Classical Algorithms for Constant Approximation of the Ground State Energy of Local Hamiltonians

François Le Gall
Graduate School of Mathematics, Nagoya University, Japan

Abstract

We construct classical algorithms computing an approximation of the ground state energy of an arbitrary k-local Hamiltonian acting on n qubits.

We first consider the setting where a good "guiding state" is available, which is the main setting where quantum algorithms are expected to achieve an exponential speedup over classical methods. We show that a *constant* approximation (i.e., an approximation with *constant relative accuracy*) of the ground state energy can be computed classically in $\mathrm{poly}\,(1/\chi, n)$ time and $\mathrm{poly}(n)$ space, where χ denotes the overlap between the guiding state and the ground state (as in prior works in dequantization, we assume sample-and-query access to the guiding state). This gives a significant improvement over the recent classical algorithm by Gharibian and Le Gall (SICOMP 2023), and matches (up to a polynomial overhead) both the time and space complexities of quantum algorithms for constant approximation of the ground state energy. We also obtain classical algorithms for higher-precision approximation.

For the setting where no guided state is given (i.e., the standard version of the local Hamiltonian problem), we obtain a classical algorithm computing a constant approximation of the ground state energy in $2^{O(n)}$ time and $\mathrm{poly}(n)$ space. To our knowledge, before this work it was unknown how to classically achieve these bounds simultaneously, even for constant approximation. We also discuss complexity-theoretic aspects of our results.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis; Theory of computation → Quantum computation theory

Keywords and phrases approximation algorithms, quantum computing, dequantization

Digital Object Identifier 10.4230/LIPIcs.ESA.2025.73

Related Version *Full Version*: `https://arxiv.org/abs/2410.21833` [45]

Funding *François Le Gall*: JSPS KAKENHI grant No. 24H00071, MEXT Q-LEAP grant No. JPMXS0120319794, JST ASPIRE grant No. JPMJAP2302 and JST CREST grant No. JPMJCR24I4.

1 Introduction

1.1 Statement of our main results

Estimating the ground state energy of Hamiltonians is a central problem in both many-body physics and quantum complexity theory. Consider a k-local Hamiltonian

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

acting on n qubits, with $k = O(1)$. Here each term H_i acts non-trivially on only k qubits (but does not need to obey any geometric locality). Let $\mathcal{E}(H)$ denote the ground state energy of H, i.e., its smallest eigenvalue. For any $\varepsilon > 0$, we say that an estimate $\hat{E}$ is an ε-approximation of $\mathcal{E}(H)$ if

$$\left|\hat{E} - \mathcal{E}(H)\right| \leq \varepsilon \sum_{i=1}^{m} \|H_i\|.$$

It is well known that computing a $1/\mathrm{poly}(n)$-approximation of $\mathcal{E}(H)$ is QMA-hard, even for $k = 2$ and for geometrically local Hamiltonians [24, 37, 41, 56, 58]. The Quantum PCP conjecture [4, 5] posits that there exists a constant $\varepsilon > 0$ such that computing an ε-approximation of $\mathcal{E}(H)$ is QMA-hard as well.

Despite these hardness results, efficient quantum algorithms for ground state energy estimation can be constructed when a good "guiding state" is available, i.e., when a quantum state $|\psi\rangle$ that has a good overlap $|\langle\psi|\psi_0\rangle|$ with a ground state $|\psi_0\rangle$ of H is given as an additional input or can be constructed easily (this problem has been called the "guided local Hamiltonian problem" in the recent literature [21, 28, 68]). More precisely, quantum phase estimation [40, 55] and more advanced techniques [67, 26, 30, 38, 51, 53, 54, 59] lead to the following result:

▶ **Fact 1.** *Given a quantum state with overlap χ with a ground state of H, there exists a quantum algorithm that computes with high probability an ε-approximation of $\mathcal{E}(H)$ in $\mathrm{poly}\big(\frac{1}{\chi}, \frac{1}{\varepsilon}, n\big)$ time and $O\big(n + \log(\frac{1}{\varepsilon})\big)$ space.*

When $\chi = 1/\mathrm{poly}(n)$ and $\varepsilon = 1/\mathrm{poly}(n)$, both the running time and the space complexity (i.e., the number of bits and qubits needed for the computation) are polynomial in n. Even for larger values of χ, the performance of this quantum algorithm can be significantly better than the performance of classical algorithms (which typically have running time exponential in n – see later for a detailed discussion). Combined with the fact that for several important applications (e.g., quantum chemistry) good candidates for guiding states can be efficiently constructed, ground state energy estimation is one of the most promising and most anticipated applications of quantum computers (we refer to, e.g., [1, 3, 7, 10, 49, 50, 60] for discussions of these applications).

In this work we investigate the classical complexity of this guided local Hamiltonian problem. A first issue is how to present the guiding state (which is a quantum state, i.e., an exponential-dimension vector) to a classical computer. As in prior works in dequantization [8, 22, 23, 25, 28, 29, 35, 46, 64, 65], we consider sample-and-query access:

- **(i)** for any $j \in [2^n]$ we can efficiently compute $\langle j|\psi\rangle$;
- **(ii)** we can efficiently sample from the probability distribution $p\colon [2^n] \to [0,1]$ that outputs j with probability $|\langle j|\psi\rangle|^2$.

The motivation for (ii), which is the central assumption in dequantized algorithms, is as follows: since measuring the quantum state $|\psi\rangle$ in the computational basis gives a sample from the probability p, it is natural (or "fair") to assume that in the classical setting this distribution is efficiently samplable as well.

Recently, Gharibian and Le Gall [28] constructed a classical algorithm computing an ε-approximation of $\mathcal{E}(H)$ in $n^{O(\log(1/\chi)/\varepsilon)}$ time by dequantizing quantum algorithms based on the Quantum Singular Value Transformation. Here is our main result:

▶ **Theorem 1** (Simplified version)**.** *Given sample-and-query access to a quantum state with overlap χ with a ground state of H, there exists a classical algorithm that computes with high probability an ε-approximation of $\mathcal{E}(H)$ in $\mathrm{poly}\big(\frac{1}{\chi^{1/\varepsilon}}, n\big)$ time and $\mathrm{poly}\big(n, \frac{1}{\varepsilon}\big)$ space.*

Our result significantly improves the running time of the algorithm from [28]. For instance, if $\chi = \Omega(1)$, i.e., if we have a good guiding state, our algorithm has time complexity $\mathrm{poly}(2^{1/\varepsilon}, n)$ instead of $n^{O(1/\varepsilon)}$ in [28]. If $\varepsilon = \Omega(1)$, i.e., if we want only constant precision, our algorithm has time complexity $\mathrm{poly}(1/\chi, n)$ instead of $n^{O(\log(1/\chi))}$ in [28]. Additionally, our approach only uses polynomial space. Comparing Theorem 1 with the bounds of Fact 1 shows that for *constant* precision, there exist a classical algorithm matching (up to a polynomial overhead) the performance of quantum algorithms.

Using the same technique, we also obtain the following result for the case where no guided state is given (i.e., the standard version of the local Hamiltonian problem):

▶ **Theorem 2** (Simplified version). *For any constant $\varepsilon > 0$, there exists a classical algorithm that computes with high probability an ε-approximation of $\mathcal{E}(H)$ in $2^{O(n)}$ time and* $\mathrm{poly}(n)$ *space.*

To our knowledge, before this work it was unknown how to achieve simultaneously running time $2^{O(n)}$ and space complexity $\mathrm{poly}(n)$ for ground state energy estimation of arbitrary local Hamiltonians (even for constant ε). We will further discuss the implications of our results in Section 1.3 after reviewing known classical approaches for ground state energy estimation in the next subsection.

1.2 Background on classical approaches for ground state energy estimation

There are two main classical approaches for estimating the ground state energy of a local Hamiltonian:

- The power method or its variant the Lanczos method [42, 43], which estimates the ground state using matrix-vector multiplications. Since the Hamiltonian is a (sparse) matrix of dimension 2^n, the time complexity is $O^*(2^n)$.[1] There are two main issues with this approach. First, it requires storing explicit vectors in memory, which leads to space complexity $\Omega(2^n)$ and significantly reduces its applicability. Second, it is unclear how a guiding state would help significantly reduce the time complexity (having a good guiding state does reduce the number of iterations, but each iteration still requires matrix-vector multiplications of matrices and vectors of dimension 2^n).
- Quantum Monte Carlo methods, which use sampling arguments to estimate the ground state without having to store explicitly the quantum state. This approach is especially useful for "stoquastic" Hamiltonians, i.e., Hamiltonians for which all the off-diagonal elements are real and non-positive, and has lead to the design of classical algorithms as well as as complexity-theoretic investigations of the complexity of the local Hamiltonian problem for stochastic Hamiltonians [17, 19, 15, 52]. While some of these techniques have been extended to a few classes of non-stoquastic local Hamiltonians, such as gapped local Hamiltonians [16] or arbitrary Hamiltonians with succinct ground state [36], for ground state energy estimation the "sign-problem" significantly limits its applications to arbitrary local Hamiltonians [32, 66]. It is also unclear how the guiding state would help reduce the time complexity.

A third approach is direct classical simulation of the quantum circuit used in Fact 1. There are several techniques for simulating quantum circuits on a classical computer. If the circuit acts on n qubits and has m gates, the Schrödinger method stores the entire state vector in memory and performs successive matrix-vector multiplications, using roughly $m2^n$ time and 2^n space. While the space complexity can in several cases be significantly reduced using matrix product states or more general tensor networks [13], those representations also require exponential space in the worst case. On the other hand, the Feynman method calculates an amplitude as a sum of terms, using roughly 4^m time and $m+n$ space (this approach was used by Bernstein and Vazirani [12] to prove the inclusion $\mathsf{BQP} \subseteq \mathsf{P}^{\#\mathsf{P}}$). Aaronson and Chen [2] have introduced a recursive version of the Feynman method, inspired by the proof of Savitch theorem [61], that works in $2^{O(n \log m)}$ time and $\mathrm{poly}(n, m)$ space.

[1] In this paper the notation $O^*(\cdot)$ suppresses the $\mathrm{poly}(n)$ factors.

Other prior works. Several ground state energy estimation classical algorithms have also been developed for special classes of Hamiltonians, such as one-dimensional gapped local Hamiltonians [44], quantum analogues of Max Cut [6, 31, 39, 48, 57], or Hamiltonians defined on structured graphs [9, 11, 14]. Additionally, there are a few works [18, 27, 33] achieving weaker (but still nontrivial) approximation ratios of the ground state energy, and a recent work [20] achieving a constant approximation ratio for any local Hamiltonian in time slightly better than 2^n (but not space-efficiently).

1.3 Implication of our results

We now discuss several implications of our results.

Better understanding of the quantum advantage. As already mentioned, Theorem 1 implies that for any constant precision parameter ε, we can construct classical ground state energy estimation algorithms with performance matching (up to a polynomial overhead) the performance of the best known quantum algorithms. While Ref. [28] already showed this for $\chi = O(1)$, Theorem 1 proves this result for any value $\chi \in (0, 1]$. This implies that (under the assumption that current quantum algorithms are optimal) there is no superpolynomial quantum advantage for the constant-precision guided local Hamiltonian problem and gives another strong evidence that exponential quantum advantage for ground state energy estimation (and applications to, e.g., quantum chemistry) comes from the improved precision achievable in the quantum setting.

Space-efficient classical ground state energy estimation algorithms. The second main contribution of this work is the design of space-efficient algorithms for ground state energy estimation. In particular, for the case where no guiding state is available, Theorem 2 gives a $2^{O(n)}$-time $\mathrm{poly}(n)$-space classical algorithm. As already mentioned, to our knowledge before this work it was unknown how to achieve simultaneously running time $2^{O(n)}$ time and space complexity $\mathrm{poly}(n)$ for arbitrary local Hamiltonians: even for constant ε, the best running time was $2^{O(n \log n)}$ by the approach by Aaronson and Chen [2].

Potentially practical classical ground state energy estimation algorithms. From a practical perspective, the potential of Theorem 1 is even more striking. In particular, for $\chi = \Omega(1)$, i.e., when we have access to a fairly good guided state (which is the case in some applications to quantum chemistry), we obtain a $\mathrm{poly}(2^{1/\varepsilon}, n)$-time $\mathrm{poly}(n)$-space classical algorithm. Note that the running time is polynomial even for $\varepsilon = 1/\log n$. Additionally, the running time is better than 2^n, which is the typically running time of other classical methods, even for precision as low as $\varepsilon = c/n$, for a small enough constant $c > 0$. While evaluating the practicality of our algorithm is beyond the scope of this paper, we hope our algorithms find applications in many-body physics.

Complexity-theoretic implications. In order to formally discuss complexity theoretic aspects of our results and their relations with standard complexity classes as BPP, BQP and QMA, we first need to introduce decision versions of our problems. As standard in Hamiltonian complexity theory, we add the promise that either (i) $\mathcal{E}(H) \leq a$ or (ii) $\mathcal{E}(H) > b$ holds, for some values $a, b \in [0, 1]$ such that $b - a > \varepsilon$, and ask to decide which of (i) or (ii) holds. This leads to the following (standard) decision version of the local Hamiltonian problem:

$\mathsf{LH}(\varepsilon)$ (Local Hamiltonian problem – decision version)

Input: $*$ a $O(1)$-local Hamiltonian H as in Equation (1) acting on n qubits
$*$ two numbers $a, b \in [0,1]$ such that $b - a > \varepsilon$

Promise: either (i) $\mathcal{E}(H) \leq a$ or (ii) $\mathcal{E}(H) > b$ holds

Goal: decide which of (i) or (ii) holds

As already mentioned, the problem $\mathsf{LH}(\varepsilon)$ is QMA-complete for $\varepsilon = 1/\mathrm{poly}(n)$. On the other hand, for any $\varepsilon = O(1)$ Theorem 2 leads to the inclusion

$$\mathsf{LH}(\varepsilon) \in \mathsf{BPTimeSpace}\left(2^{O(n)}, \mathrm{poly}(n)\right),$$

where $\mathsf{BPTimeSpace}(t(n), s(n))$ denotes the class of (promise) decision problems that can be solved with probability at least $2/3$ by a probabilistic Turing machine running in $t(n)$ time and using $s(n)$ space.

For the guided local Hamiltonian problem, another subtle issue is how to access the guiding state $|\psi\rangle$. So far we have assumed sample-and-query access to $|\psi\rangle$ when considering classical algorithms. While satisfactory when discussing algorithmic aspects of the problem (as we did so far), such "oracle" access to the input is problematic if we want to discuss relations with standard complexity classes such as BPP, BQP and QMA. Instead, we make the following assumptions: in the quantum setting, the description of a quantum polynomial-size circuit creating $|\psi\rangle$ is given as input; in the classical setting, the description of a classical polynomial-size circuit implementing sample-and-query access to $|\psi\rangle$ is given as input. This leads to the following decision version of the guided local Hamiltonian problem:

$\mathsf{GLH}(\varepsilon, \chi)$ (Guided Local Hamiltonian problem – decision version)

Input: $*$ a $O(1)$-local Hamiltonian H as in Equation (1) acting on n qubits
$*$ the description of a $\mathrm{poly}(n)$-size circuit implementing access to a quantum state $|\psi\rangle$ with overlap at least χ with the ground state of H
$*$ two numbers $a, b \in [0,1]$ such that $b - a > \varepsilon$

Promise: either (i) $\mathcal{E}(H) \leq a$ or (ii) $\mathcal{E}(H) > b$ holds

Goal: decide which of (i) or (ii) holds

Prior results on the hardness of the guided local Hamiltonian problem [21, 28] combined with Fact 1 imply that $\mathsf{GLH}(\varepsilon, \chi)$ is BQP-complete for $\varepsilon = 1/\mathrm{poly}(n)$ and constant χ. Ref. [28] also showed that this problem is in the class BPP for constant ε and constant χ. Theorem 1 enables us to strengthen this result and show that for constant ε, the inclusion $\mathsf{GLH}(\varepsilon, \chi) \in \mathsf{BPP}$ holds for $\chi = 1/\mathrm{poly}(n)$ as well.

These complexity-theoretic implications are summarized in Table 1.

1.4 Overview of the techniques

The full version of the paper [45] includes a 3-page overview of the techniques, which is omitted here due to space constraints.

2 Preliminaries

In this section we introduce notations and definitions, and present some useful lemmas.

Table 1 The complexity of the problems $\mathsf{LH}(\varepsilon)$ and $\mathsf{GLH}(\varepsilon, \chi)$.

	$\mathsf{LH}(\varepsilon)$	$\mathsf{GLH}(\varepsilon, \Theta(1))$	$\mathsf{GLH}\left(\varepsilon, \frac{1}{\mathrm{poly}(n)}\right)$
$\varepsilon = \Theta(1)$	in $\mathsf{BPTimeSpace}(2^{O(n)}, \mathrm{poly}(n))$ (Th. 2)	in BPP (Ref. [28])	in BPP (Th. 1)
$\varepsilon = \frac{1}{\mathrm{poly}(n)}$	QMA-complete (Refs. [37, 41])	BQP-complete (Refs. [21, 28])	

2.1 Notations

General notations. For any integer N we write $[N] = \{1, \ldots, N\}$. Define $\mathbb{R}[\mathrm{poly}(n)]$ as the set of all real numbers with binary expansion of polynomial length. More precisely, for any function $f\colon \mathbb{N} \to \mathbb{N}$, we define the set

$$\mathbb{R}[f(n)] = \left\{ \pm \left(a_0 + \sum_{i=1}^{f(n)} a_i 2^i + \sum_{i=1}^{f(n)} b_i 2^{-i} \right) \;\middle|\; a_0, \ldots, a_{f(n)}, b_1, \ldots, b_{f(n)} \in \{0,1\} \right\} \subseteq \mathbb{R}$$

of all real numbers with binary expansion of polynomial length $2f(n) + 2$ (including one bit for encoding the sign). Then $\mathbb{R}[\mathrm{poly}(n)]$ is the union of the $\mathbb{R}[f(n)]$ for all polynomial functions $f(n)$. We define $\mathbb{C}[f(n)]$ and $\mathbb{C}[\mathrm{poly}(n)]$ similarly, by requiring that both the real part and the imaginary part are in $\mathbb{R}[f(n)]$ and $\mathbb{R}[\mathrm{poly}(n)]$, respectively.

Vectors and matrices. In this paper we consider vectors in $\mathbb{C}^N$ and matrices in $\mathbb{C}^{N \times N}$, for some integer N, and write $n = \lceil \log_2(N) \rceil$. Note that this notation is consistent with the notation of Section 1, where we considered the special case $N = 2^n$. We usually write quantum states (i.e., unit-norm vectors) using Greek letters and Dirac notation, e.g., we use $|\psi\rangle$ or $|\varphi\rangle$. We write arbitrary vectors (i.e., vectors of arbitrary norm) using Roman letters, e.g., we use v or w.

For a matrix $A \in \mathbb{C}^{N \times N}$ and any $\ell \in [N]$, we denote the ℓ-th row of A by $A[\ell, \cdot]$. The matrix A is normal if it can be written $A = UDU^{-1}$ where D is a diagonal matrix with real entries and U is a unitary matrix. We use $\|A\|$ to denote the spectral norm of A, which is defined for a normal matrix as the maximum magnitude of the eigenvalues of A (and defined as square root of the maximum eigenvalue of A^*A, where A^* denotes the conjugate transpose of A, in general). This norm is submultiplicative, i.e., the inequality $\|AB\| \leq \|A\|\|B\|$ holds for any matrices $A, B \in \mathbb{C}^{N \times N}$. We also have $|u^* A v| \leq \|A\|\|u\|\|v\|$ for any $u, v \in \mathbb{C}^N$, where $\|u\|$ and $\|v\|$ denote the Euclidean norms of u and v, respectively.

Eigenvalues and overlap. Consider a normal matrix $A \in \mathbb{C}^{N \times N}$. Let $A = \sum_{i=1}^{2^n} \lambda_i |u_i\rangle\langle u_i|$ be its eigenvalue decomposition, with eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{2^n}$ and corresponding orthonormal eigenvectors $|u_1\rangle, \ldots, |u_{2^n}\rangle$. We denote by $\mathcal{E}(A) = \lambda_1$ the smallest eigenvalue of A. For any $\sigma \geq 0$, let us write $S(A, \sigma) = \{i \in [N] \mid \lambda_i(A) \leq \mathcal{E}(A) + \sigma\}$. For any vector $w \in \mathbb{C}^N$, let

$$\Gamma_\sigma(A, w) = \sqrt{\sum_{i \in S(A,\, \sigma)} |\langle u_i | w \rangle|^2}$$

denote the overlap of w with the eigenspace corresponding to eigenvalues in $[\mathcal{E}(A)\,, \mathcal{E}(A) + \sigma]$. Note that the standard definition of the overlap (used in Section 1) corresponds to the case $\sigma = 0$.

2.2 Access to vectors and matrices

We now define the notions of access to vectors and matrices needed for this work. These notions are similar to prior works on dequantization [23, 46, 64], but we need to precisely discuss the encoding length and the space complexity.

We start with query access to a vector.

▶ **Definition 3.** *We have query access to a vector $w \in \mathbb{C}^N$ with encoding length $\mathsf{len}(w)$ and costs $\mathsf{qt}(w)$ and $\mathsf{qs}(w)$ if*

1. *for each $i \in [N]$, we have $w_i \in \mathbb{C}[\mathsf{len}(w)]$;*
2. *for any $i \in [N]$, the coordinate w_i can be obtained in $\mathsf{qt}(w)$ time and $\mathsf{qs}(w)$ space.*

If $\mathsf{len}(w), \mathsf{qt}(w), \mathsf{qs}(w) \leq \mathrm{poly}(n)$, we simply say that we have query access to w.

Next, we introduce the stronger notion of sample-and-query access to a vector.

▶ **Definition 4.** *We have sample-and-query access to a vector $w \in \mathbb{C}^N$ if*

1. *we have query access to w;*
2. *we can compute in $\mathrm{poly}(n)$ time[2] a sample from the distribution $p\colon [N] \to [0,1]$ such that $p(i) = \frac{|w_i|^2}{\|w\|^2}$ for each $i \in [N]$.*

When $\|w\| = 1$, Item 2 in Definition 4 states that we can efficiently sample from the same distribution as the distribution obtained when measuring the quantum state $\sum_{i=1}^N w_i|i\rangle$ in the computational basis.

We extend the notion of query access to matrices as follows:

▶ **Definition 5.** *We have query access to a matrix $B \in \mathbb{C}^{N\times N}$ if*

1. *for each $(i,j) \in [N]\times[N]$, we have $B[i,j] \in \mathbb{C}[\mathrm{poly}(n)]$;*
2. *for any $(i,j) \in [N]\times[N]$, the entry $B[i,j]$ can be obtained in $\mathrm{poly}(n)$ time;*
3. *for any $i \in [N]$, the number s_i of nonzero entries in $B[i,\cdot]$ can be obtained in $\mathrm{poly}(n)$ time;*
4. *for any $i \in [N]$ and any $\ell \in [s_i]$, the ℓ-th nonzero entry of $B[i,\cdot]$ can be obtained in $\mathrm{poly}(n)$ time.*

Items 3 and 4 in Definition 5 are needed to deal with sparse matrices.

2.3 Local Hamiltonians and matrix decompositions

We give below technical details about the description of local Hamiltonians and the matrix decompositions introduced in this work.

Description of local Hamiltonians. A k-local Hamiltonian acting on n qubits is a Hermitian matrix $H \in \mathbb{C}^{N\times N}$ with $N = 2^n$ that can be written as $H = \sum_{i=1}^m H_i$ with $m = \mathrm{poly}(n)$, where each term H_i is an Hermitian matrix acting non-trivially on at most k qubits. Each H_i can be described by a $2^k \times 2^k$ matrix representing its action on the k qubits on which it acts non-trivially. We assume that each entry of this description is in $\mathbb{C}[\mathrm{poly}(n)]$. This description is given as input. For convenience we also assume that we know $\|H_i\|$ for each $i \in [N]$.[3]

[2] Since a polynomial upper bound on the time complexity implies a polynomial upper bound on the space complexity, hereafter we omit to explicitly mention that the space complexity is $\mathrm{poly}(n)$ as well.

[3] Note that each $\|H_i\|$ can be computed from its description as a $2^k \times 2^k$ matrix. The computation is efficient when k is small, e.g., for $k = O(\log n)$, which is the most interesting regime for the local Hamiltonian problem.

Matrix decomposition. Here is the complete definition of the matrix decomposition we consider.

▶ **Definition 6.** *For a matrix A, an integer $s \geq 0$ and a real number $\kappa \in \mathbb{R}[\mathrm{poly}(n)]$, an (s, κ)-decomposition of A is a decomposition*

$$A = \sum_{i=1}^{m} A_i \quad \textit{with} \quad \sum_{i=1}^{m} \|A_i\| \leq \kappa$$

in which A_i is an s-sparse matrix for each $i \in [m]$. We always (implicitly) assume the following:

- *for each $i \in [m]$, we have query access to the matrix A_i;*
- *we know bounds $\kappa_1, \ldots, \kappa_m \in \mathbb{R}[\mathrm{poly}(n)]$ such that $\|A_i\| \leq \kappa_i$ for $i \in [m]$ and $\sum_{i=1}^{m} \kappa_i = \kappa$.*

If $\kappa = 1$, we simply call the decomposition an s-decomposition.

2.4 Lemmas

We present four lemmas that are needed to prove our results.

The first lemma is the "powering lemma" from [34] to amplify the success probability of probabilistic estimators (the formulation below for complex numbers is from [46, Lemma 3]):

▶ **Lemma 7** (Powering lemma)**.** *Consider a randomized algorithm that produces an estimate $\tilde{\mu}$ of a complex-valued quantity μ such that $|\tilde{\mu} - \mu| \leq \varepsilon$ holds with probability at least $3/4$. Then, for any $\delta > 0$, it suffices to repeat $O(\log(1/\delta))$ times the algorithm and take both the median of the real parts and the median of the imaginary parts to obtain an estimate $\hat{\mu}$ such that $|\hat{\mu} - \mu| \leq \sqrt{2}\varepsilon$ holds with probability at least $1 - \delta$.*

To perform eigenvalue estimation we will need a low-degree polynomial that approximates well the "rectangle" function. We will use the following result from [30].[4]

▶ **Lemma 8** (Lemma 29 in [30])**.** *For any $\xi \in (0, 1]$, any $\tau \in [0, 1)$ and any $\theta \in (0, 1 - \tau]$, there exists an efficiently computable polynomial $P \in \mathbb{R}[x]$ of degree $O\big(\frac{1}{\theta}\log(1/\xi)\big)$ such that $|P(x)| \in [0, 1]$ for all $x \in [-1, 1]$ and*

$$\begin{cases} P(x) \in [1 - \xi, 1] & \text{if } x \in [0, \tau], \\ P(x) \in [0, \xi] & \text{if } x \in [\tau + \theta, 1]. \end{cases} \tag{2}$$

We will use the following result from [63] that gives an upper bound on the coefficients of polynomials bounded in the interval $[-1, 1]$ (such as the polynomial from Lemma 8).

▶ **Lemma 9** (Lemma 4.1 in [63])**.** *Let $P(x) = \sum_{i=0}^{d} a_i x^i$ be a univariate polynomial of degree d such that $|P(x)| \leq 1$ for all $x \in [-1, 1]$. Then $\sum_{i=0}^{d} |a_i| \leq 4^d$.*

Finally, we discuss how to classically estimate the inner product $\langle \psi | w \rangle$ given sample-and-query access to a quantum state $|\psi\rangle$ and query access to a vector w. More precisely, we are considering the following problem:

[4] Lemma 8 follows by taking $t = \tau + \theta/2$ and $\delta' = \theta/2$ in Lemma 29 of [30]. The computability of the polynomial is discussed explicitly in [47, Appendix A.3].

$\mathsf{IP}(\varepsilon, \delta)$ (Estimation of Inner Product)

Input: $*$ sample-and-query access to a quantum state $|\psi\rangle \in \mathbb{C}^N$
$*$ query access to a vector $w \in \mathbb{C}^N$ with encoding length $\mathsf{len}(w)$ and costs $\mathsf{qt}(w)$ and $\mathsf{qs}(w)$

Output: an estimate $a \in \mathbb{C}$ such that

$$|a - \langle\psi|w\rangle| \leq \varepsilon \|w\|$$

holds with probability at least $1 - \delta$

Prior works on dequantization [23, 46, 64] have shown how to solve this problem efficiently. It can be easily checked that these approaches are space-efficient as well, leading to the following statement. For completeness we give a proof in the full version of the paper [45].

▶ **Lemma 10.** *For any $\varepsilon \in (0, 1]$ and any $\delta \in (0, 1]$, the problem $\mathsf{IP}(\varepsilon, \delta)$ can be solved classically in time $O^*\big(\mathsf{qt}(w)\,\varepsilon^{-2}\log(1/\delta)\big)$ and space $\mathsf{qs}(w) + O^*\Big(\Big(\mathsf{len}(w) + \log(1/\varepsilon)\Big)\log(1/\delta)\Big)$.*

3 Iterated Matrix Multiplication

In this section we show how to classically estimate the inner product $\langle\psi|B_r \cdots B_1|\varphi\rangle$ for sparse matrices $B_1, \ldots, B_r$ and two quantum states $|\psi\rangle$ and $|\varphi\rangle$ to which we have classical access. More precisely, we consider the following problem.

$\mathsf{IMM}(s, r, \varepsilon, \delta)$ (Estimation of Iterated Matrix Multiplication)

Input: $*$ query access to s-sparse matrices $B_1, \ldots, B_r \in \mathbb{C}^{N\times N}$
$*$ query access to a quantum state $|\varphi\rangle \in \mathbb{C}^N$
$*$ sample-and-query access to a quantum state $|\psi\rangle \in \mathbb{C}^N$

Output: an estimate $\hat{E} \in \mathbb{C}$ such that

$$\left|\hat{E} - \langle\psi|B_r \cdots B_1|\varphi\rangle\right| \leq \varepsilon \|B_1\| \cdots \|B_r\|$$

holds with probability at least $1 - \delta$

Here is the main result of this section.

▶ **Proposition 11.** *For any $s \geq 1$, any $r \geq 1$, any $\varepsilon \in (0, 1]$ and any $\delta \in (0, 1]$, the problem $\mathsf{IMM}(s, r, \varepsilon, \delta)$ can be solved classically in time*

$$O^*\big(s^r \varepsilon^{-2}\log(1/\delta)\big)$$

and space

$$O^*\big(r^2 + \big(r + \log(1/\varepsilon)\big)\log(1/\delta)\big).$$

The proof of Proposition 11 is based on the following lemma, which can be seen as a space-efficient version of the approach for iterated matrix multiplication used in [28, 62].

▶ **Lemma 12.** *There is a classical algorithm that implements query access to $B_r \cdots B_1|\varphi\rangle$ with encoding length $\mathsf{len}(B_r \cdots B_1|\varphi\rangle) = O^*(r)$ and costs $\mathsf{qt}(B_r \cdots B_1|\varphi\rangle) = O^*(s^r)$ and $\mathsf{qs}(B_r \cdots B_1|\varphi\rangle) = O^*\big(r^2\big)$.*

Proof. Here is the main idea: to obtain the ℓ-th entry of $B_r \cdots B_1|\varphi\rangle$, we only need to know the s nonzero entries of the ℓ-th row of B_r, which can be queried directly, together with the corresponding entries in the vector $B_{r-1} \cdots B_1|\varphi\rangle$, which can be computed recursively. The algorithm is described in pseudocode below.

```
Algorithm 𝒜(ℓ, r)      // computes ⟨ℓ|B_r ··· B_1|φ⟩
1  if r = 0 then
2  |  return ⟨ℓ|φ⟩ ;
3  else
4  |  z ← 0 ;
5  |  get the number of nonzero entries of the row B_r[ℓ, ·] and write it s' ;
6  |  for t from 1 to s' do
7  |  |  get the index of the t-th nonzero entry of B_r[ℓ, ·] and write it j ;
8  |  |  x ← B_r[ℓ, j] ;          // queries ⟨ℓ|B_r|j⟩
9  |  |  y ← 𝒜(j, r − 1) ;     // computes recursively ⟨j|B_{r−1} ··· B_1|φ⟩
10 |  |  z ← z + x · y ;
11 |  return z ;
```

We first analyze the correctness of the algorithm. Let $j_1, \ldots, j_{s'}$ represent the indices of the nonzero entries of the row $B_r[\ell, \cdot]$. Since

$$\langle \ell | B_r \cdots B_1 | \varphi \rangle = \sum_{t=1}^{s'} \langle \ell | B_r | j_t \rangle \langle j_t | B_{r-1} \cdots B_1 | \varphi \rangle,$$

Algorithm $\mathcal{A}(\ell, r)$ outputs $\langle \ell | B_r \cdots B_1 | \varphi \rangle$.

Let $\mathcal{T}(r)$ denote the running time of this procedure. We have $\mathcal{T}(r) \leq s\mathcal{T}(r-1) + O^*(s)$, and thus $\mathcal{T}(r) = O^*(s^r)$. For each $j \in [N]$, the entry $\langle j|\psi\rangle$ has a poly(n)-bit binary expansion. Each entry of the matrices $B_1, \ldots, B_r$ also has a poly(n)-bit binary expansion. This implies that $\mathsf{len}(B_r \cdots B_1|\varphi\rangle) = O^*(r)$. We finally consider the space complexity. The recursion tree has depth r. At each level of the recursion, the values x, y and z at Steps 8, 9 and 10 can be stored in $O^*(r)$ bits, and we need one $O(\log s)$-bit counter for storing the current value of t. The overall space complexity of the algorithm is thus $O^*(r(r + \log s)) = O^*\big(r^2\big)$. ◀

Proposition 11 is obtained by applying Lemma 10 to the vector $w = B_r \cdots B_1|\varphi\rangle$, for which we can implement query access from Lemma 12:

Proof of Proposition 11. From Lemma 12 we have query access to $w = B_r \cdots B_1|\varphi\rangle$ with encoding length $\mathsf{len}(w) = O^*(r)$ and costs $\mathsf{qt}(w) = O^*(s^r)$ and $\mathsf{qs}(w) = O^*\big(r^2\big)$. Using Lemma 10, we can then compute an estimate $a \in \mathbb{C}$ such that

$$|a - \langle \psi | B_r \cdots B_1 | \varphi \rangle| \leq \varepsilon \, \|B_r \cdots B_1 |\varphi\rangle\| \leq \varepsilon \, \|B_1\| \cdots \|B_r\|$$

holds with probability at least $1 - \delta$. The time and space complexity are

$$O^*\big(s^r \varepsilon^{-2} \log(1/\delta)\big)$$

and

$$O^*\big(r^2 + \big(r + \log(1/\varepsilon)\big) \log(1/\delta)\big),$$

respectively. ◀

4 Polynomial Transformations of Decomposable Matrices

In this section we show how to classically estimate the inner product $\langle\psi|P(A)|\varphi\rangle$ for a matrix A with an s-decomposition, a polynomial P, and two quantum states $|\psi\rangle$ and $|\varphi\rangle$ to which we have classical access. More precisely, we consider the following problem.

$\mathsf{PT}(s,d,\eta)$ (Estimation of Polynomial Transformation)

Input: $*$ a matrix $A \in \mathbb{C}^{N\times N}$ with an s-decomposition
$*$ a polynomial $P \in \mathbb{R}[x]$ of degree d with $|P(x)| \leq 1 \quad \forall x \in [-1,1]$
$*$ query access to quantum state $|\varphi\rangle \in \mathbb{C}^N$
$*$ sample-and-query access to a quantum state $|\psi\rangle \in \mathbb{C}^N$

Output: an estimate $\hat{E} \in \mathbb{C}$ such that

$$\left|\hat{E} - \langle\psi|P(A)|\varphi\rangle\right| \leq \eta \tag{3}$$

holds with probability at least $1 - 1/\exp(n)$

Here is the main result of this section.

▶ **Proposition 13.** *For any $s \geq 2$, any $d \geq 1$ and any $\eta \in (0,1] \cap \mathbb{R}[\mathrm{poly}(n)]$, the problem $\mathsf{PT}(s,d,\eta)$ can be solved classically in time $O^*\left(s^{c\cdot d}\eta^{-4}\right)$ time, for some universal constant $c > 0$, and space $O^*\left(d^2\right)$.*

The proof of Proposition 13 is based on the following lemma, whose proof is given after the proof of the proposition.

▶ **Lemma 14.** *For any $r \in \{0,\ldots,d\}$, any $\eta \in (0,1] \cap \mathbb{R}[\mathrm{poly}(n)]$ and any $\delta \in (0,1]$, there is a classical algorithm that computes an estimate $\hat{E}_r \in \mathbb{C}$ such that $\left|\hat{E}_r - \langle\psi|A^r|\varphi\rangle\right| \leq \frac{\eta}{4^d}$ holds with probability at least $1-\delta$ in time $O^*\left(s^r 2^{8d}\eta^{-4} d\log(1/\delta)\right)$ and space $O^*\left(d^2 + d\log(1/\delta)\right)$.*

Proof of Proposition 13. Let us write the polynomial P as $P(x) = \sum_{r=0}^{d} a_r x^d$. For any $\delta' \in (0,1]$, we describe how to compute an estimate $\hat{E}$ such that Equation (3) holds with probability at least $1-\delta'$. Taking $\delta' = 1/\exp(n)$ then proves the proposition.

For each $r \in \{0,\ldots,d\}$ such that $a_r \neq 0$, we apply Lemma 14 with $\delta = \frac{\delta'}{d+1}$ to obtain an approximation $\hat{E}_r$ of $\langle\psi|A^r|\varphi\rangle$ such that

$$\Pr\left[\left|\hat{E}_r - \langle\psi|A^r|\varphi\rangle\right| \leq \frac{\eta}{4^d}\right] \geq 1 - \frac{\delta'}{d+1}.$$

For each $r \in \{0,\ldots,d\}$ such that $a_r = 0$, we set $\hat{E}_r = 0$. We then output $\hat{E} = \sum_{r=0}^{d} a_r \hat{E}_r$. From the union bound and the triangle inequality, with probability at least $1-\delta'$ we have

$$\left|\hat{E} - \langle\psi|P(A)|\varphi\rangle\right| \leq \sum_{r=0}^{d} |a_r|\,|E_r - \langle\psi|A^r|\varphi\rangle| \leq \eta,$$

where we used Lemma 9 to derive the last inequality.

The time complexity is

$$O^*\left(\left(\sum_{r=0}^{d} s^r\right) 2^{8d}\eta^{-4} d\log\left(\frac{d}{\delta'}\right)\right) = O^*\left(s^d 2^{8d}\eta^{-4} d\log\left(\frac{d}{\delta'}\right)\right) = O^*\left(s^{c\cdot d}\eta^{-4}\right),$$

for some universal constant $c > 0$. The space complexity is

$$O^*\left(d^2 + d\log(d/\delta')\right) = O^*\left(d^2\right),$$

as claimed. ◀

Proof of Lemma 14. As in Definition 6, we write the s-decomposition of A as $A = \sum_{i=1}^{m} A_i$, where each A_i is an s-sparse matrix such that $\|A_i\| \leq \kappa_i$, with $\kappa_1 + \cdots + \kappa_m = 1$. Consider the probability distribution $p\colon [m]^r \to [0,1]$ defined as $p(x) = \kappa_{x_1} \cdots \kappa_{x_r}$ for any $x \in [m]^r$ (the condition $\kappa_1 + \cdots + \kappa_m = 1$ guarantees that this is a probability distribution). Define a random variable X as follows: sample a vector x from the distribution p, and set

$$X = \frac{\langle\psi|A_{x_1} \cdots A_{x_r}|\varphi\rangle}{p(x)}.$$

Repeat the above procedure $t = \left\lceil 64 \cdot 4^{2^d}/\eta^2 \right\rceil$ times and output the mean. Let Y denote the corresponding complex random variable. We have

$$\begin{aligned}
\mathbb{E}[Y] = \mathbb{E}[X] &= \sum_{x\in[m]^r} \langle\psi|A_{x_1} \cdots A_{x_r}|\varphi\rangle = \langle\psi|A^r|\varphi\rangle \\
\mathbb{V}[Y] &\leq \frac{1}{t}\mathbb{E}\left[|X|^2\right] \\
&= \frac{1}{t} \sum_{x\in[m]^r} \frac{|\langle\psi|A_{x_1} \cdots A_{x_r}|\varphi\rangle|^2}{\kappa_{x_1} \cdots \kappa_{x_r}} \\
&\leq \frac{1}{t} \sum_{x\in[m]^r} \frac{\|A_{x_1} \cdots A_{x_r}\|^2}{\kappa_{x_1} \cdots \kappa_{x_r}} \\
&\leq \frac{1}{t} \sum_{x\in[m]^r} \kappa_{x_1} \cdots \kappa_{x_r} \\
&= \frac{1}{t} \left(\kappa_1 + \cdots + \kappa_m\right)^r \\
&= \frac{1}{t}.
\end{aligned}$$

From Chebyshev's inequality, we thus obtain:

$$\Pr\left[|Y - \langle\psi|A^r|\varphi\rangle| \geq \frac{\eta}{2\sqrt{2}\cdot 4^d}\right] \leq \frac{8\cdot 4^{2d}}{\eta^2 t} \leq \frac{1}{8}. \tag{4}$$

We cannot directly use this strategy since we do not know $\langle\psi|A_{x_1} \cdots A_{x_r}|\varphi\rangle$. Instead, we estimate this quantity using Proposition 11. This leads to the following algorithm.

Algorithm $\mathcal{B}(\eta)$ `// estimates` $\langle\psi|A^r|\varphi\rangle$ `with precision` $\frac{\eta}{\sqrt{2}\cdot 4^d}$

$t \leftarrow \left\lceil 64 \cdot 4^{2d}/\eta^2 \right\rceil$;
$z \leftarrow 0$;
for i **from** 1 **to** t **do**
  Take a vector x according to the distribution p.
  Use Proposition 11 for the problem $\mathsf{IMM}\left(s, r, \frac{\eta}{2\sqrt{2}\cdot 4^d}, \frac{1}{8t}\right)$ to compute an estimate $\alpha \in \mathbb{C}$ of $\langle\psi|A_{x_1} \cdots A_{x_r}|\varphi\rangle$;
  $z \leftarrow z + \frac{\alpha}{t\cdot p(x)}$;
return z;

The complexity of Algorithm $\mathcal{B}(\eta)$ is dominated by the computation at Step 5, which is done t times. From Proposition 11, we obtain the upper bounds

$$O^*\big(t \cdot s^r 2^{4d} \eta^{-2} \log(8t)\big) = O^*\big(s^r 2^{8d} \eta^{-4} d\big)$$

and

$$O^*\left(r^2 + \left(r + \log\left(\frac{2\sqrt{2}\cdot 4^d}{\eta}\right)\right)\log(8t)\right) = O^*\big(r^2 + d^2\big) = O^*\big(d^2\big)$$

on the time and space complexities, respectively.

We now analyze the correctness of Algorithm $\mathcal{B}(\eta)$. Let Z be the random variable corresponding to the output of Step 7 when at Step 4 the vectors x's are the same vectors as in the random variable Y. For any choice of x at Step 4, the estimate α of Step 5 satisfies

$$|\alpha - \langle\psi|A_{x_1}\cdots A_{x_r}|\varphi\rangle| \leq \frac{\eta\,\kappa_{x_1}\cdots\kappa_{x_r}}{2\sqrt{2}\cdot 4^d} \tag{5}$$

with probability at least $1-1/(8t)$. Under the condition that Inequality (5) is always satisfied during the t repetitions, we have

$$|Z-Y| \leq \sum_x \frac{1}{t\cdot p(x)} \frac{\eta\,\kappa_{x_1}\cdots\kappa_{x_r}}{2\sqrt{2}\cdot 4^d} = \frac{\eta}{2\sqrt{2}\cdot 4^d},$$

where the sum is over the t vectors x chosen at Step 4. By the union bound, we thus have

$$\Pr\left[|Z-Y| > \frac{\eta}{2\sqrt{2}\cdot 4^d}\right] \leq \frac{1}{8}. \tag{6}$$

Combining Equation (4) and Equation (6) gives

$$\begin{aligned}\Pr\left[|Z - \langle\psi|A^r|\varphi\rangle| \leq \frac{\eta}{\sqrt{2}\cdot 4^d}\right] &\geq \Pr\left[|Z-Y| \leq \frac{\eta}{2\sqrt{2}\cdot 4^d} \text{ and } |Y - \langle\psi|A^r|\varphi\rangle| \leq \frac{\eta}{2\sqrt{2}\cdot 4^d}\right] \\ &\geq 1 - \Pr\left[|Z-Y| > \frac{\eta}{2\sqrt{2}\cdot 4^d}\right] - \Pr\left[|Y - \langle\psi|A^r|\varphi\rangle| > \frac{\eta}{2\sqrt{2}\cdot 4^d}\right] \\ &\geq \frac{3}{4}.\end{aligned}$$

We can then use Lemma 7 to obtain an estimate $\hat{E}_r \in \mathbb{C}$ such that $\left|\hat{E}_r - \langle\psi|A^r|\varphi\rangle\right| \leq \frac{\eta}{4^d}$ holds with probability at least $1-\delta$. Lemma 7 introduces a $\log(1/\delta)$ factor in the time complexity and an additive $O^*(d\log(1/\delta))$ term in the space complexity since for the computation of the medians we need to store $O(\log(1/\delta))$ values, each requiring $O^*(d)$ bits.[5] ◀

5 Eigenvalue Estimation

In this section we use the results proved in Section 4 to estimate the smallest eigenvalue of a normal matrix. We describe the most general problem we are solving in Section 5.1 and then prove Theorems 1 and 2 in Section 5.2.

[5] Here we are using the assumption $\kappa_i \in \mathbb{C}[\text{poly}(n)]$ for all $i \in [m]$, see Definition 6. This implies that $p(x)$ can be encoded in $O^*(r)$ bits and the output of Procedure $\mathcal{B}(\eta)$ in $O^*(d)$ bits.

5.1 General result

We consider the problem of estimating the smallest eigenvalue of a normal matrix with an (s, κ)-decomposition, given classical access to a guiding state. Here is the formal description of the problem:

$\mathsf{SE}(s, \chi, \varepsilon)$ `(Estimation of the Smallest Eigenvalue)`

Input: $*$ a normal matrix $A \in \mathbb{C}^{N \times N}$ with an (s, κ)-decomposition (for any κ)

$*$ sample-and-query access to a quantum state $|\psi\rangle \in \mathbb{C}^N$ with $\Gamma_{\frac{\varepsilon}{2}\kappa}(A, |\psi\rangle) \geq \chi$

Output: an estimate $E^* \in \mathbb{R}$ such that

$$|E^* - \mathcal{E}(A)| \leq \varepsilon\,\kappa \tag{7}$$

holds with probability at least $1 - 1/\exp(n)$

We prove the following theorem:

▶ **Theorem 15.** *For any $s \geq 2$ and any $\varepsilon, \chi \in (0, 1] \cap \mathbb{R}[\mathrm{poly}(n)]$, the problem $\mathsf{SE}(s, \chi, \varepsilon)$ can be solved classically in time $O^*\left(s^{\frac{c' \log(1/\chi)}{\varepsilon}}\right)$ time, for some universal constant $c' > 0$, and space $O^*\left(1/\varepsilon^2\right)$.*

Proof of Theorem 15. Let us define $A' = \frac{1}{2}(I + \frac{A}{\kappa})$ and write $T = \lceil 4/\varepsilon \rceil$. The (s, κ)-decomposition of A gives an $(s+1)$-decomposition of A'. Observe that A' has eigenvalues in the interval $[0, 1]$. The main idea is to divide this interval into T subintervals of length at most $\varepsilon/4$ and find in which subinterval $\mathcal{E}(A')$ lies in. Since $\mathcal{E}(A) = 2\kappa\mathcal{E}(A') - \kappa$, this will give an estimate of $\mathcal{E}(A)$.

Concretely, for any $t \in \{0, \ldots, T-1\}$, we consider the following test that checks if $\mathcal{E}(A')$ is "approximately" smaller than $t\frac{\varepsilon}{4}$. The approximation comes from the use of an estimator at Step 2 – details of the implementation of this step are discussed later.

$\mathsf{Test}(t)$ `// checks if` $\mathcal{E}(A')$ `is (approximately) smaller than` $t\frac{\varepsilon}{4}$

1 Let P be the polynomial of Lemma 8 with $\tau = t\frac{\varepsilon}{4}$, $\theta = \frac{\varepsilon}{4}$ and $\xi = \frac{\chi^2}{12}$;
2 Compute an estimate $\hat{E} \in \mathbb{C}$ such that $\left|\hat{E} - \langle\psi|P(A')|\psi\rangle\right| \leq \frac{\chi^2}{4}$;
3 **if** $\left|\hat{E}\right| \geq \frac{\chi^2}{2}$ **then** output "yes";
4 **else** output "no";

Let t^* be the smallest value of $t \in \{0, \ldots, T-1\}$ such that $\mathsf{Test}(t)$ outputs "yes". Define $E^* = t^*\frac{\varepsilon}{2}\kappa - \kappa$. The following claim, whose proof is given in the full version of the paper [45], guarantees that E^* is a correct estimate of $\mathcal{E}(A)$.

▷ Claim 16. $|E^* - \mathcal{E}(A)| \leq \varepsilon\,\kappa$.

We now discuss the implementation of Step 2. For any $\delta \in (0, 1]$, we describe how to compute an estimate E^* such that Equation (7) holds with probability at least $1 - \delta$. Taking $\delta = 1/\exp(n)$ then proves the theorem.

We use the algorithm of Proposition 13 for the problem $\mathsf{PT}(s, \deg(P), \chi^2/4)$ in order to obtain an estimator $\hat{E}$ such that $\left|\hat{E} - \langle\psi|P(A')|\psi\rangle\right| \leq \chi^2/4$ holds with probability at least $1 - \delta/T$. Since $\mathsf{Test}(t)$ is called at most T times, the union bound guarantees that with probability at least $1 - \delta$ no error occur during these tests. This implies that the output E^* satisfies the bound of Claim 16 with probability at least $1 - \delta$.

The overall time complexity is $O^*\big(T \cdot (s+1)^{c \cdot \deg(P)} \chi^{-8}\big) = O^*\Big(s^{\frac{c' \log(1/\chi)}{\varepsilon}}\Big)$, for some universal constant $c' > 0$. Since $\deg(P) = O^*(1/\varepsilon)$, the space complexity is $O^*\big(1/\varepsilon^2\big)$. ◀

5.2 Consequences: Theorems 1 and 2

We are now ready to give the full statements of Theorems 1 and 2 and prove them.

▶ **Theorem 14** (Full version). *Consider any $\varepsilon, \chi \in (0,1] \cap \mathbb{R}[\mathrm{poly}(n)]$. For any k-local Hamiltonian H on n qubits, given sample-and-query access to a quantum state $|\psi\rangle$ with*

$$\Gamma_{\frac{\varepsilon}{2}\kappa}(H, |\psi\rangle) \geq \chi,$$

where $\kappa = \sum_{i=1}^m \|H_i\|$, there is a classical algorithm that computes in $\mathrm{poly}\big(\frac{1}{\chi^{k/\varepsilon}}, n\big)$ time and $O^*\big(\frac{1}{\varepsilon^2}\big)$ space an estimate $\hat{E}$ such that

$$\left|\hat{E} - \mathcal{E}(H)\right| \leq \varepsilon \sum_{i=1}^m \|H_i\|$$

holds with probability at least $1 - 1/\exp(n)$.

Proof of Theorem 1. We apply Theorem 15 with $A = H$, $s = 2^k$ and $\kappa = \sum_{i=1}^m \|H_i\|$. ◀

▶ **Theorem 15** (Full version). *Consider any $\varepsilon \in (0,1] \cap \mathbb{R}[\mathrm{poly}(n)]$. For any k-local Hamiltonian H on n qubits, there is a classical algorithm that computes in $2^{O(kn/\varepsilon)}$ time and $O^*\big(\frac{1}{\varepsilon^2}\big)$ space an estimate $\hat{E}$ such that*

$$\left|\hat{E} - \mathcal{E}(H)\right| \leq \varepsilon \sum_{i=1}^m \|H_i\|$$

holds with probability at least $1 - 1/\exp(n)$.

Proof of Theorem 2. Let us write $H = \sum_{i=1}^{2^n} \lambda_i |u_i\rangle\langle u_i|$ the spectral decomposition of H, with $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{2^n}$ and corresponding orthonormal eigenvectors $|u_1\rangle, \ldots, |u_{2^n}\rangle$, where $\lambda_1 = \mathcal{E}(H)$. We apply Theorem 1 with the Hamiltonian $H' = H \otimes I$ acting on $2n$ qubits (here I is the identity matrix acting on n qubits) and guiding state $|\Phi\rangle = \frac{1}{\sqrt{2^n}} \sum_{i=1}^{2^n} |i\rangle|i\rangle$, for which it is trivial to implement sample-and-query access. This is a maximally entangled state, which can also be written as $|\Phi\rangle = \frac{1}{\sqrt{2^n}} \sum_{i=1}^{2^n} |u_i\rangle|v_i\rangle$, for another orthonormal basis $\{|v_1\rangle, \ldots, |v_{2^n}\rangle\}$.

Let $t \in [2^n]$ denote the multiplicity of the ground state energy of H. The eigenspace corresponding to the ground state energy of H' is thus $\mathrm{span}\,\{|u_i\rangle|j\rangle \mid i \in [t], j \in [2^n]\}$. We have

$$\Gamma_0(H, |\Phi\rangle) = \sqrt{\sum_{i=1}^{t} \sum_{j=1}^{2^n} |(\langle u_i|\langle j|)|\Phi\rangle|^2} \geq \frac{1}{\sqrt{2^n}} \sqrt{\sum_{j=1}^{2^n} |\langle j|v_1\rangle|^2} = \frac{1}{\sqrt{2^n}}.$$

The conclusion follows from Theorem 1 with $\chi = 2^{-n/2}$. ◀

References

1 Scott Aaronson. Why quantum chemistry is hard. *Nature Physics*, 5:707–708, 2009. doi:10.1038/nphys1415.

2 Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the 32nd Computational Complexity Conference (CCC 2017)*, pages 22:1–22:67, 2017. doi:10.5555/3135595.3135617.

3 Daniel S. Abrams and Seth Lloyd. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Physical Review Letters*, 83:5162–5165, 1999. doi:10.1103/PhysRevLett.83.5162.

4 Dorit Aharonov, Itai Arad, and Thomas Vidick. Guest column: the quantum PCP conjecture. *SIGACT News*, 44(2):47–79, 2013. doi:10.1145/2491533.2491549.

5 Dorit Aharonov and Tomer Naveh. Quantum NP - a survey, 2002. arXiv:quant-ph/0210077.

6 Anurag Anshu, David Gosset, and Karen Morenz. Beyond product state approximations for a quantum analogue of Max Cut. In *Proceedings of the 15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, pages 7:1–7:15, 2020. doi:10.4230/LIPIcs.TQC.2020.7.

7 Alán Aspuru-Guzik, Anthony D. Dutoi, Peter J. Love, and Martin Head-Gordon. Simulated quantum computation of molecular energies. *Science*, 309(5741):1704–1707, 2005. doi:10.1126/science.1113479.

8 Ainesh Bakshi and Ewin Tang. An improved classical singular value transformation for quantum machine learning. In *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms (SODA 2024)*, pages 2398–2453. SIAM, 2024. doi:10.1137/1.9781611977912.86.

9 Nikhil Bansal, Sergey Bravyi, and Barbara M. Terhal. Classical approximation schemes for the ground-state energy of quantum and classical Ising spin Hamiltonians on planar graphs. *Quantum Information and Computation*, 9(7):701–720, 2009. doi:10.26421/QIC9.7-8-12.

10 Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan. Quantum algorithms for quantum chemistry and quantum materials science. *Chemical Reviews*, 120(22):12685–12717, 2020. doi:10.1021/acs.chemrev.9b00829.

11 Thiago Bergamaschi. Improved product-state approximation algorithms for quantum local Hamiltonians. In *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, pages 20:1–20:18, 2023. doi:10.4230/LIPICS.ICALP.2023.20.

12 Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. doi:10.1137/S0097539796300921.

13 Jacob Biamonte and Ville Bergholm. Tensor networks in a nutshell, 2017. arXiv:1708.00006.

14 Fernando G.S.L. Brandao and Aram W. Harrow. Product-state approximations to quantum ground states. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC 2013)*, pages 871–880, 2013. doi:10.1145/2488608.2488719.

15 Sergey Bravyi. Monte Carlo simulation of stoquastic Hamiltonians. *Quantum Information and Computation*, 15(13–14):1122–1140, 2015. doi:10.26421/QIC15.13-14-3.

16 Sergey Bravyi, Giuseppe Carleo, David Gosset, and Yinchen Liu. A rapidly mixing Markov chain from any gapped quantum many-body system. *Quantum*, 7:1173, 2023. doi:10.22331/q-2023-11-07-1173.

17 Sergey Bravyi, David P. Divincenzo, Roberto Oliveira, and Barbara M. Terhal. The complexity of stoquastic local Hamiltonian problems. *Quantum Information and Computation*, 8(5):361–385, 2008. doi:10.26421/QIC8.5-1.

18 Sergey Bravyi, David Gosset, Robert König, and Kristan Temme. Approximation algorithms for quantum many-body problems. *Journal of Mathematical Physics*, 60(3):032203, March 2019. doi:10.1063/1.5085428.

19 Sergey Bravyi and Barbara Terhal. Complexity of stoquastic frustration-free Hamiltonians. *SIAM Journal on Computing*, 39(4):1462–1485, 2010. doi:10.1137/08072689X.

20 Harry Buhrman, Sevag Gharibian, Zeph Landau, François Le Gall, Norbert Schuch, and Suguru Tamaki. Beating Grover search for low-energy estimation and state preparation, 2024. doi:10.48550/arXiv.2407.03073.

21 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 *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, pages 32:1–32:19, 2023. doi:10.4230/LIPIcs.ICALP.2023.32.

22 Nai-Hui Chia, András Gilyén, Han-Hsuan Lin, Seth Lloyd, Ewin Tang, and Chunhao Wang. Quantum-inspired algorithms for solving low-rank linear equation systems with logarithmic dependence on the dimension. In *Proceedings of the 31st International Symposium on Algorithms and Computation (ISAAC 2020)*, pages 47:1–47:17, 2020. doi:10.4230/LIPIcs.ISAAC.2020.47.

23 Nai-Hui Chia, András Pal Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. *Journal of the ACM*, 69(5):33:1–33:72, 2022. doi:10.1145/3549524.

24 Toby Cubitt and Ashley Montanaro. Complexity classification of local Hamiltonian problems. *SIAM Journal on Computing*, 45(2):268–316, 2016. doi:10.1137/140998287.

25 Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, and Dacheng Tao. Quantum-inspired algorithm for general minimum conical hull problems. *Physical Review Research*, 2:033199, 2020. doi:10.1103/PhysRevResearch.2.033199.

26 Yimin Ge, Jordi Tura, and J. Ignacio Cirac. Faster ground state preparation and high-precision ground energy estimation with fewer qubits. *Journal of Mathematical Physics*, 60(2):022202, 2019. doi:10.1063/1.5027484.

27 Sevag Gharibian and Julia Kempe. Approximation algorithms for QMA-complete problems. *SIAM Journal on Computing*, 41(4):1028–1050, 2012. doi:10.1137/110842272.

28 Sevag Gharibian and François Le Gall. Dequantizing the quantum singular value transformation: Hardness and applications to quantum chemistry and the quantum PCP conjecture. *SIAM Journal on Computing*, 52(4):1009–1038, 2023. doi:10.1137/22M1513721.

29 András Gilyén, Zhao Song, and Ewin Tang. An improved quantum-inspired algorithm for linear regression. *Quantum*, 6:754, 2022. doi:10.22331/q-2022-06-30-754.

30 András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204, 2019. doi:10.1145/3313276.3316366.

31 Sean Hallgren, Eunou Lee, and Ojas Parekh. An approximation algorithm for the MAX-2-local Hamiltonian problem. In *Proceedings of the International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2020)*, pages 59:1–59:18, 2020. doi:10.4230/LIPICS.APPROX/RANDOM.2020.59.

32 Dominik Hangleiter, Ingo Roth, Daniel Nagaj, and Jens Eisert. Easing the Monte Carlo sign problem. *Science Advances*, 6(33):2375–2548, 2020. doi:10.1126/sciadv.abb8341.

33 Aram W. Harrow and Ashley Montanaro. Extremal eigenvalues of local Hamiltonians. *Quantum*, 1:6, 2017. doi:10.22331/q-2017-04-25-6.

34 Mark Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986. doi:10.1016/0304-3975(86)90174-X.

35 Dhawal Jethwani, François Le Gall, and Sanjay K. Singh. Quantum-inspired classical algorithms for singular value transformation. In *Proceedings of the 45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*, pages 53:1–53:14, 2020. doi:10.4230/LIPIcs.MFCS.2020.53.

36 Jiaqing Jiang. Local Hamiltonian problem with succinct ground state is MA-complete, 2024. arXiv:2309.10155.

37 Julia Kempe, Alexei Y. Kitaev, and Oded Regev. The complexity of the local Hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006. doi:10.1137/S0097539704445226.

38 Alex Kerzner, Vlad Gheorghiu, Michele Mosca, Thomas Guilbaud, Federico Carminati, Fabio Fracas, and Luca Dellantonio. A square-root speedup for finding the smallest eigenvalue. *Quantum Science and Technology*, 6(4):045025, 2023. doi:10.1088/2058-9565/ad6a36.

39 Robbie King. An improved approximation algorithm for quantum Max-Cut on triangle-free graphs. *Quantum*, 7:1180, 2023. doi:10.22331/q-2023-11-09-1180.

40 Alexei Yu. Kitaev. Quantum measurements and the Abelian stabilizer problem, 1995. arXiv:quant-ph/9511026.

41 Alexei Yu. Kitaev, Alexander H. Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, 2002.

42 Jacek Kuczyński and Henryk Woźniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM Journal on Matrix Analysis and Applications*, 13(4):1094–1122, 1992. doi:10.1137/0613066.

43 Cornelius Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of the National Bureau of Standards*, 45(4):255–282, 1950. doi:10.6028/jres.045.026.

44 Zeph Landau, Umesh Vazirani, and Thomas Vidick. A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians. *Nature Physics*, 11:566–569, 2015. doi:10.1038/nphys3345.

45 François Le Gall. Classical algorithms for approximating the ground energy of local Hamiltonians (full version of this paper), 2025. arXiv:2410.21833.

46 François Le Gall. Robust dequantization of the quantum singular value transformation and quantum machine learning algorithms. *computational complexity*, 34:2, 2025. doi:10.1007/s00037-024-00262-3.

47 François Le Gall, Yupan Liu, and Qisheng Wang. Space-bounded quantum state testing via space-efficient quantum singular value transformation, 2023. doi:10.48550/arXiv.2308.05079.

48 Eunou Lee and Ojas Parekh. An improved quantum Max Cut approximation via maximum matching. In *Proceedings of the 51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, pages 105:1–105:11, 2024. doi:10.4230/LIPIcs.ICALP.2024.105.

49 Joonho Lee, Dominic W. Berry, Craig Gidney, William J. Huggins, Jarrod R. McClean, Nathan Wiebe, and Ryan Babbush. Even more efficient quantum computations of chemistry through tensor hypercontraction. *PRX Quantum*, 2:030305, 2021. doi:10.1103/PRXQuantum.2.030305.

50 Seunghoon Lee, Joonho Lee, Huanchen Zhai, Yu Tong, Alexander M. Dalzell, Ashutosh Kumar, Phillip Helms, Johnnie Gray, Zhi-Hao Cui, Wenyuan Liu, Michael Kastoryano, Ryan Babbush, John Preskill, David R. Reichman, Earl T. Campbell, Edward F. Valeev, Lin Lin, and Garnet Kin-Lic Chan. Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nature Communications*, 14:1952, 23. doi:10.1038/s41467-023-37587-6.

51 Lin Lin and Yu Tong. Near-optimal ground state preparation. *Quantum*, 4:372, 2020. doi:10.22331/q-2020-12-14-372.

52 Yupan Liu. StoqMA meets distribution testing. In *Proceedings of the 16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, pages 4:1–4:22, 2021. doi:10.4230/LIPIcs.TQC.2021.4.

53 Guang Hao Low and Yuan Su. Quantum eigenvalue processing. In *Proceedings of the 2024 Symposium on Foundations of Computer Science (FOCS 2024)*, pages 1051–1062, 2024. doi:10.1109/FOCS61266.2024.00070.

54 John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. Grand unification of quantum algorithms. *PRX Quantum*, 2:040203, 2021. doi:10.1103/PRXQuantum.2.040203.

55 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

56 Roberto Oliveira and Barbara M. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Information and Computation*, 8(10):900–924, 2008. doi:10.26421/QIC8.10-2.

57 Ojas Parekh and Kevin Thompson. Application of the level-2 quantum Lasserre hierarchy in quantum approximation algorithms. In *Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, pages 102:1–102:20, 2021. doi:10.4230/LIPIcs.ICALP.2021.102.

58 Stephen Piddock and Ashley Montanaro. The complexity of antiferromagnetic interactions and 2D lattices. *Quantum Information and Computation*, 17(7-8):636–672, 2017. doi:10.26421/QIC17.7-8-6.

59 David Poulin and Pawel Wocjan. Preparing ground states of quantum many-body systems on a quantum computer. *Physical Review Letters*, 102(13):130503, 2009. doi:10.1103/PHYSREVLETT.102.130503.

60 Markus Reiher, Nathan Wiebe, Krysta M. Svore, Dave Wecker, and Matthias Troyer. Elucidating reaction mechanisms on quantum computers. *Proceedings of the National Academy of Sciences*, 114(29):7555–7560, 2017. doi:10.1073/pnas.1619152114.

61 Walter J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2):177–192, 1970. doi:10.1016/S0022-0000(70)80006-X.

62 Martin Schwarz and Maarten Van den Nest. Simulating quantum circuits with sparse output distributions, 2013. arXiv:1310.6749.

63 Alexander A. Sherstov. Making polynomials robust to noise. *Theory of Computing*, 9:593–615, 2013. doi:10.4086/TOC.2013.V009A018.

64 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, 2019. doi:10.1145/3313276.3316310.

65 Ewin Tang. Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions. *Physical Review Letters*, 127:060503, 2021. doi:10.1103/PhysRevLett.127.060503.

66 Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to Fermionic quantum Monte Carlo simulations. *Physical Review Letters*, 94:170201, 2005. doi:10.1103/PhysRevLett.94.170201.

67 Joran van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf. Quantum SDP-Solvers: Better upper and lower bounds. *Quantum*, 4:230, 2020. doi:10.22331/q-2020-02-14-230.

68 Jordi Weggemans, Marten Folkertsma, and Chris Cade. Guidable local Hamiltonian problems with implications to heuristic ansatz state preparation and the quantum PCP conjecture. In *Proceedings of the 19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024)*, pages 10:1–10:24, 2024. doi:10.4230/LIPIcs.TQC.2024.10.