answer would have been 1, i.e. we failed to "guess" a good proof for this query in Step 1).
- **(Strongly incorrect query strings)** Let $C = \{0,1\}^m \setminus (A \cup B)$ denote the set of query strings such that for any $y \in C$, at least one bit corresponding to an incorrect query answer is set to 1 (whereas the correct query answer would have been 0). Such an error can only arise due to the bounded-error of our QMA verifier in Step 1(b).

Let Y be a random variable corresponding to the query string y obtained at the end of Step 3. To show correctness, we claim that it suffices to show that $\Delta := \Pr[Y \in A] - \Pr[Y \in B \cup C] > 0$. To see this, let p_1, p_2, and p_3 denote the probability that after Step 3, $y = \#$, $y \in A$, and $y \in B \cup C$, respectively. Then, $p_1 + p_2 + p_3 = 1$, and let $p_2 - p_3 = \Delta > 0$. Suppose now that the input to Π is a YES instance. Then, our protocol outputs 1 with probability at least $\frac{p_1}{2} + p_2 = \frac{1-p_2-p_3}{2} + p_2 = \frac{1+\Delta}{2} > \frac{1}{2}$. If the input is a NO instance, the protocol outputs

[5] Under the assumptions that V has zero error and Π makes only valid queries, $y^* = 1^m$ can only be obtained by this procedure if all queries are for YES instances of 2-LH. If, on the other hand, query i is a NO query, then a correct proof cannot be guessed (since it does not exist), and so $y_i^* = 0$ necessarily.

1 with probability $\leq \frac{p_1}{2} + p_3 = \frac{1-\Delta}{2} < \frac{1}{2}$. We hence have a PQP computation, as desired. We thus now show that $\Delta > 0$.

To ease the presentation, we begin by making two assumptions (to be removed later): (i) V is zero-error and (ii) Π makes only valid queries. In this case, assumption (i) implies $C = \emptyset$ (i.e. all incorrect query strings belong to B), and (ii) implies A is a singleton (i.e. there is a unique correct query string y^*). Thus, here $\Delta = \Pr[Y \in A] - \Pr[Y \in B]$.

To begin, note that for any $y \in \{0,1\}^m$, we have

$$\Pr[Y = y] = \Pr[y \text{ chosen in Step 2 }] \cdot \left(\frac{1}{2^q}\right)^{(n^c-1)-|y|}, \tag{8}$$

where $|y|$ denotes the non-negative integer represented by string y. Let $\mathrm{HW}(x)$ denote the Hamming weight of $x \in \{0,1\}^m$. Since each query corresponds to a verifier on M proof qubits, we have for (the unique) $y^* \in A$ that

$$\Pr[y^* \text{ chosen in Step 2 }] \geq 2^{-M \cdot \mathrm{HW}(y^*)} \geq 2^{-Mm} \tag{9}$$

(recall from Section 2 that setting $\rho = I/2^M$ simulates "guessing" a correct proof with probability at least $1/2^M$). It follows by Equations (8) and (9) that

$$\begin{aligned} \Delta &\geq \left(\frac{1}{2^q}\right)^{(n^c-1)-|y^*|} \left[\frac{1}{2^{Mm}} - \sum_{y \in B} \left(\frac{1}{2^q}\right)^{|y^*|-|y|}\right] \\ &\geq \left(\frac{1}{2^q}\right)^{(n^c-1)-|y^*|} \left[\frac{1}{2^{Mm}} - (2^m)\left(\frac{1}{2^q}\right)\right] \geq \left(\frac{1}{2^q}\right)^{(n^c-1)} \frac{1}{2^{Mm}} \left[1 - \frac{1}{2^m}\right], \end{aligned} \tag{10}$$

where the first inequality follows since $\Pr[y \text{ chosen in Step 2 }] \leq 1$, the second since $y \in B$ if and only if $|y| < |y^*|$, and the third since $q = (M+2)m$. Thus, $\Delta > 0$ as desired.

Removing assumption (i). We now remove the assumption that V is zero error. In this case, A is still a singleton; let $y^* \in A$. We can now also have strongly incorrect query strings, i.e. $C \neq \emptyset$ necessarily. Assume without loss of generality that V acts on M proof qubits, and by strong error reduction [23] has completeness $c := 1 - 2^{-p(n)}$ and soundness $s := 2^{-p(n)}$, for p a polynomial to be chosen as needed. Then, since V can err, Equation (9) becomes

$$\begin{aligned} \Pr[y^* \text{ chosen in Step 2 }] &\geq \left(\frac{c}{2^M}\right)^{\mathrm{HW}(y^*)} (1-s)^{m-\mathrm{HW}(y^*)} = \frac{1}{2^M}^{\mathrm{HW}(y^*)} e^{m \ln(1-\frac{1}{2^p})} \\ &\geq \frac{1}{2^{Mm}} \left(1 - \frac{m}{2^p - 1}\right), \end{aligned} \tag{11}$$

where the equality follows by the definitions of c and s, and the second inequality by applying the Maclaurin series expansion of $\ln(1+x)$ for $|x| < 1$ and the fact that $e^t \geq 1 + t$ for all $t \in \mathbb{R}$. Thus, the analysis of Equation (10) yields that

$$\Pr[Y \in A] - \Pr[Y \in B] \geq \left(\frac{1}{2^q}\right)^{(n^c-1)} \frac{1}{2^{Mm}} \left[1 - \frac{1}{2^m} - \frac{m}{2^p - 1}\right], \tag{12}$$

i.e. the additive error introduced when assumption (i) is dropped scales as $\approx 2^{-p}$. Crucially, Equation (12) holds for all $y \in B$ even with assumption (i) dropped since the analysis of Equation (10) used only the trivial bound $\Pr[y \text{ chosen in Step 2 }] \leq 1$ for any $y \in B$.

Next, we upper bound the probability of obtaining $y \in C$ in Step 2. For any fixed $y \in C$, suppose the first bit on which y and y^* disagree is bit j. Then, bits j of y and y^* must be

1 and 0, respectively. This means 0 is the correct answer for query j. By the soundness property of V, the probability of obtaining 1 on query j (and hence that of obtaining y in Step 2) is at most 2^{-p}. Thus,

$$\Delta \geq \left(\frac{1}{2^q}\right)^{(n^c-1)} \frac{1}{2^{Mm}} \left[1 - \frac{1}{2^m} - \frac{m}{2^p - 1}\right] - \frac{2^m}{2^p}. \tag{13}$$

We conclude that setting p to a sufficiently large fixed polynomial ensures $\Delta > 0$, as desired.

Removing assumption (ii). We now remove the assumption that Π only makes valid queries, which is the most involved step. Here, A is no longer necessarily a singleton. The naive approach would be to let y^* denote the *lexicographically largest* string in A, and attempt to run a similar analysis as before. Unfortunately, this no longer necessarily works for the following reason. For any invalid query i, we do not have strong bounds on the probability that V accepts in Step 1(b); in principle, this value can lie in the range $(2^{-p}, 1 - 2^{-p})$. Thus, running the previous analysis with the lexicographically largest $y^* \in A$ may cause Equation (13) to yield a negative quantity. We hence require a more delicate analysis.

We begin by showing the following lower bound.

▶ **Lemma 5.1.** *Define* $\Delta' := \Pr[Y \in A] - \Pr[Y \in B]$. *Then,*

$$\Delta' \geq \left(\frac{1}{2^q}\right)^{(n^c-1)} \frac{1}{2^{Mm}} \left[1 - \frac{1}{2^m} - \frac{m}{2^p - 1}\right].$$

Proof of Lemma 5.1. For any string $y \in \{0,1\}^m$, let $I_y \subseteq \{1, \ldots, m\}$ denote the indices of all bits of y set by invalid queries. We call each such $i \in I_y$ a *divergence point*. Let $p_{y,i}$ denote the probability that (invalid) query i (defined given answers to queries 1 through $i-1$) outputs bit y_i, i.e. $p_{y,i}$ denotes the probability that at divergence point i, we go in the direction of bit y_i. We define the *divergence probability* of $y \in \{0,1\}^m$ as $p_y = \Pi_{i \in I_y} p_{y,i}$, i.e. p_y is the probability of answering all invalid queries as y did.

The proof now proceeds by giving an iterative process, $\Gamma(i)$, where $1 \leq i \leq |A|$ denotes the iteration number. Each iteration defines a 3-tuple $(y^*_{i-1}, y^*_i, B_{y^*_i}) \in \{0,1\}^m \times \{0,1\}^m \times \mathcal{P}(B)$, where $\mathcal{P}(X)$ denotes the power set of set X. Set $\Delta'_i := \Pr[Y \in \{y^*_1, \ldots, y^*_i\}] - \Pr[Y \in B_{y^*_1} \cup \cdots \cup B_{y^*_i}]$, where it will be the case that $\{B_{y^*_i}\}_{i=1}^{|A|}$ is a partition of B. Thus, we have $\Delta' \geq \Delta'_{|A|}$, implying that a lower bound on $\Delta'_{|A|}$ suffices to prove our claim. We hence prove via induction that for all $1 \leq i \leq |A|$, $\Delta'_i \geq \left(\frac{1}{2^q}\right)^{(n^c-1)} \frac{1}{2^{Mm}} \left[1 - \frac{1}{2^m} - \frac{m}{2^p-1}\right]$. The definition of process $\Gamma(i)$ is integrated into the induction proof below.

Base case (i=1). In this case y^*_0 is undefined. Set y^*_1 to any string in A with divergence probability at least

$$p^*_1 = \prod_{i \in I_{y^*_1}} p_{y^*_1, i} \geq 2^{-\left|I_{y^*_1}\right|}. \tag{14}$$

Such a string must exist, since at each divergence point i, at least one of the outcomes in $\{0,1\}$ occurs with probability at least $1/2$. (Note: Queries are not being made to a QMA oracle here, but to a QMA verifier V with a maximally mixed proof as in Step 1(a). Whereas in the former case the output of the oracle on an invalid query does not have to consistently output a value with any particular probability, in the latter case, there is some fixed probability p with which V outputs 1 each time it is run on a fixed proof.) Finally, define $B_{y^*_1} := \{y \in B \mid |y| < |y^*_1|\}$ (recall $|y|$ is the non-negative integer with binary encoding y).

Let k_* denote the number of divergence points of y_1^* (i.e. $k_* = \left|I_{y_1^*}\right|$), and k_0 (k_1) the number of zeroes (ones) of y_1^* arising from valid queries. Thus, $k_* + k_0 + k_1 = m$. Then, Equation (11) becomes

$$\begin{aligned}\Pr[y_1^* \text{ in Step 2 }] &\geq \left(\frac{c}{2^M}\right)^{k_1} (1-s)^{k_0}\, p_1^* \geq \left(\frac{1}{2^M}\right)^{k_1} \left(\frac{1}{2}\right)^{k_*} \left(1 - \frac{m-k_*}{2^p-1}\right) \\ &\geq \frac{1}{2^{Mm}} \left(1 - \frac{m}{2^p-1}\right), \end{aligned} \tag{15}$$

where the second inequality follows from Equation (14), and the third since $k_* \geq 0$ and $k_1 + k_* \leq m$. Thus, Δ_1' is lower bounded by the expression in Equation (12) via an analogous analysis for y_1^* and $B_{y_1^*}$.

Inductive step. Assume the claim holds for $1 \leq i-1 < |A|$. We show it holds for i. Let y_{i-1}^* be the choice of y^* in the previous iteration $i-1$ of our process. Define $A_{y_i^*} := \{\, y \in A \mid |y| > \left|y_{i-1}^*\right| \,\}$. Partition $A_{y_i^*}$ into sets S_k for $k \in [m]$, such that S_k is the subset of strings in $A_{y_i^*}$ which agrees with y_{i-1}^* on the first $k-1$ bits, but disagrees on bit k. Note that if $S_k \neq \emptyset$, then bit k of y_{i-1}^* is 0 and bit k of any string in S_k is 1. For each $S_k \neq \emptyset$, choose arbitrary representative $z_k \in S_k$, and define *bounded* divergence probability $q_i(k) := \prod_{t \in I_{z_k}^{\leq k}} p_{z_k,t}$ where $I_{z_k}^{\leq k} := \{\, t \in I_{z_k} \mid t \leq k \,\}$. Note that $q_i(k) > 0$ (since $S_k \neq \emptyset$). Else if $S_k = \emptyset$, set $q_i(k) = 0$. Let q_i^* be the max such bounded divergence probability:

$$q_i^* = \max_{k \in [m]} q_i(k) \qquad \text{and} \qquad k_i^* = \arg\max_{k \in [m]} q_i(k). \tag{16}$$

Let y_i^* be the lexicographically largest query string in $S_{k_i^*}$ with divergence probability p_i^* s.t.:

$$p_i^* \geq q_i^* \cdot 2^{-\left|I_{y_i^*}\right| + \left|I_{y_i^*}^{\leq k_i^*}\right|}. \tag{17}$$

That such a $y_i^* \in S_{k_i^*}$ exists follows from an argument similar to Equation (14): By definition, q_i^* denotes the bounded divergence probability for all invalid queries up to and including query k_i^*, and the term exponential in $\left(-\left|I_{y_i^*}\right| + \left|I_{y_i^*}^{\leq k_i^*}\right|\right)$ is obtained by greedily choosing, for all invalid queries of y_i^* *after* query k_i^*, the outcome which occurs with probability at least $1/2$. Set $B_{y_i^*} := \{\, y \in B \mid \left|y_{i-1}^*\right| < |y| < |y_i^*| \,\}$. The following is proved in the full version.

▶ **Lemma 5.2.** *For any $y \in B_{y_i^*}$, $\Pr[y \text{ chosen in Step 2}] \leq q_i^*$.*

To continue with the inductive step, again consider k_*, k_0, and k_1, now corresponding to y_i^*. Then, an argument similar to Equation (15) says $\Pr[y_i^* \text{ chosen in Step 2 }]$ is at least

$$\begin{aligned}\left(\frac{c}{2^M}\right)^{k_1} (1-s)^{k_0}\, p_i^* &\geq \left(\frac{1}{2^M}\right)^{k_1} \left(1 - \frac{m-k_*}{2^p-1}\right) q_i^* \left(\frac{1}{2}\right)^{\left|I_{y_i^*}\right| - \left|I_{y_i^*}^{\leq k_i^*}\right|} \\ &\geq \frac{q_i^*}{2^{Mm}} \left(1 - \frac{m}{2^p-1}\right), \end{aligned} \tag{18}$$

where the first inequality follows from Equation (17), and the second since $\left|I_{y_i^*}\right| - \left|I_{y_i^*}^{\leq k_i^*}\right| \leq k_*$. Now, define $\zeta_i := \Pr[Y = y_i^*] - \Pr[Y \in B_{y_i^*}]$. Applying the argument of Equation (10) yields $\zeta_i \geq \left(\frac{1}{2^q}\right)^{(n^c-1)-|y_i^*|} \left[\frac{q_i^*}{2^{Mm}}\left(1 - \frac{m}{2^p-1}\right) - q_i^* \sum_{y \in B_{y_i^*}} \left(\frac{1}{2^q}\right)^{|y_i^*|-|y|}\right]$, where the first q_i^* is due

to Equation (18), and the second q_i^* to Lemma 5.2. Thus, similar to Equation (12), $\zeta_i \geq \left(\frac{1}{2^q}\right)^{(n^c-1)} \frac{q_i^*}{2^{Mm}} \left[1 - \frac{1}{2^m} - \frac{m}{2^{p-1}}\right] > 0$. Observing the recurrence that for all i, $\Delta_i' \geq \Delta_{i-1}' + \zeta_i$, unrolling this recurrence yields $\Delta_i' \geq \Delta_1$, which by the base case yields the claim. ◀

We require one last lemma (proof in the full version).

▶ **Lemma 5.3.** $\Pr(Y \in C) \leq \frac{2^m}{2^p}$.

Finally, combining Lemmas 5.1 and 5.3 yields that $\Pr[Y \in A] - \Pr[Y \in B \cup C]$ is lower bounded by $\Pr[Y \in A] - \Pr[Y \in B] - \Pr[Y \in C] \geq \left(\frac{1}{2^q}\right)^{(n^c-1)} \frac{1}{2^{Mm}} \left[1 - \frac{1}{2^m} - \frac{m}{2^p}\right] - \frac{2^m}{2^p}$. For sufficiently large fixed p, this quantity is strictly positive, yielding Theorem 1.3. ◀

6 Estimating spectral gaps

We now prove Theorem 1.4 on SPECTRAL-GAP. UQMA is defined in Appendix A.

▶ **Definition 6.1** (SPECTRAL-GAP(H, ϵ) (Ambainis [3]))**.** Given a Hamiltonian H and a real number $\alpha \geq n^{-c}$ for n the number of qubits H acts on and $c > 0$ some constant, decide:

- If $\lambda_2 - \lambda_1 \leq \alpha$, output YES.
- If $\lambda_2 - \lambda_1 \geq 2\alpha$, output NO.

where λ_2 and λ_1 denote the second and first smallest eigenvalues of H, respectively.

For clarity, if the ground space of H is degenerate, then we define its spectral gap as 0.

We now discuss Theorem 1.4. Previously, Ambainis [3] showed that SPECTRAL-GAP $\in \mathrm{P}^{\mathrm{QMA[log]}}$, and gave a claimed proof that SPECTRAL-GAP is $\mathrm{P}^{\mathrm{UQMA[log]}}$-hard for $O(\log)$-local Hamiltonians under mapping reductions. ($\mathrm{P}^{\mathrm{UQMA[log]}}$ is defined as $\mathrm{P}^{\mathrm{QMA[log]}}$, except with a UQMA oracle in place of a QMA oracle.) As discussed in Section 1, however, Ambainis' proof of the latter result does not hold if the $\mathrm{P}^{\mathrm{UQMA[log]}}$ machine makes invalid queries (which in general is the case). Here, we build on Ambainis' approach [3] to show $\mathrm{P}^{\mathrm{UQMA[log]}}$-hardness of SPECTRAL-GAP under Turing reductions even when invalid queries are allowed, and we also improve the hardness to apply to $O(1)$-local Hamiltonians.

We begin by showing the following modified version of Lemma 3.2 tailored to UQMA. In contrast to Lemma 3.2, the lemma below only proves the *existence* of a Hamiltonian H; it does not give an *efficient* procedure for computing it. The proof is in the full version; roughly, it replaces invalid queries with "dummy" NO queries to obtain the desired spectral gap. The reason why the mapping is not efficient is that generally a polynomial-time machine alone cannot identify such invalid queries.

▶ **Lemma 6.2.** *For any $x \in \{0,1\}^m$, let $\hat{x}$ denote its unary encoding. Then, for any $\mathrm{P}^{\mathrm{UQMA[log]}}$ circuit U acting on n bits and making m queries to a* UQMA *oracle, there exists a 4-local Hamiltonian H acting on space $(\mathbb{C}^2)^{\otimes 2^m - 1} \otimes \mathcal{Y}$ such that there exists a correct query string $x = x_1 \cdots x_m$ such that:*

1. *The unique ground state of H lies in subspace $|\hat{x}\rangle\langle\hat{x}| \otimes \mathcal{Y}$.*
2. *The spectral gap of H is at least $(\epsilon - \delta)/4^m$ for inverse polynomial ϵ, δ with $\epsilon - \delta \geq 1/\mathrm{poly}(n)$.*
3. *For all strings $x' \in \{0,1\}^m$, H acts invariantly on subspace $|\hat{x}'\rangle\langle\hat{x}'| \otimes \mathcal{Y}$.*

Proof sketch of Theorem 1.4. The key idea is to show how to use an *oracle* for SPECTRAL-GAP polynomially many times to efficiently identify invalid queries, and hence efficiently compute H in Lemma 6.2 given U. (It is these *multiple* uses of the oracle which yield a Turing reduction, rather than a many-one reduction.) Roughly, this is done by using

the SPECTRAL-GAP oracle in conjunction with binary search to estimate the spectral gap of specific Hamiltonian terms in Ambainis's original construction of [3]. Some care is required here: The naive approach, which does not work, would be to apply this spectral gap estimation technique to each 2-local Hamiltonian $H_{\mathcal{Y}_i}^{i,y_1\cdots y_{i-1}}$ corresponding to each query made by U. Rather, the terms we apply this technique to exploit the structure of Ambainis's construction. Finally, with H in hand, we apply Ambainis's [3] original construction to obtain the desired result. The full proof is given in the full version of this article. ◀

Acknowledgements. We thank Xiaodi Wu for stimulating discussions which helped motivate this project, including suggesting to think about two-point correlation functions (which arose via discussions with Aram Harrow, whom we also thank). We also thank Andris Ambainis and Norbert Schuch for helpful discussions, and remark they independently conceived of some of the ideas behind Lemma 3.2 and Theorem 1.1, respectively (private communication).

References

1 D. Aharonov, M. Ben-Or, F. Brandão, and O. Sattath. The pursuit for uniqueness: Extending Valiant-Vazirani theorem to the probabilistic and quantum settings. Available at arXiv.org e-Print quant-ph/0810.4840v1, 2008.

2 D. Aharonov and T. Naveh. Quantum NP - A survey. Available at arXiv.org e-Print quant-ph/0210077v1, 2002.

3 A. Ambainis. On physical problems that are slightly more difficult than QMA. In *Proceedings of 29th IEEE Conference on Computational Complexity (CCC 2014)*, pages 32–43, 2014.

4 R. Beigel, L. A. Hemachandra, and G. Wechsung. On the power of probabilistic polynomial time: $\mathrm{P}^{\mathrm{NP}[\log]} \subseteq \mathrm{PP}$. In *Proceedings of the 4th IEEE Conference on Structure in Complexity Theory*, pages 225–227, 1989.

5 A. D. Bookatz. QMA-complete problems. *Quantum Information & Computation*, 14(5–6), 2014.

6 B. Brown, S. Flammia, and N. Schuch. Computational difficulty of computing the density of states. *Physical Review Letters*, 104:040501, 2011.

7 A. Chailloux and O. Sattath. The complexity of the separable Hamiltonian problem. In *Proceedings of 27th IEEE Conference on Computational Complexity (CCC 2012)*, pages 32–41, 2012.

8 S. Cook. The complexity of theorem proving procedures. In *Proceedings of the 3rd ACM Symposium on Theory of Computing (STOC 1972)*, pages 151–158, 1972.

9 T. Cubitt and A. Montanaro. Complexity classification of local hamiltonian problems. Available at arXiv.org e-Print quant-ph/1311.3161, 2013.

10 R. Feynman. Quantum mechanical computers. *Optics News*, 11:11, 1985.

11 S. Gharibian. *Approximation, proof systems, and correlations in a quantum world.* PhD thesis, University of Waterloo, 2013. Available at arXiv.org e-Print quant-ph/1301.2632.

12 S. Gharibian and J. Kempe. Hardness of approximation for quantum problems. In *Proceedings of 39th International Colloquium on Automata, Languages and Programming (ICALP 2012)*, pages 387–398, 2012.

13 S. Gharibian and J. Sikora. Ground state connectivity of local Hamiltonians. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015)*, pages 617–628, 2015.

14 Sevag Gharibian, Yichen Huang, Zeph Landau, and Seung Woo Shin. Quantum hamiltonian complexity. *Foundations and Trends in Theoretical Computer Science*, 10(3):159–282, 2015. doi:10.1561/0400000066.

15 J. Gill. Computational complexity of probabilistic Turing machines. *SIAM Journal on Computing*, 6(4):675–695, 1977.

16 O. Goldreich. On promise problems: A survey. *Theoretical Computer Science*, 3895:254–290, 2006.

17 R. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. New York: Plenum, 1972.

18 J. Kempe, A. Kitaev, and O. Regev. The complexity of the local Hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006.

19 A. Kitaev, A. Shen, and M. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, 2002.

20 A. Kitaev and J. Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *Proceedings of the 32nd ACM Symposium on Theory of Computing (STOC 2000)*, pages 608–617, 2000.

21 L. Levin. Universal search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.

22 Y.-K. Liu. Consistency of local density matrices is QMA-complete. In *Lecture Notes in Computer Science*, volume 4110, pages 438–449, 2006.

23 C. Marriott and J. Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005.

24 S. Piddock and A. Montanaro. The complexity of antiferromagnetic interactions and 2d lattices. Available at arXiv.org e-Print quantph/1506.04014, 2015.

25 Y. Shi and S. Zhang. Note on quantum counting classes. URL: `http://www.cse.cuhk.edu.hk/~syzhang/papers/SharpBQP.pdf`.

26 M. Vyalyi. QMA=PP implies that PP contains PH. *Electronic Colloquium on Computational Complexity*, 2003.

27 J. Watrous. *Encyclopedia of Complexity and System Science*, chapter Quantum Computational Complexity. Springer, 2009.

A Additional definitions

▶ **Definition 1.1** (Unique QMA (UQMA) (Aharonov *et al.* [1]))**.** We say a promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in Unique QMA if and only if there exist polynomials p, q and a polynomial-time uniform family of quantum circuits $\{ Q_n \}$, where Q_n takes as input a string $x \in \Sigma^*$ with $|x| = n$, a quantum proof $|y\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$, and $q(n)$ ancilla qubits in state $|0\rangle^{\otimes q(n)}$, such that:

- (Completeness) If $x \in A_{\text{yes}}$, then there exists a proof $|y\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$ such that Q_n accepts $(x, |y\rangle)$ with probability at least $2/3$, and for all $|\hat{y}\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$ orthogonal to $|y\rangle$, Q_n accepts $(x, |\hat{y}\rangle)$ with probability at most $1/3$.
- (Soundness) If $x \in A_{\text{no}}$, then for all proofs $|y\rangle \in (\mathbb{C}^2)^{\otimes p(n)}$, Q_n accepts $(x, |y\rangle)$ with probability at most $1/3$.