Supplementary Material

Quantum Complexity Classes

We start with the relevant definitions. Let x be a binary string. Then, we denote by the verifier $V \equiv V_x$ a quantum circuit of length $T = \text{poly}(|x|)$, $U = U_T \cdots U_1$ (with local gates U_t) acting on $m = \text{poly}(|x|)$ qubits, which is generated uniformly from x. The verifier takes an $n = \text{poly}(|x|)$ qubit quantum state $|\psi\rangle_I$ as an input (we will express everything in terms of n instead of $|x|$ in the following), together with $m - n$ initialized ancillas, $|\mathbf{0}\rangle_A \equiv |0 \cdots 0\rangle_A$, applies U, and finally measures the first qubit in the $\{|0\rangle_1, |1\rangle_1\}$ basis to return 1 ("proof accepted") or 0 ("proof rejected"). The acceptance probability for a proof $|\psi\rangle$ is then given by $p_{\text{acc}}(V(\psi)) := \langle\psi|\Omega|\psi\rangle$, with

$$\Omega = (\mathbb{1}_I \otimes \langle\mathbf{0}|_A) U^\dagger (|1\rangle\langle 1|_1 \otimes \mathbb{1}) U (\mathbb{1}_I \otimes |\mathbf{0}\rangle_A) \,. \tag{1}$$

Definition 1 (QMA)**.** *Let $0 \leq a, b \leq 1$ s.th. $a - b > \frac{1}{p(n)}$ for some polynomial $p(n)$. A language L is in $\mathsf{QMA}(a, b)$ if there exists a verifier s.th.*

$$\begin{aligned} x \in L &\Rightarrow \lambda_{\max}(\Omega) > a \\ x \notin L &\Rightarrow \lambda_{\max}(\Omega) < b \,. \end{aligned}$$

Definition 2 (#BQP)**.** *Let $0 \leq a, b \leq 1$ s.th. $a - b > 1/\text{poly}(n)$, and let Ω be a verifier map with no eigenvalues between a and b. Then, the class $\#\mathsf{BQP}(a, b)$ consists of all problems of the form "compute the dimension of the space spanned by all eigenvectors of Ω with eigenvalues $\geq a$" .*

An alternative definition for #BQP is the following:

Definition 3 (#BQP, alternate definition)**.** *Consider a verifier Ω with the property that there exist subspaces $\mathcal{A} \oplus \mathcal{R} = \mathbb{C}^N$ ($N = 2^n$) such that $\langle\psi|\Omega|\psi\rangle \geq a$ for all $|\psi\rangle \in \mathcal{A}$, and $\langle\psi|\Omega|\psi\rangle \leq b$ for all $|\psi\rangle \in \mathcal{R}$, where again $a - b > 1/\text{poly}(n)$. Then $\#\mathsf{BQP}(a, b)$ consists of all problems of the form "compute $\dim \mathcal{A}$".*

Note that $\dim \mathcal{A}$ is well-defined: Consider two decompositions $\mathbb{C}^N = \mathcal{A} \oplus \mathcal{R}$ and $\mathbb{C}^N = \mathcal{A}' \oplus \mathcal{R}'$. Without loss of generality, if we assume $\dim \mathcal{A} > \dim \mathcal{A}'$, it follows $\dim \mathcal{A} + \dim \mathcal{R}' > N$, and thus there exists a non-trivial $|\mu\rangle \in \mathcal{A} \cap \mathcal{R}'$, which contradicts the definition.

Theorem 4. *Definition 2 and Definition 3 are equivalent.*

Proof. To show that Definition 2 implies Definition 3, let $\mathcal{A}$ be spanned by the eigenvectors with eigenvalues $\geq a$. To show the converse, we use the minimax principle for eigenvalues [**?**], which states that the kth largest eigenvalue λ_k of a Hermitian operator Ω in an N-dimensional Hilbert space can be obtained from either of the equivalent optimizations

$$\begin{aligned} \lambda_k(\Omega) &= \max_{\mathcal{M}_k} \min_{|x\rangle \in \mathcal{M}_k} \langle x|\Omega|x\rangle & (2) \\ &= \min_{\mathcal{M}_{N-k+1}} \max_{|x\rangle \in \mathcal{M}_{N-k+1}} \langle x|\Omega|x\rangle \,, & (3) \end{aligned}$$

where $\mathcal{M}_k$ is a subspace of dimension k, and $|x\rangle$ is a unit vector. Now notice that Def. 3 implies that

$$\min_{|x\rangle \in \mathcal{A}} \langle x|\Omega|x\rangle \geq a \quad \text{and} \quad \max_{|x\rangle \in \mathcal{R}} \langle x|\Omega|x\rangle \leq b \,. \tag{4}$$

Next, consider the minimax theorem for $k = \dim \mathcal{A}$. From Eq. (2) we have

$$\begin{aligned} \lambda_{\dim \mathcal{A}} &= \max_{\mathcal{M}_{\dim \mathcal{A}}} \min_{|x\rangle \in \mathcal{M}_{\dim \mathcal{A}}} \langle x|\Omega|x\rangle \\ &\geq \min_{|x\rangle \in \mathcal{A}} \langle x|\Omega|x\rangle \\ &\geq a \,. \end{aligned}$$

Now consider the case that $k = \dim \mathcal{A} + 1$. From Eq. (3), using the fact that $N - (\dim \mathcal{A} + 1) + 1 = \dim \mathcal{R}$, we have

$$\begin{aligned} \lambda_{\dim \mathcal{A}+1} &= \min_{\mathcal{M}_{\dim \mathcal{R}}} \max_{|x\rangle \in \mathcal{M}_{\dim \mathcal{R}}} \langle x|\Omega|x\rangle \\ &\leq \max_{|x\rangle \in \mathcal{R}} \langle x|\Omega|x\rangle \\ &\leq b \,. \end{aligned}$$

Thus we have

$$\lambda_{\dim \mathcal{A}} \geq a > b \geq \lambda_{\dim \mathcal{A}+1} \,, \tag{5}$$

since $a - b \geq 1/\text{poly}(n)$. It follows that $\lambda_{\dim \mathcal{A}}$ is the smallest eigenvalue of Ω which is still larger than a, and therefore the span of the first $\dim \mathcal{A}$ eigenvectors of Ω is equal to the span of all eigenvectors with eigenvalue $\geq a$. The equivalence follows. □

The class $\#\mathsf{BQP}(a, b)$ inherits the useful property of strong error reduction from QMA: the thresholds (a, b) can be amplified to $(1 - 2^{-r}, 2^{-r})$, $r = \text{poly}(n)$ while keeping the size of the proof:

Theorem 5. *Let $\#\mathsf{BQP}_n(a, b)$ denote #BQP with an n qubit witness. Then $\#\mathsf{BQP}_n(a, b) \subset \#\mathsf{BQP}_n(1 - 2^{-r}, 2^{-r})$ for every $r \in \text{poly}(n)$.*

This follows directly from the strong amplification procedure presented in [**?**], which describes a procedure to amplify any verifier map Ω such that any eigenvalue above a (below b) is shifted above $1 - 2^{-r}$ (below 2^{-r}) at an overhead polynomial in r.

Using Thm. 5, we will write #BQP instead of $\#\mathsf{BQP}(a, b)$ from now on, where $a - b > \text{poly}(n)$, and a, b can be exponentially close to 1 and 0, respectively.

The Complexity of the Density of States

We now use the class #BQP to characterize the complexity of the density of states problem and the problem of counting the number of ground states of a local Hamiltonian. We start by defining these problems, as well as the notion of local Hamiltonian, and then show that both problems are #BQP-complete.

Definition 6 (k-local Hamiltonian)**.** *Given a set of* poly(n) *quantum spins each with dimension bounded by a constant, a Hamiltonian H for the system is said to be k-local if $H = \sum_i H_i$ is a sum of at most* poly(n) *Hermitian operators H_i, $\|H_i\| \leq 1$, each of which acts nontrivially on at most k spins.*

Note that k-local does not imply any geometric locality, only that each spin interacts with at most $k-1$ other spins for any given interaction term. However, we restrict ourselves to $k = O\big(\log(n)\big)$ so that each H_i can be specified by an efficient classical description.

Definition 7 (DOS)**.** *Let $E_2 - E_1 > 1/\text{poly}(n)$, $\Delta = (E_2 - E_1)/\text{poly}(n)$, and let $H = \sum_i H_i$ be a k-local Hamiltonian such that H has no eigenvalues in the intervals $[E_1 - \frac{\Delta}{2}, E_1 + \frac{\Delta}{2}]$ and $[E_2 - \frac{\Delta}{2}, E_2 + \frac{\Delta}{2}]$. Then, the problem* DOS *(density of states) is to compute the number of orthogonal eigenstates with eigenvalues in the interval $[E_1, E_2]$.*

Definition 8 (#LH)**.** *Let $E_0 \leq E_1 < E_2$, with $E_2 - E_1 > 1/\text{poly}(n)$, and let $H = \sum_i H_i$ be a k-local Hamiltonian s.th. $H \geq E_0$, and H has no eigenvalues between E_1 and E_2. Then, the problem $\#\text{LH} \equiv \#\text{LH}(E_1 - E_0)$ (sharp local Hamiltonian) is to compute the dimension of the eigenspace with eigenvalues $\leq E_1$.*

Note that #LH depends on the "energy splitting" $E_1 - E_0$ of the low-energy subspace. In particular, for $E_1 - E_0 = 0$, $\#\text{LH}(0)$ corresponds to computing the degeneracy of the ground state subspace. As we will see in what follows, the class $\#\text{LH}(\sigma)$ is the same for any splitting $\exp(-\text{poly}(n)) \leq \sigma \leq 1/\text{poly}(n)$.

We now show that #LH and DOS are both #BQP-complete. We do so by giving reductions from $\#\text{LH}(1/\text{poly}(n))$ to DOS, from DOS to #BQP, and from #BQP to $\#\text{LH}\big(\exp(-\text{poly}(n))\big)$; this will at the same time prove the claimed independence of $\#\text{LH}(\sigma)$ of the splitting σ.

Theorem 9. $\#\text{LH}(1/\text{poly}(n))$ *reduces to* DOS.

Proof. If we denote the parameters of the #LH problem by $\tilde{E}_0, \tilde{E}_1, \tilde{E}_2$, then we can simply relate them to the parameters E_1, E_2, Δ of a DOS problem by $\Delta = \tilde{E}_2 - \tilde{E}_1$, $E_1 = \tilde{E}_0 - \frac{1}{2}\Delta$ and $E_2 = \tilde{E}_1 + \frac{1}{2}\Delta$, and the result follows directly. □

Theorem 10. DOS *is contained in* #BQP.

Proof. We start with a k-local Hamiltonian H as in Def. 7. Now define a new Hamiltonian

$$H' := \nu(H^2 - (E_1 + E_2)H + E_1 E_2) \ .$$

H' is a $2k$-local Hamiltonian; here, $\nu = 1/\text{poly}(n)$ is chosen such that each term in H' is subnormalized. Any eigenvalue of H in the interval $[E_1 + \frac{\Delta}{2}; E_2 - \frac{\Delta}{2}]$ translates into an eigenvalue of H' which is below

$$A := -\nu\tfrac{\Delta}{2}(E_2 - E_1 - \tfrac{\Delta}{2}) \leq -1/\text{poly}(n) \ ,$$

whereas any eigenvalue outside $[E_1 - \frac{\Delta}{2}; E_2 + \frac{\Delta}{2}]$ translates into an eigenvalues of H' above

$$B := \nu\tfrac{\Delta}{2}(E_2 - E_1 + \tfrac{\Delta}{2}) \geq 1/\text{poly}(n) \ .$$

The original DOS problem now translates into counting the number of eigenstates of H' with negative energy, given a spectral gap in a $1/\text{poly}(n)$ sized interval $[A; B]$ around zero. We now use the circuit which was used in [**?**] to prove that log-local Hamiltonian is in QMA; it accepts any input state $|\psi\rangle$ with probability

$$p(|\psi\rangle) = \frac{1}{2} - \frac{\langle\psi|H'|\psi\rangle}{2m} \ ,$$

where $m = \text{poly}(n)$ is the number of terms in H'. [The idea is to randomly pick one term H'_i in the Hamiltonian and toss a coin with probability $(1 - \langle\psi|H'_i|\psi\rangle)/2$.] Computing the answer to the original DOS problem then translates to counting the number of states with acceptance probability $\geq a = \frac{1}{2} + \frac{A}{2m}$, given that there are no eigenstates between a and $b = \frac{1}{2} + \frac{B}{2m}$, and since $a - b = (A - B)/2m \geq 1/\text{poly}(n)$, this shows that this number can be computed in #BQP. □

Theorem 11. $\#\text{LH}\big(\exp(-\text{poly}(n))\big)$ *is* #BQP*-hard.*

Proof. To show #BQP-hardness of #LH, we need to start with an arbitrary QMA verifier circuit $U = U_T \dots U_1$ and construct a Hamiltonian with as many ground states as the circuit has accepting inputs. By amplification, we can assume that the acceptance and rejection thresholds for the verifier are $a = 1 - \epsilon$ and $b = \epsilon$, where we can choose $\epsilon = O(\exp(-p(n)))$ for any polynomial $p(n)$. As before, let $\mathcal{A}$ and $\mathcal{R}$ be the eigenspaces of Ω with eigenvalues $\geq a$ and $\leq b$, respectively. Define

$$U[\mathcal{R}] := \{U|\psi\rangle_I|\mathbf{0}\rangle_A : |\psi\rangle_I \in \mathcal{R}\} \tag{6}$$

and denote the projector onto this space by $\Pi_{U[\mathcal{R}]}$. Notice that for any state $|\chi\rangle \in U[\mathcal{R}]$, due to our rejection threshold $b = \epsilon$, we have

$$\langle\chi|\big(|1\rangle\langle 1|_1 \otimes \mathbb{1}\big)|\chi\rangle \leq \epsilon \, . \tag{7}$$

We now follow Kitaev's original construction to encode a QMA verifier circuit into a Hamiltonian which has a "proof history" as its ground state for any proof $|\phi\rangle_I \in \mathcal{A}$ [**?** **?**]. That is, the ground states of the Hamiltonian are given by

$$|\Phi\rangle = \sum_{t=0}^{T} U_t \dots U_1 |\phi\rangle_I |\mathbf{0}\rangle_A |t\rangle_T \tag{8}$$

where the third register is used as a "clock". The Hamiltonian has the form

$$H = H_{\text{init}} + \sum_{t=1}^{T} H_{\text{evol}}(t) + H_{\text{final}} \tag{9}$$

where

- $H_{\text{init}} = \mathbb{1}_I \otimes (\mathbb{1} - |\mathbf{0}\rangle\langle\mathbf{0}|_A) \otimes |0\rangle\langle 0|_T$ checks that the ancilla is property initialized, penalizing states without properly initialized ancillas;
- $H_{\text{evol}}(t) = -\frac{1}{2} U_t \otimes |t\rangle\langle t-1|_T - \frac{1}{2} U_t^\dagger \otimes |t-1\rangle\langle t|_T + \frac{1}{2}\mathbb{1} \otimes |t\rangle\langle t|_T + \frac{1}{2}\mathbb{1} \otimes |t-1\rangle\langle t-1|_T$
  checks that the propagation from time $t-1$ to t is correct, penalizing states with erroneous propagation;
- $H_{\text{final}} = \Pi_{U[\mathcal{R}]} \otimes |T\rangle\langle T|_T$ causes each state $|\phi\rangle$ built from an input $|\psi\rangle_I \in \mathcal{R}$ (but which nonetheless has a correctly initialized ancilla) to receive an energy penalty.

As we show in Lemma 12, the total Hamiltonian H has a spectral gap $1/\text{poly}(n)$ above the $\dim \mathcal{A}$–dimensional ground state subspace. However, H_{final} is not a local Hamiltonian, but as we argue in the following, it can be replaced by the usual term $H^{\text{std}}_{\text{final}} = |0\rangle\langle 0|_1 \otimes \mathbb{1} \otimes |T\rangle\langle T|_T$ while keeping the ground space dimension (up to small splitting in energies) and the $1/\text{poly}(n)$ spectral gap. As we prove in Lemma 13,

$$H^{\text{std}}_{\text{final}} \geq H_{\text{final}} - \sqrt{\epsilon}\mathbb{1}\,. \tag{10}$$

Thus, replacing H_{final} by $H^{\text{std}}_{\text{final}}$ will decrease the energy of any excited state by at most $\sqrt{\epsilon} = O(\exp(-p(n)/2))$. (The energy of the ground states is already minimal and cannot decrease.) On the other hand, the energy of any proper proof history for H cannot increase by more than

$$\langle\chi|H^{\text{std}}_{\text{final}}|\chi\rangle = \langle\chi||0\rangle\langle 0|_1 \otimes \mathbb{1}|\chi\rangle \leq \epsilon = O(\exp(-p(n)))\,,$$

due to our choice of acceptance threshold, i.e., the low energy subspace has dimension $\dim \mathcal{A}$. We thus obtain a Hamiltonian with a $\dim \mathcal{A}$ dimensional ground state subspace with energy splitting $\leq \epsilon = \exp(-p(n))$ and a $1/\text{poly}(n)$ spectral gap above. □

The following two lemmas are used in the preceding proof of Theorem 11.

Lemma 12. *H has a spectral gap of size $1/\text{poly}(n)$.*

Proof. Our proof follows closely the discussion in Ref. [**?**], cf. also [**?**]. We can block diagonalize H by the (conjugate) action of the following unitary operator,

$$W = \sum_{j=0}^{T} U_j \cdots U_1 \otimes |j\rangle\langle j|_T\,, \tag{11}$$

which maps $H \to H' = W^\dagger H W$. As this has no effect on the spectrum, we can work with the simpler H' from now on.

Let us explicitly write the effect of conjugation by W on the terms of H. The first term is unaffected, $H'_{\text{init}} = H_{\text{init}}$. The final term becomes $H'_{\text{final}} = \Pi_{\mathcal{R}} \otimes |\mathbf{0}\rangle\langle\mathbf{0}|_A \otimes |T\rangle\langle T|_T$, where $\Pi_{\mathcal{R}}$ is simply the projector onto the space $\mathcal{R}$, and the ancillas are in the correct initial state.

We can now conjugate each of the terms in $H_{\text{evol}}(t)$ separately. For example, the first term gives

$$W^\dagger (U_t \otimes |t\rangle\langle t-1|) W = \mathbb{1} \otimes |t\rangle\langle t-1|\,. \tag{12}$$

The other terms are exactly analogous, and we find that

$$H'_{\text{evol}}(t) = \mathbb{1} \otimes \tfrac{1}{2}\big[|t-1\rangle\langle t-1| + |t\rangle\langle t| - |t\rangle\langle t-1| - |t-1\rangle\langle t|\big]\,.$$

The total evolution Hamiltonian is then block-diagonal a matrix which looks like a hopping Hamiltonian in the clock register,

$$\sum_t H'_{\text{evol}}(t) = \mathbb{1} \otimes E\,, \tag{13}$$

where the $(T+1)$-by-$(T+1)$–dimensional tri-diagonal matrix E is given by

$$E = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & & & \\ -\frac{1}{2} & 1 & -\frac{1}{2} & & \\ & -\frac{1}{2} & \ddots & & \\ & & & 1 & -\frac{1}{2} \\ & & & -\frac{1}{2} & \frac{1}{2} \end{pmatrix}. \tag{14}$$

To discuss the spectrum of H', let

$$\mathcal{S}_1 = \mathcal{A} \otimes |0\rangle\langle 0|_A \otimes \mathbb{C}^{T+1}\,,$$
$$\mathcal{S}_2 = \mathcal{R} \otimes |0\rangle\langle 0|_A \otimes \mathbb{C}^{T+1}\,,$$

and $\mathcal{S}_3$ the orthogonal complement of $\mathcal{S}_1 \oplus \mathcal{S}_2$; i.e., $\mathcal{S}_1$ corresponds to evolutions starting from good proofs, $\mathcal{S}_2$ to those starting from wrong proofs, and $\mathcal{S}_3$ to evolutions with wrongly initialized ancillas. Note that $H' = H|_{\mathcal{S}_1} \oplus H|_{\mathcal{S}_2} \oplus H|_{\mathcal{S}_3}$, thus, we can analyze the spectrum for $H|_{\mathcal{S}_p}$ separately. Note that the restriction to $\mathcal{S}_p$ does not affect $H'_{\text{evol}}(t)$. Since we expect the ground state subspace to

occur on $\mathcal{S}_1$, we need to compute ground state energy and gap of $H'|_{\mathcal{S}_1}$, as well as lower bound the ground state energies of $H'|_{\mathcal{S}_2}$ and $H'_{\mathcal{S}_3}$.

First, $H'_{\text{init}}|_{\mathcal{S}} = H'_{\text{final}}|_{\mathcal{S}} = 0$, i.e., the spectrum of $H'|_{\mathcal{S}}$ equals the spectrum of E, which can be straightforwardly determined to be $1-\cos\omega_n$, with $\omega_n = n\pi/(T+1)$, and eigenfunctions $(\cos\frac{1}{2}\omega_n, \cos\frac{3}{2}\omega_n, \dots)$; the ground state degeneracy is indeed $\dim\mathcal{A}$ as desired.

On the other hand, $H'_{\text{final}}|_{\mathcal{S}_2} = 1$, and $H'_{\text{init}}|_{\mathcal{S}_3} \geq 1$, i.e., in both cases the ground state energy of $H'|_{\mathcal{S}_p}$ is lower bounded by the ground state energy of

$$E' = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & & & \\ -\frac{1}{2} & 1 & -\frac{1}{2} & & \\ & -\frac{1}{2} & \ddots & & \\ & & & 1 & -\frac{1}{2} \\ & & & -\frac{1}{2} & \frac{3}{2} \end{pmatrix},$$

which has eigenvalues $1-\cos\vartheta_n$, with $\vartheta_n = (n+\frac{1}{2})\pi/(T+\frac{3}{2})$, and eigenfunctions $(\cos\frac{1}{2}\vartheta_n, \cos\frac{3}{2}\vartheta_n, \dots)$.

It follows that H' (and thus H) has a ground state energy of $1-\cos\omega_0 = 0$, and a gap $1-\cos\frac{\pi}{2T+3} = O(1/T^2) = O(1/\text{poly}(n))$ above. □

It remains to prove Eq. (10), which follows from the following Lemma by choosing $P = |0\rangle\langle 0|_1 \otimes \mathbb{1}$, $Q = \Pi_{U[\mathcal{R}]}$, and using Eq. (7).

Lemma 13. *Let P and Q be projectors such that $\|Q(\mathbb{1}-P)Q\|_\infty \leq \epsilon$. Then*

$$P - Q \geq -\sqrt{\epsilon}\mathbb{1}\,. \tag{15}$$

Proof. We begin by recalling the result due to Jordan [**?**] (see Ref. [**?**] for a more modern treatment) for the simultaneous canonical form of two projectors. In the subspace where P and Q commute, both operators are diagonal in a common basis and the spectrum is either $(0,0), (0,1), (1,0)$, or $(1,1)$, and direct sums of those terms. In the subspace where they don't commute, the problem decomposes into a direct sum of two-by-two blocks given by

$$P_j = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad , \quad Q_j = \begin{pmatrix} c^2 & cs \\ cs & s^2 \end{pmatrix}, \tag{16}$$

where $s = \sin(\theta_j)$ for some angle θ_j, $c^2 + s^2 = 1$, and the subscript j just labels a generic block.

In fact, this two-by-two block form is completely general if we allow embedding our projectors into a larger space while preserving their rank. The rank-preserving condition guarantees that our bound is unchanged, since we are only appending blocks of zeros, and so we will consider this two-by-two form without loss of generality.

The constraint that $\|Q(\mathbb{1}-P)Q\|_\infty \leq \epsilon$ implies constraints on the values that $\sin(\theta_j)$ can take. In particular, we can directly compute this operator norm in each block separately, and we find that for all j

$$\|Q_j(\mathbb{1}-P_j)Q_j\|_\infty = \sin^2(\theta_j) \leq \epsilon\,. \tag{17}$$

We can also directly compute in each block that

$$P_j - Q_j = \begin{pmatrix} 1-c^2 & 1-cs \\ -cs & -s^2 \end{pmatrix}. \tag{18}$$

The spectrum of this operator is easily computed to be $\pm|\sin(\theta_j)|$. Thus, the least eigenvalue of $P-Q$ is bounded from below by $-\sqrt{\epsilon}$. □

Quantum vs. Classical Counting Complexity

We finally want to relate #BQP to the classical counting class #P. It is clear that any #P problem can be solved in #BQP by using a classical circuit. We will now show that conversely, #BQP can be reduced to #P under *weakly parsimonious reductions*: That is, for any function $f \in$ #BQP there exist polynomial-time computable functions α and β, and a function $g \in$ #P, such that $f = \alpha \circ g \circ \beta$. This differs from Karp reductions where no postprocessing is allowed, $\alpha = \text{Id}$, but still only requires a single call to a #P oracle, in contrast to Turing reductions.

Theorem 14. *There exists a weakly parsimonious reduction from #BQP to #P.*

Proof. We start from a verifier operator Ω for a #BQP problem. First, we use strong error reduction to let $a = 1-2^{-(n-2)}$ and $b = 2^{-(n+2)}$. It follows that

$$|\dim\mathcal{A} - \text{tr}\,\Omega| \leq 2^n 2^{-(n+2)} = \frac{1}{4} \tag{19}$$

and thus we need to compute $\text{tr}\,\Omega$ to accuracy $1/4$. This can be done using the "path integral" method previously used to show containments of quantum classes inside PP and #P [**?**]. We rewrite $\text{tr}\,\Omega = \sum_\zeta f(\zeta)$, where the sum is over products of transition probabilities along a path, which we label

$$\zeta \equiv (i_0, \dots, i_N, j_1, \dots, j_N)\,, \tag{20}$$

so that

$$\begin{aligned} f(\zeta) = &\langle i_0|_I \langle \mathbf{0}|_A U_1^\dagger |j_1\rangle\langle j_1| U_1^\dagger \cdots U_T^\dagger |j_T\rangle \times \\ &\langle i_T|\big[|0\rangle\langle 0|_1 \otimes \mathbb{1}\big]|i_T\rangle\langle i_T|U_T\cdots U_1|i_0\rangle_I |\mathbf{0}\rangle_A\,. \end{aligned} \tag{21}$$

(cf. Fig. 1 in the main manuscript for an illustration).

Since any quantum circuit can be recast in terms of real gates at the cost of doubling the number of qubits [**?**

], we can simplify the proof by assuming $f(\zeta)$ to be real. To achieve the desired accuracy it is sufficient to approximate f up to $|\zeta|+2$ digits, where $|\zeta| = \mathrm{poly}(n)$ is the number of bits in ζ. Now define

$$g(\zeta) := \mathrm{round}\big[2^{|\zeta|+2}(f(\zeta)+1)\big] \tag{22}$$

and note that $g(\zeta)$ is a positive and integer-valued function satisfying

$$\Big|\big[2^{-|\zeta|-2}\sum_{\zeta} g(\zeta) - 1\big] - \sum_{\zeta} f(\zeta)\Big| \leq \tfrac{1}{4}\,. \tag{23}$$

Finally, by defining a boolean indicator function,

$$h(\zeta,\xi) = \begin{cases} 1 & \text{if } 0 \leq \xi < g(\zeta) \\ 0 & \text{otherwise} \end{cases}$$

we can write $g(\zeta) = \sum_{\xi\geq 0} h(\zeta,\xi)$, and thus,

$$\sum_{\zeta} g(\zeta) = \sum_{\xi,\zeta} h(\zeta,\xi)\ .$$

This shows that $\mathrm{tr}\,\Omega$ can be approximated to accuracy $\frac{1}{4}$, and thus $\dim\mathcal{A}$ can be determined by counting the number of satisfying assignments of a single boolean function $h(\zeta,\xi)$ that can be efficiently constructed from Ω, i.e., by a single query to a black box solving #P problems, together with the efficient postprocessing described by Eq. (23). □