# Dequantizing the Quantum Singular Value Transformation: Hardness and Applications to Quantum Chemistry and the Quantum PCP Conjecture 

Sevag Gharibian<br>Paderborn University<br>sevag.gharibian@upb.de

François Le Gall<br>Nagoya University<br>legall@math.nagoya-u.ac.jp


#### Abstract

The Quantum Singular Value Transformation (QSVT) is a recent technique that gives a unified framework to describe most quantum algorithms discovered so far, and may lead to the development of novel quantum algorithms. In this paper we investigate the hardness of classically simulating the QSVT. A recent result by Chia, Gilyén, Li, Lin, Tang and Wang (STOC 2020) showed that the QSVT can be efficiently "dequantized" for low-rank matrices, and discussed its implication to quantum machine learning. In this work, motivated by establishing the superiority of quantum algorithms for quantum chemistry and making progress on the quantum PCP conjecture, we focus on the other main class of matrices considered in applications of the QSVT, sparse matrices.

We first show how to efficiently "dequantize", with arbitrarily small constant precision, the QSVT associated with a low-degree polynomial. We apply this technique to design classical algorithms that estimate, with constant precision, the singular values of a sparse matrix. We show in particular that a central computational problem considered by quantum algorithms for quantum chemistry (estimating the ground state energy of a local Hamiltonian when given, as an additional input, a state sufficiently close to the ground state) can be solved efficiently with constant precision on a classical computer. As a complementary result, we prove that with inverse-polynomial precision, the same problem becomes BQP-complete. This gives theoretical evidence for the superiority of quantum algorithms for chemistry, and strongly suggests that said superiority stems from the improved precision achievable in the quantum setting. We also discuss how this dequantization technique may help make progress on the central quantum PCP conjecture.


## 1 Introduction

### 1.1 Motivations

The Quantum Singular Value Transformation. The Quantum Singular Value Transformation (QSVT) is a recent technique developed by Gilyén, Su , Low and Wiebe [GSLW19], which both generalizes the Hamiltonian simulation results by Low and Chuang [LC17] and exploits the framework of "qubitization" introduced in [LC19] (see also [CGJ19]). Informally, the QSVT associated to a function $f:[0, \infty) \rightarrow \mathbb{R}$ can be interpreted as a process that implements, given access to a matrix $A \in \mathbb{C}^{M \times N}$, a unitary operator "representing" the action of the (generally non-unitary) matrix $f(A)$. The matrix $f(A)$ is defined via the singular value decomposition of $A$, which explains why this transformation is called the quantum singular value transformation. The QSVT gives a unified framework to describe most quantum algorithms discovered so far, such as Hamiltonian simulation [Llo96], Grover search [Gro96], amplitude amplification [BHMT02], the HHL algorithm for solving linear systems of equations [HHL09] and some quantum walks [MNRS11, Sze04]. As shown in [GSLW19], the QSVT can lead to the development of novel quantum algorithms as well. We also refer to [MRTC21] for a survey of all these techniques.

Chia, Gilyén, Li, Lin, Tang and Wang [CGL+20a] and Jethwani, Le Gall and Singh [JGS20] have recently shown that under mild smoothness conditions on the function $f$ (which are satisfied, for instance, if $f$ is a low-degree polynomial function), quantum algorithms computing the QSVT of a matrix $A$ can be "dequantized" if $A$ is a low-rank matrix 1 and $A$ is accessible via $\ell_{2}$-norm sampling. The latter assumption, dating back to works on randomized algorithms for linear algebra [FKV04], means that rows and columns of $A$ can be sampled according to their $\ell_{2}$-norm, so that rows and columns that are heavier in $\ell_{2}$-norm will be sampled more frequently. Starting from Tang's breakthrough on dequantizing quantum algorithms for recommendation systems [Tan19], this kind of access to the data has now become a standard assumption when considering dequantization of quantum algorithms for linear algebraic problems. Assuming that the input matrix $A$ has low rank, on the other hand, is a strong assumption. This assumption is motivated by applications to machine learning. Indeed, based on this dequantization technique, Chia, Gilyén, Li, Lin, Tang and Wang [CGL+ 20a] were able to dequantize most of the quantum algorithms for machine learning developed so far.

Another central class of matrices considered in quantum algorithms, and especially in quantum algorithms for the simulation of quantum systems, is sparse matrices. Such matrices play a fundamental role in quantum complexity theory as well: the quantum PCP conjecture [AAV13], which is one of the central conjectures in quantum complexity theory, and specifically in the field of Hamiltonian complexity (see, e.g., [Osb12, GHLS15]), mainly focuses on a subclass of sparse matrices called local Hamiltonians (Hamiltonians where each local term only acts on a small number of $k$ of quantum particles, say $k \in O(1))$. In general, it is not expected that the QSVT can be dequantized for sparse matrices and all (smooth) functions $f$ - this is because this setting contains the HHL [HHL09] algorithm for inverting well-conditioned matrices, a problem known to be BQP-complete for (high-rank) sparse matrices, even for constant precision [HHL09]. It is thus fundamental to understand more thoroughly for which functions $f$ the QSVT can be dequantized, and for which $f$ we have a provable quantum advantage.

[^0]Quantum chemistry and eigenvalue estimation. Quantum chemistry is considered one of the most promising applications of quantum computers Aar09, AGDLHG05, BBMC20, LBG ${ }^{+}$21, RWS $^{+}$17]. From a theoretical perspective, the main goal in quantum chemistry algorithms is clearly defined: compute an estimate of the ground state energy of a local Hamiltonian representing the quantum system ${ }^{2}$ or, in more mathematical terms, compute the smallest eigenvalue of a sparse matrix. The most rigorous quantum approach to speed up such calculations, first proposed in AL99, AGDLHG05], utilizes the quantum phase estimation technique [Kit95]. Other approaches such as variational quantum algorithms, expected to be easier than phase estimation to implement on near-term quantum computers, have also been developed (see $\left[\mathrm{CAB}^{+} 21\right]$ for a survey of recent works on variational quantum algorithms). These other approaches, however, are mostly heuristic-based and thus their performance is much more difficult than quantum phase estimation to analyze rigorously.

Quantum Phase Estimation and quantum chemistry. The standard problem quantum phase estimation (QPE) solves is roughly as follows: given (1) an efficient quantum implementation of a unitary matrix $U \in \mathbb{C}^{N \times N}$ and (2) an eigenvector $|u\rangle \in \mathbb{C}^{N}$ of $U$ encoded as a physical quantum state, output an estimate of the eigenvalut $\sqrt{3}^{3} \lambda_{u}$ satisfying $U|u\rangle=\lambda_{u}|u\rangle$. Since any unitary can be written $U=e^{i H}$ for some Hermitian matrix $H$, one can also apply QPE to estimate eigenvalues of Hermitian $H$ (assuming one can efficiently implement powers of the unitary $e^{i H}$, which is possible for sparse $H$ [ATS03, BCC ${ }^{+} 14$, LC19]).

Thus, in quantum chemistry, $U=e^{i H}$ encodes the Hamiltonian $H$ of the system, which is a sparse matrix, and we are interested in the smallest eigenvalue (the ground state energy) of $H$. In general, this problem (denoted the local Hamiltonian problem) is well-known to be QMAcomplete [KSV02], for Quantum Merlin-Arthur (QMA) a quantum analogue of NP. Thus, an additional central working assumption $\sqrt[4]{4}$ is needed in quantum chemistry: One assumes the ability to efficiently prepare a quantum state "sufficiently close" to the true ground state. This assumption is motivated by the observation that in practice, computationally-efficient classical techniques in computational quantum chemistry often give a fairly good initial approximation of the ground state. For instance, the Hartree-Fock method [EA07] typically 5 recovers $99 \%$ of

[^1]the total energy [WLAG13].
With this assumption in place, the question is now about precision in QPE. As discussed in [LBG $\left.{ }^{+} 21\right]$, the holy grail in quantum chemistry is to estimate the ground state energy with precision less than the "chemical accuracy" (about 1.6 millihartree), where the norm of the Hamiltonian is allowed to increase polynomially in the number of particles and local dimension per particle. However, if the Hamiltonian is normalized (as in the current paper), then the required precision is inverse-polynomial in these parameters. The key point is now that QPE has two advantages: (1) It can produce an eigenvalue estimate within inverse polynomial precision efficiently, and (2) this works (with polynomial overhead) even when QPE is given only a "rough" approximation $|\tilde{u}\rangle$ of the ground state $|u\rangle$. Thus, quantum algorithms can solve the desired chemistry problem to the desired accuracy in poly-time, given a "sufficiently good" initial state $|\tilde{u}\rangle$.

Connecting QPE with the QSVT. QPE can be cast as a special case of the QSVT (see, for instance, Sections IV and V of [MRTC21]). Thus, given the quantum chemistry application above, a natural question is whether the QSVT, when applied to a local Hamiltonian $H$ or, more generally, to a sparse matrix $A \in \mathbb{C}^{N \times N}$, can be dequantized. Somewhat unexpectedly, this is (trivially) possible under mild conditions if access to the exact eigenvector is provided: if we have classical random-access to the ground state $u \in \mathbb{C}^{N}$ and can efficiently find a non-zero entry of $u$ (which can be done, for instance, if $u$ is also accessible via $\ell_{2}$-norm sampling), then the ground state energy can be computed exactly, classically, by computing only one entry of the vector $A u$ (for instance, if $u_{1} \neq 0$ then $(A u)_{1} / u_{1}$ is the eigenvalue corresponding to $u$ ), which can be done efficiently when $A$ is sparse, e.g., when $A$ has only $O(\operatorname{poly}(\log N))$ non-zero entries per row. Of course, having an exact description of $u$ in the context of quantum chemistry is not a reasonable assumption, so the interesting and practically relevant question is: Do similar efficient classical strategies exist given just an approximation $\tilde{u}$ to $u$ ? To our knowledge, this question, closely related to dequantizing the QSVT and central to establishing theoretical foundations of quantum algorithms for quantum chemistry, has not been studied in the literature 6

### 1.2 Main results

We first describe our results concerning estimating the ground state energies of local Hamiltonians, which are easier to state and most closely related to quantum chemistry and quantum PCP applications (see Section 1.4 for applications). We then describe our more general technique to dequantize the QSVT and estimate singular values of sparse matrices. To begin, we define the notions of "query-access" and "sampling-access".

Query-access and sampling-access. In this work, as in prior dequantization works7 [VDN11, SdN13, Tan19, CGL ${ }^{+}$20a, CGL ${ }^{+}$20b, GST20, JGS20, RWC $^{+}$20], we consider two ways to access
dense local Hamiltonian systems can be efficiently solved to arbitrary constant relative error $\epsilon$ via a product state ansatz [GK12a, BH13, Ber23], even though the actual ground state may be highly-entangled, and thus far from any product state.
${ }^{6}$ Wocjan and Zhang [WZ06] showed the BQP-completeness of sampling with precision $1 / \operatorname{poly}(\log (N))$ from the eigenvalues of the matrix $A$. This problem is different to our problem: we are targeting one specific eigenvalue (the smallest one) and we are given as input a vector that has large overlap with the corresponding eigenvector (this is a non-trivial assumption which may significantly aid the estimation of the eigenvalue).
${ }^{7}$ Rudi, Wossnig, Ciliberto, Rocchetto, Pontil, and Severini $\mathrm{RWC}^{+} 20$ also consider "dequantization" of simulating quantum Hamiltonian via sampling assumptions. A key difference to our work is that their runtime scales
vectors and matrices. For a vector $u \in \mathbb{C}^{N}$ for some integer $N \geq 1$, we say we have query-access to $u$ if for any index $i \in\{1, \ldots, N\}$, we can compute coordinate $u_{i} \in \mathbb{C}$ in $O(\operatorname{poly}(\log N))$ time (we assume in this paper that scalars, i.e., complex numbers, can be manipulated at unit cost). Query-access thus corresponds to standard polylogarithmic-cost RAM access to the vector $u$. Query-access to a matrix $A \in \mathbb{C}^{N \times N}$ is defined similarly (if the matrix is sparse, we assume that we can directly access the non-zero entries - see Definition [2.3 in Section (2). We say that we have sampling-access to $u$ if in addition to having query-access to $u$, we can also efficiently sample from the probability distribution that outputs index $j$ with probability $\left|u_{j}\right|^{2} /\|u\|^{2}$, where $\|\cdot\|$ is the Euclidean norm, for each $j \in\{1, \ldots, N\}$. For technical reasons, we also require knowledge of a good estimate of $\|u\|$. Our actual definition of sampling-access is even more general, since it only requires the ability to sample from a distribution close to the ideal distribution, and also only requires the knowledge of an approximation of $\|u\|$. We refer to Definition 2.4 for the precise definition.

Estimating the ground state energy of a local Hamiltonian. Given a Hamiltonian $H$ acting on $n$ qubits (i.e., a Hermitian matrix $H \in \mathbb{C}^{2^{n} \times 2^{n}}$ ), let $\lambda_{H}$ denote its ground state energy (i.e., the smallest eigenvalue of $H$ ) and $V_{H}$ the vector space spanned by its ground states (i.e., the vector space spanned by all the eigenvectors corresponding to $\lambda_{H}$ ). We denote by $\Pi_{H}$ the orthogonal projection onto $V_{H}$. One says $H$ is $k$-local if it can be written as a sum of poly $(n)$ terms $H=\sum_{i=1}^{\text {poly }(n)} H_{i}$, where each term $H_{i}$ acts non-trivially on at most $k$ qubits. In this paper, we consider the problem of computing the ground state energy $\lambda_{H}$ of a local Hamiltonian when "guided" by a vector $u$ that has some overlap with the space $V_{H}$. Since in quantum algorithms for this problem the vector $u$ is given as a quantum state, in this paper we consider the natural classical analog: we assume we have sampling-access to $u$. We denote this problem by $\operatorname{GLH}^{\text {est }}(k, \epsilon, \delta)$, where $k \geq 2$ is an integer and $\epsilon, \delta$ are real numbers in the interval $(0,1]$. The description of this problem is as follows 8

```
\(\mathrm{GLH}^{\text {est }}(k, \epsilon, \delta) \quad\) (guided local Hamiltonian)
    Input: a \(k\)-local Hamiltonian \(H\) acting on \(n\) qubits such that \(\|H\| \leq 1\)
        sampling-access to a vector \(u \in \mathbb{C}^{2^{n}}\) such that \(\|u\| \leq 1\)
    Promise: \(\left\|\Pi_{H} u\right\| \geq \delta\)
    Output: an estimate \(\hat{\lambda}\) such that \(\left|\hat{\lambda}-\lambda_{H}\right| \leq \epsilon\)
```

Quantum algorithms based on QPE (or the QSVT) can solve GLH ${ }^{\text {est }}$ in poly ( $n$ ) time for $k=O(\log n)$ even for $\epsilon=1 / \operatorname{poly}(n)$ and $\delta=1 / \operatorname{poly}(n)$, under the condition that a quantum state representing $u$ can be constructed efficiently (see for instance [MRTC21, Section IV] for a description of the decision version of the problem, which can then be converted into an

[^2]algorithm for the estimation version by binary search). In this paper, we first show that with constant precision and constant overlap, this problem can actually be solved in polynomial time even classically.
Theorem 1.1. For any constants $\epsilon, \delta \in(0,1]$ and any $k=O(\log n)$, the problem $\operatorname{GLH}^{\text {est }}(k, \epsilon, \delta)$ can be solved classically with probability at least $1-1 / \exp (n)$ in $O(\operatorname{poly}(n))$ time.
As described in the remark at the end of Section 4.3, the approach we develop to show Theorem 1.1 actually also solves the version of the problem GLH ${ }^{\text {est }}$ in which only approximate sampling-access to $u$ (as described in Definition 2.4) is provided.

We complement this result by a hardness result. We actually show hardness even for a fairly simple class of guiding vectors, which we call semi-classical states 9 and now define. First, recall the concept of subset state introduced by Grilo, Kerenidis and Sikora [GKS16]: for any non-empty subset $S \subseteq\{0,1\}^{n}$, the subset state associated with $S$ is the unit-norm vector

$$
\frac{1}{\sqrt{|S|}} \sum_{w \in S}|w\rangle
$$

i.e., the uniform superposition over the basis states corresponding to the strings in $S$. We say that a unit-norm vector $u \in \mathbb{C}^{2^{n}}$ is a semi-classical state if $u$ is a subset state associated with a subset $S \subseteq\{0,1\}^{n}$ with $|S|=\operatorname{poly}(n)$. Note that in this case the set $S$ can be represented using a polynomial number of bits. Thus, the description of the semi-classical state $u$ is defined as this polynomial-length description of $S$, i.e., an explicit enumeration of the (polynomially many) elements of $S$.

Given the description of a semi-classical state $u$, sampling-access to $u$ can trivially be implemented. This observation motivates the definition of the following non-oracle version of (a gapped decision version of) the problem GLH ${ }^{\text {est }}$, which we denote by $\operatorname{GLH}^{*}(k, a, b, \delta)$, where $k \geq 2$ is an integer, $a, b \in[0,1]$ are two real numbers such that $a<b$, and $\delta \in(0,1]$.

```
\(\mathrm{GLH}^{*}(k, a, b, \delta) \quad\) (guided local Hamiltonian with semi-classical guiding vector)
```

Input: $k$-local Hamiltonian $H$ acting on $n$ qubits such that $\|H\| \leq 1$ description of a semi-classical state $u \in \mathbb{C}^{2^{n}}$
Promises: (i) $\left\|\Pi_{H} u\right\| \geq \delta$
(ii) either $\lambda_{H} \leq a$ or $\lambda_{H} \geq b$ holds

Goal: decide which $\lambda_{H} \leq a$ or $\lambda_{H} \geq b$ holds
Our hardness result is as follows.
Theorem 1.2. For any $\delta \in(0,1 / \sqrt{2}-\Omega(1 / \operatorname{poly}(n)))$, there exist parameters $a, b \in[0,1]$ with $b-a=\Omega(1 / \operatorname{poly}(n))$ such that $\mathrm{GLH}^{*}(6, a, b, \delta)$ is BQP-hard.
Theorem 1.2]shows that the decision version of the guided local Hamiltonian problem (and thus the estimation version as well) is BQP-hard for inverse-polynomial precision, even for constant $k$ and $\delta$, and even when the guiding vector $u$ is a semi-classical state given explicitly. Combined with Theorem 1.1, this gives compelling evidence that for the guided local Hamiltonian problem, the advantage of quantum computation mainly comes from the improved precision it provides.

[^3]Dequantized algorithms for the Singular Value Transformation. We now describe our more general technique to dequantize the QSVT and estimate singular values of sparse matrices. For simplicity, here we discuss only the case of square matrices - the general case of rectangular matrices is dealt with in Section 4 .

Any matrix $A \in \mathbb{C}^{N \times N}$ can be written in terms of its Singular Value Decomposition (SVD),

$$
A=\sum_{i=1}^{N} \sigma_{i} u_{i} v_{i}^{\dagger}
$$

for real singular values $\sigma_{i} \geq 0$, and singular vectors $\left\{u_{i}\right\}_{i}$ and $\left\{v_{i}\right\}_{i}$, where each set forms an orthonormal basis for $\mathbb{C}^{N}$. For any $a, b \in[0, \infty)$ such that $a<b$, let $\Pi_{A}^{[a, b]}$ denote the orthogonal projector onto the subspace generated by the vectors $v_{i}$ such that $\sigma_{i} \in[a, b]$.

The Singular Value Transformation (SVT) of $A$ associated to an even polynomial $P \in \mathbb{R}[x]$ (i.e., a polynomial such that $P(x)=P(-x)$ ) is defined as ${ }^{10}$

$$
\begin{equation*}
P\left(\sqrt{A^{+} A}\right)=\sum_{i=1}^{N} P\left(\sigma_{i}\right) v_{i} v_{i}^{\dagger} \tag{1}
\end{equation*}
$$

We next say that a matrix is $s$-sparse if it contains at most $s$ non-zero entries per row and column. For an $s$-sparse matrix $A \in \mathbb{C}^{N \times N}$ with $\|A\| \leq 1$, a unit-norm vector $u \in \mathbb{C}^{N}$ given as a quantum state, and a polynomial $P$ of degree $d$ such that $P(x) \in[0,1]$ for all $x \in[0,1]$, the framework for QSVT developed by Gilyén, Su , Low and Wiebe [GSLW19] gives a quantum algorithm that constructs a quantum representation of the vector $P\left(\sqrt{A^{+} A}\right) u$ in $O(d s \cdot \operatorname{poly}(\log N, \log s))$ time (assuming quantum query-access to $A)$. For an arbitrary vector $v \in \mathbb{C}^{N}$ with $\|v\| \leq 1$ given via sampling-access, this quantum algorithm can then be used to estimate with high probability the complex number $v^{\dagger} P\left(\sqrt{A^{+} A}\right) u$ with additive precision $\epsilon$ in poly $(\log N, d, s, 1 / \epsilon)$ time.

In this paper, we show the following result.
Theorem 1.3 (Informal version of Theorem4.1). Let $P \in \mathbb{R}[x]$ be a constant-degree polynomial. For any constant $\epsilon>0$, there is a $O(\operatorname{poly}(\log N, s))$-time classical algorithm that, given query-access to an s-sparse matrix $A \in \mathbb{C}^{N \times N}$ with $\|A\| \leq 1$, query-access to a vector $u \in \mathbb{C}^{N}$, and sampling-access to a vector $v \in \mathbb{C}^{N}$, outputs with high probability an estimate of the quantity $v^{\dagger} P\left(\sqrt{A^{\dagger} A}\right)$ u with additive precision $\epsilon$.

In words, Theorem [1.3]efficiently dequantizes the QVST of [GSLW19] for constant-degree polynomials and constant precision, in the sense that both the quantum algorithm and our classical algorithm run in time poly $(\log N, s)$. (Naturally, exponents in the quantum and classical complexities differ. As shown in the formal statement of Theorem 1.3(Theorem4.1), in the classical case the degree of $P$ appears in the exponent, which is not the case in the quantum case).

We then show how this result can be used to compute a good estimate of the singular values of a sparse matrix, given a vector that has constant overlap with the relevant subspace.
Theorem 1.4 (Informal version of Theorem 4.3). For any interval $\left[t_{1}, t_{2}\right] \subseteq[0,1]$, any constant $\theta>0$ and any s-sparse matrix $A \in \mathbb{C}^{N \times N}$ with $\|A\| \leq 1$ given by query-access, consider the problem of deciding whether

[^4](i) A has a singular value in interval $\left[t_{1}, t_{2}\right]$, or
(ii) A has no singular value in interval $\left(t_{1}-\theta, t_{2}+\theta\right)$.

In addition, we are given as input sampling-access to a vector $u \in \mathbb{C}^{N}$, which we are promised in case $(i)$ satisfies $\left\|\Pi_{A}^{\left[t_{1}, t_{2}\right]} u\right\|=\Omega(1)$. Then, there is a $O(\operatorname{poly}(\log N, s))$-time classical algorithm that decides with high probability which of (i) and (ii) holds.
As shown in Section 4.3. Theorem 1.1 follows from Theorem 1.4 Note the condition $\left\|\Pi_{A}^{\left[t_{1}, t_{2}\right]} u\right\|=$ $\Omega(1)$ is only required for case (i); in case (ii) there is no condition imposed on $u$.

For context, quantum algorithms for the same problem have been considered implicitly in [GSLW19, Section 3.5] and explicitly in [LT20, Section 3] and [MRTC21, Section IV]. These quantum algorithms have complexity $O(\operatorname{poly}(\log N, s, 1 / \theta))$. We thus efficiently dequantize these results for constant $\theta$, in the sense that again both the classical and quantum versions have complexity polynomial in $\log N$ and $s$.

### 1.3 Overview of our techniques

BQP-hardness for inverse-polynomial-precision. Given a BQP circuit $U=U_{m} \cdots U_{1}$, the natural approach is to apply Kitaev's [KSV02] circuit-to-Hamiltonian construction to map $U$ to 5-local Hamiltonian $H=H_{\text {in }}+H_{\text {prop }}+H_{\text {out }}+H_{\text {stab }}$ (a quantum analogue of the CookLevin construction [Coo72, Lev73]), so that the ground space of $H$ "simulates" $U$ as follows: $H_{\text {in }}$ ensures the computation is initialized correctly at time $t=0, H_{\text {prop }}$ that each time step $t>0$ correctly applies $U_{t}$, and $H_{\text {out }}$ penalizes any computation which rejects. Then, one can show [KSV02] that if $U$ accepts with high (resp., low) probability, then $\lambda_{H}$ is small (resp., large). The catch is that $\mathrm{GLH}^{*}$ also requires us to produce a semi-classical state $u$ (as defined in Section(1.2) with $\left\|\Pi_{H} u\right\| \geq \delta$. This is the challenge - in general, one does not know what the exact eigenvectors of $H$ are.

To get around this, we apply three tricks. First, in the YES case (U accepts with high probability) it is known that the "history state" $\left|\psi_{\text {hist }}\right\rangle$ (Equation (5) in Section 3) has "low energy" against $H$. By the standard trick of weighting $H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}$ by a large polynomial $\Delta$, we can force $\left|\psi_{\text {hist }}\right\rangle$ to be $1 / \operatorname{poly}(\Delta)$ close to a ground state $|\phi\rangle$. To next ensure $|\phi\rangle$ can be approximated by a semi-classical state, we "pre-idle ${ }^{11}$ " $U$ by adding a large number $N$ of identity gates to the start of $U$ - this ensures the first $N$ time steps of $\left|\psi_{\text {hist }}\right\rangle$ have large overlap with a semi-classical state (roughly, the all-zeroes initial state of the computation, coupled with an equal superposition over the first $N$ time steps).

However, this alone does not suffice - in the NO case ( $U$ accepts with low probability), we know nothing about the ground space of $H$ (i.e., we do not even have $\left|\psi_{\text {hist }}\right\rangle$ to guide us). To resolve this, we take inspiration from an arguably unexpected source — the study of $\mathrm{PQMA}[\log ] .12$ There, Ambainis [Amb14] introduced a block-encoding technique to embed QMA query answers into a desired qubit of the ground space. While our work is unrelated to $P^{\mathrm{QMA}}[\log ]$, what this technique allows one to do is encode YES and NO instances of $U$ into orthogonal subspaces, affording us a high degree of control in engineering what the "NO subspace" looks

[^5]like. This, in turn, allows us to ensure the semi-classical state $u$ we design has large overlap with the ground state, even in the NO case.

Dequantizing the QSVT. We focus on explaining the main ideas to prove Theorem 1.1, which is a special case of our most general result (Theorem (1.4) that "dequantizes" the QSVT. The main idea is best illustrated by considering classical algorithms for the decision version of the guided local Hamiltonian problem, which asks to decide which of $\lambda_{H} \leq a$ and $\lambda_{H} \geq b$ holds, for $b-a=\Omega(1)$, given sampling-access to a vector such that $\left\|\Pi_{H} u\right\| \geq \delta=\Omega(1)$. Assume for simplicity that (1) all eigenvalues of $H$ are positive (i.e., $H$ is positive-definite), so that the concepts of singular value decomposition and spectral decomposition coincide (i.e., singular values equal eigenvalues), and that (2) the ground space of $H$ is non-degenerate, i.e., $\lambda_{H}$ has multiplicity 1 . Then, one has (unknown) spectral decomposition

$$
H=\sum_{i=1}^{N} \sigma_{i} u_{i} u_{i}^{\dagger}
$$

where $0<\sigma_{1}<\sigma_{2} \leq \sigma_{3} \leq \cdots \leq \sigma_{N} \leq 1$, with $\sigma_{1}=\lambda_{H}$. Write the vector $u$ with respect to the same (unknown) basis $u=\sum_{i=1}^{N} \alpha_{i} u_{i}$. From the promise on $u$, we know that $\left|\alpha_{1}\right| \geq \delta=\Omega(1)$.

Conceptually, the QSVT-based approach for this problem works as follows. Consider a smooth-enough function $f:[0,1] \rightarrow[0,1]$ such as $f(x) \approx 1$ if $x \in[0, a]$ and $f(x) \approx 0$ if $x \in[b, 1]$. Applying the function $f$ to $H$ gives

$$
f(H)=\sum_{i=1}^{N} f\left(\sigma_{i}\right) u_{i} u_{i}^{\dagger} \approx \sum_{\substack{i \in\{1, \ldots, N\} \\ \sigma_{i} \leq a}} u_{i} u_{i}^{\dagger},
$$

from which we conclude that we have

$$
\begin{cases}\|f(H) u\| \approx 0 & \text { if } \lambda_{H} \geq b  \tag{2}\\ \|f(H) u\| \geq \delta f\left(\lambda_{H}\right) \approx \delta & \text { if } \lambda_{H} \leq a .\end{cases}
$$

The objective is to use (2) to decide whether $\lambda_{H} \leq a$ or $\lambda_{H} \geq b$. There are now two difficulties: how to choose the threshold function $f$, and how to evaluate the quantity $\|f(H) u\|$.

To enable efficient evaluation of $f(H)$, the function $f$ is chosen as a polynomial of the smallest possible degree. As shown in prior works on quantum algorithms for problems related to singular value estimation [LC17, GSLW19, MRTC21], it turns out that since $b-a=\Omega(1)$, there exists a constant-degree polynomial that approximates our desired threshold function well.

The key technical issue is to classically efficiently evaluate the quantity $\|f(H) u\|$. For this, we first show how to recursively efficiently compute, $\forall i \in\{1, \ldots, N\}$, the $i$-th entry of the vector $f(H) u$ in poly $(\log N, s)$ time given query-access to the vector $u$ (which we have). This crucially uses the facts that $f$ is a constant-degree- $d$ polynomial and $H$ an $s$-sparse matrix. Then, the key technical insight is that given sampling-access to a vector $v$, it is possible to compute a good approximation of $v^{\dagger} f(H) u$ in poly $(\log N, s)$ time as well. To prove this result, we combine techniques from prior works on dequantization [Tan19] with estimation techniques tailored to our application (this is the main contribution of Section 4.1). Roughly, the approach is natural - to estimate an inner product $v^{\dagger} f(H) u$ of exponentially long vectors, we exploit our sampling assumption to "pick out" the "heaviest/most important" terms in the (sum comprising the) inner product. By rescaling and averaging these appropriately, one can show via the

Chebyshev inequality that $\operatorname{poly}\left(s^{d}, 1 / \epsilon\right)$ random samples suffice to obtain $\epsilon$ precision. Finally, taking $v=u$, which is possible since we have sampling-access to $u$, enables us to calculate an approximation of $u^{\dagger} f(H) u$, which leads to an approximation of $\|f(H) u\|$, and then enables us to use (2) to decide which of $\lambda_{H} \leq a$ and $\lambda_{H} \geq b$ holds.

### 1.4 Implications to quantum chemistry and the quantum PCP conjecture

We now describe applications of our dequantization technique to quantum chemistry and to the quantum PCP conjecture.

Towards a quantum advantage for quantum chemistry. The problem GLH ${ }^{\text {est }}$ may be viewed as a theoretical abstraction of the core problem considered in quantum algorithmic proposals for quantum chemistry based on quantum phase estimation: computing an estimate of the ground state energy of a local Hamiltonian, given access to a guiding state sufficiently close to the ground state. Indeed, essentially all those proposals use as guiding state a quantum state that has an efficient classical description that enables sampling-access (e.g., the state obtained by the Hartree-Fock method). With respect to this viewpoint, Theorem 1.2 shows that computing an inverse-polynomial estimate, which is the level of precision needed for quantum chemistry (Section 1.1), is BQP-hard ${ }^{13}$ Theorem 1.1, on the other hand, shows that this problem can be solved efficiently with constant precision on a classical computer. This gives theoretical evidence that the level of precision required for quantum chemistry problems can only be achieved on quantum computers, not classical ones. Thus, quantum chemistry may indeed be a practical application at which quantum computers demonstrate "superiority" over classical ones.

Of course, there are two important caveats to keep in mind: (1) By definition, GLH ${ }^{\text {est }}$ assumes possession of a good "guiding state" $u$. As discussed in Section 1.1, although in practice such a guiding state can often be found, in the worst case, even the existence of such a state is not generally guaranteed. This is because GLH ${ }^{\text {est }}$ without $u$ is essentially the $k$-local Hamiltonian problem (defined below), which is QMA-complete [KSV02] - and whether ground states of QMA-hard local Hamiltonians have efficient quantum preparation circuits remains a major open question (i.e., is $\mathrm{QCMA}=\mathrm{QMA}$ ?). (2) In practice, the particular properties of a given local Hamiltonian (such as geometric constraints) could in principle be exploited to bypass the worst-case BQP-hardness result of Theorem 1.2 in special cases. Indeed, one of the main appeals of recent "quantum supremacy/advantage" frameworks such as Boson Sampling [AA11] (or Random Circuit Sampling [AAB ${ }^{+}$19, BFNV19]) is that, in the long-term setting where "noise-free" quantum computers may become a reality (i.e., the setting in which full-scale quantum phase estimation for quantum chemistry is feasible), one can rigorously prove [AA11] not just worst-case hardness of Boson Sampling for classical computers, but even average-case hardness! Whether a similar average-case hardness result could hold for GLH ${ }^{\text {est }}$ is left as an open question. Nevertheless, the appeal of GLH ${ }^{\text {est }}$ (versus, say, Boson Sampling) is that it is a practically motivated task, whereas all existing near-term "quantum advantage" proposals (such as Boson Sampling and Random Circuit Sampling) are primarily "proofs-ofprinciple", i.e., do not have strong practical applications.

[^6]Finally, note that our classical simulation (Theorem 1.1) applies to general sparse Hamiltonians. ${ }^{14}$ Thus, if one did have a good guiding state for a particular fermionic Hamiltonian $H_{f}$ from quantum chemistry, one could in principle apply a fermion-to-qubit mapping [JW28, BK02. VC05, WHT16, SW19, SBMW19, JMBN19, DKBC21] to obtain a sparse Hamiltonian $H$, and then apply Theorem 1.1

Quantum PCP conjecture. The standard (i.e., non-guided) Local Hamiltonian problem considered in quantum complexity theory, which we denote by $\operatorname{LH}(k, a, b)$, where $k \geq 2$ is an integer and $a, b \in[-1,1]$ are two real numbers such that $a<b$, is defined as follows (see, e.g., [GHLS15]).

```
LH ( \(k, a, b\) ) ( \(k\)-local Hamiltonian problem)
    Input: a \(k\)-local Hamiltonian \(H\) acting on \(n\) qubits such that \(\|H\| \leq 1\)
    Promise: \(\lambda_{H} \leq a\), or \(\lambda_{H} \geq b\)
    Goal: decide which of \(\lambda_{H} \leq a\) or \(\lambda_{H} \geq b\) holds
```

The standard version of the quantum PCP conjecture (see, e.g., [AAV13, Gri18]) states that this problem is QMA-hard even for $k=O(1)$ and $b-a=\Omega(1)$. (Note the norm constraint $\|H\| \leq 1$ in LH is crucial for this to be a meaningful statement, as otherwise one could simply multiply $H$ by a large enough polynomial to achieve $b-a \in \Omega(1)$.)

Conjecture 1 (quantum PCP conjecture). There exist $k=O(1)$ and $a, b \in[-1,1]$ satisfying $b-a=$ $\Omega(1)$ such that $\mathrm{LH}(k, a, b)$ is QMA-hard.

This conjecture is one of the central conjectures in quantum complexity theory.
In Section5, we introduce the concept of quantum states (i.e., unit complex vectors) having a succinct representation that allows sampling-access (Definition 5.1) ${ }^{15}$ This concept is natural, in view of the correspondence between quantum states and sampling-access to vectors that has now become standard, and is used both in the present work and in prior works on dequantization [CGL ${ }^{+} 20 \mathrm{a}, \mathrm{CGL}^{+} 20 \mathrm{~b}$, GST20, JGS20, Tan19]. A consequence of our dequantization result is the following result (a formal version is stated and proved in Section 5).

Theorem 1.5 (Informal version of Theorem [5.2). Consider problem $\mathrm{LH}(k, a, b)$ with the following additional promise: the ground state of $H$ has constant overlap with a complex vector that has a succinct representation allowing sampling-access. For any $k=O(\log n)$, any $a, b \in[-1,1]$ such that $b-a=$ $\Omega(1)$, this variant of $\mathrm{LH}(k, a, b)$ is in the complexity class MA.

We stress that Theorem 1.5 crucially requires $b-a=\Omega(1)$. Under the widely believed assumption QMA $\neq M A$, this theorem can thus be interpreted either as evidence that the quantum

[^7]PCP conjecture (Conjecture 1) is false, or as evidence that the ground states of the Hamiltonians involved in the quantum PCP conjecture do not have good approximations as states with succinct representations allowing sampling-access.

To discuss the latter possibility further, recall that assuming $Q C M A \neq Q M A$, the only way the quantum PCP conjecture could hold is if ground states of QMA-hard Hamiltonians cannot be reasonably approximated by (even non-uniform) poly-size quantum circuits. Thus, a first step towards proving the quantum PCP conjecture would be to argue that local Hamiltonians with such "complex" low-energy spaces even exist in the first place. This is essentially the premise of Freedman and Hasting's No Low-Energy Trivial States (NLTS) conjecture ${ }^{16}$ [FH14, EH17]. To state the latter, one defines the concept of trivial states: a family of quantum states $\left\{\left|\varphi_{n}\right\rangle\right\}_{n \in \mathbb{N}}$ is trivial if the states in the family have low-complexity entanglement, in the sense that they can be generated by a family of constant-depth quantum circuits.

Conjecture 2 (NLTS conjecture). There exist a family of $O(1)$-local Hamiltonians $\left\{H_{n}\right\}_{n \in \mathbb{N}}$, where each $H_{n}$ acts on $n$ qubits, and a constant $\epsilon>0$ such that for any family of trivial states $\left\{\left|\varphi_{n}\right\rangle\right\}_{n \in \mathbb{N}}$, where each $\left|\varphi_{n}\right\rangle$ is an $n$-qubit state, we have for any sufficiently large $n$ :

$$
\left\langle\varphi_{n}\right| H_{n}\left|\varphi_{n}\right\rangle>\lambda_{H_{n}}+\epsilon .
$$

Note that NLTS is based on the presumption that NP $\neq$ QMA (since expectation values for local observables against states preparable by constant-depth circuits can be classically efficiently evaluated via a light-cone argument ${ }^{17}$ ), Theorem 1.5 is based on the presumption MA $\neq$ QMA. In words, if the quantum PCP conjecture holds and MA $\neq \mathrm{QMA}$, Theorem 1.5 implies that, roughly, sampling access to any low-energy state of a QMA-hard local Hamiltonian should not be possible. We formalize this via the following variant of the NLTS conjecture, which we call the No Low-Energy Samplable States (NLSS) conjecture.

Conjecture 3 (NLSS conjecture - informal version of Conjecture 4). There exist a family of O(1)local Hamiltonians $\left\{H_{n}\right\}_{n \in \mathbb{N}}$, where each $H_{n}$ acts on $n$ qubits, and a constant $\epsilon>0$ such that for any family of quantum states $\left\{\left|\varphi_{n}\right\rangle\right\}_{n \in \mathbb{N}}$ that has a succinct representation allowing sampling-access, where each $\left|\varphi_{n}\right\rangle$ is an $n$-qubit state, we have for any sufficiently large $n$ :

$$
\left\langle\varphi_{n}\right| H_{n}\left|\varphi_{n}\right\rangle>\lambda_{H_{n}}+\epsilon .
$$

(The formal statement of Conjecture 3 is given in Section 5) In principle, since the assumption MA $\neq$ QMA is stronger than NP $\neq$ QMA, the NLSS conjecture "should" be "easier" to prove than NLTS. Conversely, if the NLSS conjecture can be shown false, then the quantum PCP conjecture is false. Morally, this raises the interesting possibility that the "straw which breaks the quantum PCP camel's back" is the same as that used to refute quantum advantages for (say) recommender systems and certain quantum machine learning algorithms - Tang's classical sampling-access assumption [Tan19].

[^8]
## 2 Preliminaries

In this paper we assume that arithmetic operations in C (i.e., addition, subtraction, multiplication and division of two scalars) can be done at unit cost ${ }^{18}$ For any integer $s \geq 1$, we say that a matrix is $s$-sparse if each row and column has at most $s$ non-zero entries. We say that a polynomial $P \in \mathbb{R}[x]$ is even if $P(x)=P(-x)$ for all $x \in \mathbb{R}$, and say that $P$ is odd if $P(x)=-P(-x)$ for all $x \in \mathbb{R}$.

For $\epsilon \in(0,1)$, we say that a real number $\tilde{\mu}$ is an $\epsilon$-additive-approximation of a real number $\mu$ if $|\tilde{\mu}-\mu| \leq \epsilon$. As in [Mon15], we will use the following "powering lemma" from [JVV86] to amplify the success probability of probabilistic estimators 10

Lemma 2.1 (Lemma 6.1 in [JVV86]). Let $\mathcal{A}$ be an algorithm that produces with probability at least $3 / 4$ an $\epsilon$-additive-approximation of a quantity $\mu$. Then, for any $\delta>0$, it suffices to repeat $\mathcal{A} O(\log (1 / \delta))$ times and take the median to obtain an $\epsilon$-additive-approximation of $\mu$ with probability at least $1-\delta$.

### 2.1 Query-access and sampling-access to vectors and matrices

We first define two natural ways of accessing vectors and matrices. First, we define queryaccess to a vector, which corresponds to the standard model where the input is stored in a (classical) random-access-memory with polylogarithmic-time access cost.
Definition 2.2. We say that we have query-access to a vector $u \in \mathbb{C}^{N}$ if there exists an $O(\operatorname{poly}(\log N))$ time classical algorithm $\mathcal{Q}_{u}$ that on input $i \in\{1, \ldots, N\}$, outputs the entry $u_{i}$.

We define query-access to sparse matrices in a way that allows to efficiently recover the non-zero entries.

Definition 2.3. We say that we have query-access to an s-sparse matrix $A \in \mathbb{C}^{M \times N}$ if we have access to two $O(\operatorname{poly}(\log M N))$-time classical algorithms $\mathcal{Q}_{A}^{\text {row }}$ and $\mathcal{Q}_{A}^{\text {col }}$ such that

- on input $(i, \ell) \in\{1, \ldots, M\} \times\{1, \ldots, s\}$, the algorithm $\mathcal{Q}_{A}^{\text {row }}$ outputs the $\ell$-th non-zero entry of the $i$-th row of $A$ if this row has at least $\ell$ non-zero entries, and outputs an error message otherwise;
- on input $(j, \ell) \in\{1, \ldots, N\} \times\{1, \ldots, s\}$, the algorithm $\mathcal{Q}_{A}^{\text {col }}$ outputs the $\ell$-th non-zero entry of the $j$-th column of $A$ if this column has at least $\ell$ non-zero entries, and outputs an error message otherwise.

We now define the concept of sampling-access to a vector. Compared with prior works [CGL ${ }^{+}$20a, CGL ${ }^{+}$20b, GST20, JGS20, Tan19], our definition introduces a slight generalization that we call $\zeta$-sampling-access, where $\zeta \in[0,1)$ is a parameter (for $\zeta=0$ we obtain the same notion of sampling-access as in prior works). Note that [CGL+20a] also considered a slighly different generalization, which was called $\phi$-oversampling.

Definition 2.4. For a parameter $\zeta \in[0,1)$, we say that we have $\zeta$-sampling-access to a vector $u \in \mathbb{C}^{N}$ if the following three conditions are satisfied:

[^9](i) we have access to an $O(\operatorname{poly}(\log N))$-time classical algorithm $\mathcal{Q}_{u}$ that on input $i \in\{1, \ldots, N\}$, outputs the entry $u_{i}$;
(ii) we have access to an $O(\operatorname{poly}(\log N))$-time classical algorithm $\mathcal{S Q}_{u}$ that samples from a probability distribution $p:\{1, \ldots, N\} \rightarrow[0,1]$ such that
$$
p(j) \in\left[(1-\zeta) \frac{\left|u_{j}\right|^{2}}{\|u\|^{2}}(1+\zeta) \frac{\left|u_{j}\right|^{2}}{\|u\|^{2}}\right]
$$
for all $j \in\{1, \ldots, N\} ;$
(iii) we are given a real number $m$ satisfying $|m-\|u\|| \leq \zeta\|u\|$.

We simply say that we have sampling-access to $u$ (without specifying $\zeta$ ) if we have 0 -sampling-access.

### 2.2 The Singular Value Transformation

Any matrix $A \in \mathbb{C}^{M \times N}$ can be written in the form

$$
\begin{equation*}
A=\sum_{i=1}^{\min (M, N)} \sigma_{i} u_{i} v_{i}^{\dagger} \tag{3}
\end{equation*}
$$

where the $\sigma_{i}$ 's are nonnegative real numbers, $\left\{u_{i}\right\}_{i}$ are orthonormal vectors in $\mathbb{C}^{M}$ and $\left\{v_{i}\right\}_{i}$ are orthonormal vectors in $\mathbb{C}^{N}$. Such a decomposition is called a Singular Value Decomposition of $A$. Note that while the $\sigma_{i}^{\prime}$ s are unique (up to reordering), the choice of the $u_{i}^{\prime}$ 's and $v_{i}^{\prime}$ 's is not unique. When $M \geq N$ we obviously have $\min (M, N)=N$. When $M<N$, it will be convenient to also write the decomposition as

$$
\begin{equation*}
A=\sum_{i=1}^{N} \sigma_{i} u_{i} v_{i}^{\dagger} \tag{4}
\end{equation*}
$$

by using the following convention: for each $i \in\{M+1, \ldots, N\}$ we set $\sigma_{i}=0$ and choose an arbitrary vector $u_{i}$ in $\mathbb{C}^{M}$, and we choose the $v_{i}$ 's for $i \in\{M+1, \ldots, N\}$ so that $\left\{v_{i}\right\}_{i \in\{1, \ldots, N\}}$ forms an orthonormal basis $\mathbb{C}^{N}$.

Consider an even polynomial $P \in \mathbb{R}[x]$ and write

$$
P(x)=a_{0}+a_{2} x^{2}+a_{4} x^{4}+\cdots+a_{2 d} x^{2 d}
$$

where $2 d$ is the degree of $P$ and $a_{0}, \ldots, a_{2 d}$ are real coefficients. The polynomial singular value transformation of $A$ associated to $p$ is defined as

$$
P\left(\sqrt{A^{\dagger} A}\right)=a_{0} I_{N}+a_{2} A^{\dagger} A+a_{4}\left(A^{\dagger} A\right)^{2}+\cdots+a_{2 d}\left(A^{\dagger} A\right)^{d}
$$

where $I_{N}$ denotes the identity matrix of dimension $N$. Using Eq. (4), we obtain

$$
\begin{aligned}
P\left(\sqrt{A^{\dagger} A}\right) & =a_{0} I_{t}+a_{2} \sum_{i=1}^{N} \sigma_{i}^{2} v_{i} v_{i}^{\dagger}+a_{4} \sum_{i=1}^{N} \sigma_{i}^{4} v_{i} v_{i}^{\dagger}+\cdots+a_{2 d} \sum_{i=1}^{N} \sigma_{i}^{2 d} v_{i} v_{i}^{\dagger} \\
& =\sum_{i=1}^{N} P\left(\sigma_{i}\right) v_{i} v_{i}^{\dagger} .
\end{aligned}
$$

For any $a, b \in[0,1]$ such that $a<b$, we define

$$
\Pi_{A}^{[a, b]}=\sum_{\substack{i=\{1, \ldots, N\} \\ \sigma_{i} \in[a, b]}} v_{i} v_{i}^{\dagger}
$$

Note that $\Pi_{A}^{[a, b]}$ is the orthogonal projector into the subspace of the row space corresponding to the singular values in the interval $[a, b]$.

Remark 2.5. The Singular Value Transformation can also be defined for an odd polynomial $P \in \mathbb{R}[x]$. In this case, however, the definition needs to be slightly adjusted by swapping the left singular vectors $\left\{u_{i}\right\}_{i}$ and the right singular vectors $\left\{v_{i}\right\}_{i}$. Our dequantization technique can easily be adapted to work in this case as well, but to avoid confusion in this paper we discuss only the case of even polynomials, which is enough for our application to estimating the singular values of a sparse matrix.

## 3 BQP-Hardness of Polynomial-Precision Eigenvalue Estimation

In this section we prove the BQP-Hardness of the guided local Hamiltonian problem with inverse-polynomial precision (Theorem 1.2).

Proof of Theorem 1.2 Let $\Pi=\left(\Pi_{\mathrm{yes}}, \Pi_{\mathrm{no}}\right)$ be a promise problem in BQP, and $x \in\{0,1\}^{n}$ an input. Let $U=U_{m} \cdots U_{1}$ be the poly-time uniform circuit consisting of 1- and 2-qubit gates $U_{i}$ deciding $\Pi$. Specifically, $U$ takes in an $n$-bit input register $A$, and $p(n)$-qubit (for polynomial $p$ ) ancilla register $B$ initialized to $|0\rangle^{\otimes p(n)}$. If $x \in \Pi_{\text {yes }}$ (resp., $x \in \Pi_{\mathrm{no}}$ ), then measuring U's output qubit in the standard basis yields 1 with probability at least $\alpha$ (resp., at most $\beta$ ). Via standard error reduction for BQP (via parallel repetition), we may assume $\alpha=1-2^{-n}$ and $\beta=2^{-n}$, We give a poly-time many-one reduction from $(x, U)$ to an instance of $H$ of $\operatorname{GLH}^{*}(k, a, b, \delta)$, for $k, a, b, \delta$ to be chosen as needed.

Kitaev's construction. To begin, for completeness we restate Kitaev's [KSV02] 5-local circuit-to-Hamiltonian construction elements applied to our $U$ :

$$
\begin{aligned}
H_{\text {in }} & :=(I-|x\rangle\langle x|)_{A} \otimes(I-|0 \cdots 0\rangle\langle 0 \cdots 0|)_{B} \otimes|0\rangle\left\langle\left. 0\right|_{C}\right. \\
H_{\text {out }} & :=|0\rangle\left\langle\left. 0\right|_{\text {out }} \otimes \mid m\right\rangle\left\langle\left. m\right|_{C}\right. \\
H_{\text {stab }} & :=\sum_{j=1}^{m-1}|0\rangle\left\langle\left. 0\right|_{C_{j}} \otimes \mid 1\right\rangle\left\langle\left. 1\right|_{C_{j+1}}\right. \\
H_{\text {prop }} & :=\sum_{t=1}^{m} H_{t}, \text { where } H_{t} \text { is defined as } \\
H_{t} & :=-\frac{1}{2} U_{t} \otimes|t\rangle\left\langle\left. t-\left.1\right|_{C}-\frac{1}{2} U_{t}^{+} \otimes \right\rvert\, t-1\right\rangle\left\langle\left. t\right|_{C}+\frac{1}{2} I \otimes(|t\rangle\langle t|+|t-1\rangle\langle t-1|)_{C}\right.
\end{aligned}
$$

where:

- $A, B$, and $C$ are the input, ancilla, and clock registers, respectively,
- $C$ is encoded implicitly in unary,
- $C_{j}$ denotes the $j$ th qubit of register $C$,
- $|0\rangle\left\langle\left. 0\right|_{\text {out }}\right.$ projects onto the dedicated output wire of $U$ (for example, the first qubit of $\left.B\right){ }^{20}$
- finally, all Hamiltonians above are positive semi-definite.

The circuit-to-Hamiltonian construction $H^{\prime}=H_{\text {in }}+H_{\text {prop }}+H_{\text {out }}+H_{\text {stab }}$ has the following property.
Fact $1([$ KSV02 $])$. Let $H^{\prime}=H_{\text {in }}+H_{\text {prop }}+H_{\text {out }}+H_{\text {stab }}$. Then,

- if $U$ accepts $x$ with probability at least $\alpha$, then $\lambda_{H^{\prime}} \leq \frac{1-\alpha}{m+1}=: \alpha^{\prime}$,
- if $U$ accepts $x$ with probability at most $\beta$, then $\lambda_{H^{\prime}} \geq \Omega\left(\frac{1-\sqrt{\beta}}{m^{3}}\right)=: \beta^{\prime}$.

We state a helpful lemma lower bounding the non-zero eigenvalues on the frustration-free component of this construction.
Lemma 3.1 (Lemma 2.2 of [GY19] (based on ${ }^{21}$ Lemma 3 of [GK12c])). The smallest non-zero eigenvalue of $H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}$ is at least $\pi^{2} /\left(64 m^{3}\right)$ for $m \geq 1$.

Since $\alpha=1-2^{-n}$ and $\beta=2^{-n}$, we have that estimating $\lambda_{H^{\prime}}$ to inverse-polynomial precision suffices to decide $x$. In the YES case, the upper bound $\alpha^{\prime}$ will be attained via the history state, which we now state:

$$
\begin{equation*}
\left|\psi_{\text {hist }}\right\rangle=\frac{1}{\sqrt{m+1}} \sum_{t=0}^{m} U_{t} \cdots U_{1}|x\rangle_{A}|0 \cdots 0\rangle_{B}|t\rangle_{C} . \tag{5}
\end{equation*}
$$

Our construction. The rub is that $\mathrm{GLH}^{*}$ also requires us to produce, as part of the reduction, a semi-classical state $|u\rangle$ (as defined in Section (1.2) with $\| \Pi_{H^{\prime}}|u\rangle \| \geq \delta$. Unfortunately, in the YES case, if $\alpha \neq 1$ then the history state $\left|\psi_{\text {hist }}\right\rangle$ is not necessarily an actual ground state of $H^{\prime}$. Even worse, the NO case analysis goes via Kitaev's Geometric Lemma [KSV02], which says nothing about the ground state of $H$ (i.e., we do not even have the history state to guide us).

To get around this, we apply three tricks. First, we apply a standard trick of weighing $H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}$ by a large polynomial pre-factor $\Delta$, to allow us to argue about the structure of the low energy space of $H^{\prime}$. Second, we "pre-idle" the verifier, by which we mean we update $U$ to a $M:=(m+N)$-gate circuit, where the first $N$ gates are $I$, and the last $m$ gates are the original circuit $U$. (In contrast, "idling" usually refers to the same trick applied at the end of $U$, so as to increase the amplitude on time step $m$ of the history state. Here, we increase the weight on step 0.) At this point, we have

$$
\begin{equation*}
H^{\prime}:=\Delta H_{\text {in }}+\Delta H_{\text {prop }}+\Delta H_{\text {stab }}+H_{\text {out }} \tag{6}
\end{equation*}
$$

for $H_{\text {prop }}$ using our new $M$-step $U$, and where we update our definitions of $\alpha^{\prime}$ and $\beta^{\prime}$ to

$$
\alpha^{\prime}:=\frac{1-\alpha}{M+1} \quad \text { and } \quad \beta^{\prime}:=\Omega\left(\frac{1-\sqrt{\beta}}{M^{3}}\right) .
$$

[^10]Third, we take inspiration from Ambainis' QMA query gadget construction [Amb14], and define:

$$
\begin{aligned}
H & :=\frac{\alpha^{\prime}+\beta^{\prime}}{2} I_{A B C} \otimes|0\rangle\left\langle\left. 0\right|_{D}+H_{A B C}^{\prime} \otimes \mid 1\right\rangle\left\langle\left. 1\right|_{D^{\prime}}\right. \\
|u\rangle & :=|x\rangle_{A}|0 \cdots 0\rangle_{B}\left(\frac{1}{\sqrt{N}} \sum_{t=1}^{N}|t\rangle\right)_{C}|+\rangle_{D} .
\end{aligned}
$$

Since we will shortly set $N \in \operatorname{poly}(n),|u\rangle$ is semi-classical (i.e., an equally weighted superposition of poly ( $n$ ) strings), as desired. We remark that a similar idea of using block-encoding relative to a single qubit register (in our context, $D$ ) to force $H$ to encode QMA query answers in $D$ was used by [Amb14], albeit in the unrelated context of class $\mathrm{P}^{\mathrm{QMA}[\log ] \text {. While our aim is }}$ not to study $\mathrm{P}^{\mathrm{QMA}}[\log ]$, this type of gadget nevertheless will allow us to precisely pin down the ground space of $H$ in the NO case. (For clarity, $|u\rangle$ is specific to this paper, i.e., the concept of $|u\rangle$ is not relevant to Amb14].) Note $|u\rangle$ sums up to time step $N$, not $M$, and starts the summation at $t=1$.

Finally, set $k=6, a=\alpha^{\prime}, b=\left(\alpha^{\prime}+\beta^{\prime}\right) / 2$, and any $\delta \in(0,1 / \sqrt{2}-\Omega(1 / \operatorname{poly}(n)))$. Set $N$ as a large enough fixed polynomial so that $N /(2(M+1)) \geq 1 / 4+\Delta^{\prime}$ for some $\Delta^{\prime}$ defined later ( $\Delta^{\prime}$ will scale as $\left.1 / 4-1 / \operatorname{poly}(n)\right)$. We will also set $\Delta>0$ as needed. This completes the construction, which clearly runs in poly $(n)$ time. As an aside, as defined, $H$ will have poly $(n)$ norm, whereas $\mathrm{GLH}^{*}$ asks for $\|H\| \leq 1$. This is easily corrected in the standard manner by dividing $H$ by any polynomial bound on $\|H\|$ (obtained, e.g., by applying the triangle inequality for $\|\cdot\|$ over the local terms of $H$ ). This scales down the thresholds $\alpha^{\prime}, \beta^{\prime}$ (which is fine, since they remain inverse poly gapped), and leaves the eigenvectors invariant, which is important for our claims regarding $\| \Pi_{H}|u\rangle \| \geq \delta$.

Correctness. For clarity, recall $U$ is now a $M$-gate circuit, with $M=m+N$.
YES case. Suppose $x \in \Pi_{\text {yes }}$. We must show that $\lambda_{H} \leq a$ (to satisfy the YES case requirement in $\mathrm{GLH}^{*}$ ) and $\| \Pi_{H}|u\rangle \| \geq \delta$ (to satisfy the guiding state promise in GLH*). For the first claim, since $x \in \Pi_{\text {yes }}, U$ accepts $x$ with probability $\alpha$. Thus, we may upper bound $\lambda_{H}$ using assignment $\left|\psi_{\text {hist }}\right\rangle_{A B C}|1\rangle_{D}$ as follows:

$$
\begin{align*}
\lambda_{H} & \leq\left\langle\left.\psi_{\text {hist }}\right|_{A B C}\left\langle\left. 1\right|_{D} H \mid \psi_{\text {hist }}\right\rangle_{A B C} \mid 1\right\rangle_{D} \\
& =\left\langle\left.\psi_{\text {hist }}\right|_{A B C} H^{\prime} \mid \psi_{\text {hist }}\right\rangle_{A B C} \\
& \leq(1-\alpha) /(M+1) \\
& =\alpha^{\prime}, \tag{7}
\end{align*}
$$

where the second statement follows since $H_{\text {in }}\left|\psi_{\text {hist }}\right\rangle=H_{\text {prop }}\left|\psi_{\text {hist }}\right\rangle=H_{\text {stab }}\left|\psi_{\text {hist }}\right\rangle=0$, and the third by direct calculation (in the process re-deriving the first point of Fact 11). (Recall we updated $m$ in Kitaev's original construction at the start of this discussion to M.) Thus, in this case $\lambda_{H}=\lambda_{H^{\prime}} \leq a$, as desired.

To next show $\| \Pi_{H}|u\rangle \| \geq \delta$, observe that since we just showed $\lambda_{H} \leq \alpha^{\prime}$ (Equation (7)), and since $\beta^{\prime}>\alpha^{\prime}$, we have that $\lambda_{H}$ lies in the 1-block relative to $D$, i.e., $\lambda_{H}$ is achieved only by an eigenvector of the form ${ }^{22}|\phi\rangle_{A B C}|1\rangle_{D}$ for some $|\phi\rangle_{A B C}$. We claim that any such $|\phi\rangle_{A B C}$ must

[^11]have large overlap on the history state $\left|\psi_{\text {hist }}\right\rangle_{A B C}$, which will in turn allow us to conclude that $|\phi\rangle_{A B C}$ has large overlap with $|u\rangle$. Since $|\phi\rangle_{A B C}|1\rangle_{D}$ was a ground state of $H$, this will allow us to conclude that $\| \Pi_{H}|u\rangle \| \geq \delta$, i.e., that $|u\rangle$ is a valid guiding state.

To begin, letting $\mathcal{N}$ denote the null space $\operatorname{Null}\left(H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}\right)$, recall that $\mathcal{N}$ is 1dimensional and spanned by $\left|\psi_{\text {hist }}\right\rangle{ }^{23}$ Let $|\phi\rangle$ be an arbitrary ground state of $H^{\prime}$, implying $|\phi\rangle_{A B C}|1\rangle_{D}$ is a ground state of $H$. Write $|\phi\rangle=\eta_{1}\left|\psi_{\text {hist }}\right\rangle+\eta_{2}\left|\psi_{\text {hist }}^{\perp}\right\rangle$ for $\left|\psi_{\text {hist }}\right\rangle \in \mathcal{N},\left|\psi_{\text {hist }}^{\perp}\right\rangle \in \mathcal{N}^{\perp}$, and $\left|\eta_{1}\right|^{2}+\left|\eta_{2}\right|^{2}=1$. Then, since $H^{\prime} \succeq 0$ (since recall $H_{\text {in }}, H_{\text {prop }}, H_{\text {out }}, H_{\text {stab }} \succeq 0$ and $\Delta \geq 0$ ),

$$
\begin{aligned}
\alpha^{\prime} & \geq\langle\phi| H^{\prime}|\phi\rangle \\
& \geq\langle\phi| \Delta\left(H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}\right)|\phi\rangle \\
& =\left|\eta_{1}\right|^{2}\left\langle\psi_{\text {hist }}\right| \Delta\left(H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}\right)\left|\psi_{\text {hist }}\right\rangle+ \\
& \left|\eta_{2}\right|^{2}\left\langle\psi_{\text {hist }}^{\perp}\right| \Delta\left(H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}\right)\left|\psi_{\text {hist }}^{\perp}\right\rangle \\
& =\left|\eta_{2}\right|^{2}\left\langle\psi_{\text {hist }}^{\perp}\right| \Delta\left(H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}\right)\left|\psi_{\text {hist }}^{\perp}\right\rangle \\
& \geq\left|\eta_{2}\right|^{2} \frac{c \Delta}{M^{3}}
\end{aligned}
$$

for $c=\pi^{2} / 64=\Omega(1)$, where the second statement holds since $H_{\text {out }} \succeq 0$ (and thus $H^{\prime} \succeq$ $\left.\Delta\left(H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}\right)\right)$, the third and fourth since $\left(H_{\text {in }}+H_{\text {prop }}+H_{\text {stab }}\right)\left|\psi_{\text {hist }}\right\rangle=0$, and the last by Lemma 3.1. We conclude that

$$
\begin{equation*}
\left|\left(\left\langle\left.\left.\phi\right|_{A B C}\left\langle\left. 1\right|_{D}\right)\left(\left|\psi_{\text {hist }}\right\rangle_{A B C}|1\rangle_{D}\right)\right|^{2} \geq 1-\frac{M^{3} \alpha^{\prime}}{c \Delta}\right.\right.\right. \tag{8}
\end{equation*}
$$

i.e., $|\phi\rangle_{A B C}$ has large overlap with $\left|\psi_{\text {hist }}\right\rangle$, as claimed. At the same time, the following separate calculation shows that $\left|\psi_{\text {hist }}\right\rangle$ also has large overlap with $|u\rangle$ :

$$
\begin{equation*}
\left\lvert\,\left.\left\langle\left.\left. u\right|_{A B C D}\left(\left|\psi_{\text {hist }}\right\rangle_{A B C}|1\rangle_{D}\right)\right|^{2}=\frac{1}{2}\right| \sum_{t=1}^{N} \frac{1}{\sqrt{N(M+1)}}\right|^{2}=\frac{N}{2(M+1)} \geq \frac{1}{4}+\Delta^{\prime} .\right. \tag{9}
\end{equation*}
$$

Via Equations (8) and (9) and the triangle inequality, we obtain our desired claim, i.e., that $|\phi\rangle_{A B C}|1\rangle_{D}$ has large overlap with $|u\rangle_{A B C D}$ :

$$
\begin{equation*}
\left\lvert\,\left(\left.\left\langle\left.\phi\right|_{A B C}\left\langle\left. 1\right|_{D}\right) \mid u\right\rangle_{A B C D}\right|^{2} \geq 1-\frac{1}{4}\left(2 \sqrt{\frac{M^{3} \alpha^{\prime}}{c \Delta}}+\sqrt{3-4 \Delta^{\prime}}\right)^{2} .\right.\right. \tag{10}
\end{equation*}
$$

Here, we have used the known identity $\||v\rangle\langle v|-|u\rangle\langle u| \|_{\text {tr }}=2 \sqrt{1-|\langle u \mid v\rangle|^{2}}$, and applied the triangle inequality to the trace norm.

Finally, we observe that all parameters in Equation (10) can be set so as to yield $\| \Pi_{H}|u\rangle \| \geq$ $\delta$. Specifically, for sufficiently large fixed polynomial $\Delta$, the term $2 \sqrt{\frac{M^{3} \alpha^{\prime}}{c \Delta}}$ can be made arbitrarily inverse polynomially close to 0 . As for $\Delta^{\prime} \in(0,1 / 4)$, choosing $\Delta^{\prime}=1 / 4-1 / q$ for some sufficiently large polynomial $q$, the term $\sqrt{3-4 \Delta^{\prime}}$ can be made to scale like $\sqrt{2+\epsilon}$ for $\epsilon>0$ any desired inverse polynomial. Taking into account the square, and combining with the $\Delta$ scaling, we conclude that for sufficiently large fixed polynomials $\Delta, \Delta^{\prime}$, the right hand

[^12]side above is at least $1-\frac{1}{4}\left(2+\epsilon^{\prime}\right)$ for any desired inverse polynomial $\epsilon^{\prime}$. By choosing $\epsilon^{\prime}$ small enough, this expression is at least $\delta^{2}$, from which it follows 24 that $\| \Pi_{H}|u\rangle \| \geq \delta$, as desired.

NO case. Suppose $x \in \Pi_{\text {no }}$. We must show that $\lambda_{H} \geq b$ (to satisfy the NO case requirement in $\mathrm{GLH}^{*}$ ) and $\| \Pi_{H}|u\rangle \| \geq \delta$ (to satisfy the guiding state promise in $\mathrm{GLH}^{*}$ ). For the first claim, observe first that

$$
H^{\prime} \succeq H_{\mathrm{in}}+H_{\mathrm{prop}}+H_{\mathrm{stab}}+H_{\mathrm{out}} \succeq \Omega\left(\frac{1-\sqrt{\beta}}{M^{3}}\right)=\beta^{\prime}
$$

where the first inequality follows since $H_{\text {in }}, H_{\text {prop }}, H_{\text {stab }} \succeq 0$ and $\Delta \geq 1$, and the second from Fact 1 Since $\beta^{\prime}>\left(\alpha^{\prime}+\beta^{\prime}\right) / 2$, we conclude that $\lambda_{H}$ lies in the 0 -block relative to $D$, implying

$$
\lambda_{H}=\frac{\alpha^{\prime}+\beta^{\prime}}{2}=b
$$

as desired. This also now yields the second claim regarding $|u\rangle$ - by construction every vector in the 0 -block is a ground state of $H$, including unit vector

$$
|x\rangle_{A}|0 \cdots 0\rangle_{B}\left(\frac{1}{\sqrt{N}} \sum_{t=1}^{N}|t\rangle\right)_{C}|0\rangle_{D}
$$

which has overlap $1 / \sqrt{2}$ with $|u\rangle$, as claimed.

## 4 Dequantizing the Quantum Singular Value Transformation

In this section we show how to dequantize quantum algorithms for the Singular Value Transformation when the input matrices are sparse, and when dealing with constant precision. In Section 4.1 we present our main dequantization technique. In Section 4.2 we apply this technique to construct a classical algorithm for estimating the singular values of a matrix given sampling-access to a state that has some overlap with the corresponding subspace. Finally, in Section 4.3 we apply this singular value estimation algorithm to prove Theorem 1.1 .

In this section, the notation $O^{*}(\cdot)$ removes from the complexities factors polynomial in $\log (M N)$, which mainly arise from the cost of accessing vectors and matrices using queryaccess and sampling-access as defined in Definitions 2.2, 2.3 and 2.4.

### 4.1 Main technique

In this subsection we describe our main technical tool. Concretely, we consider the following problem, which we denote by $\operatorname{SVT}(s, \epsilon, \zeta)$, for an integer $s \geq 2$ and two parameters $\epsilon \in(0,1]$, $\zeta \in[0,1)$

[^13]$\operatorname{SVT}(s, \epsilon, \zeta)$
Input: query-access to an $s$-sparse matrix $A \in \mathbb{C}^{M \times N}$ with $\|A\| \leq 1$ query-access to a vector $u \in \mathbb{C}^{N}$ such that $\|u\| \leq 1$ $\zeta$-sampling-access to a vector $v \in \mathbb{C}^{N}$ such that $\|v\| \leq 1$ even polynomial $P \in \mathbb{R}[x]$ of degree $2 d$ with $|P(x)| \leq 1 \quad \forall x \in[-1,1]$
Output: an estimate $\hat{z} \in \mathbb{C}$ such that $\left|\hat{z}-v^{\dagger} P\left(\sqrt{A^{\dagger} A}\right) u\right| \leq \epsilon$

We show the following result.
Theorem 4.1 (Formal version of Theorem 1.3). For any $s \geq 2$ and any $\epsilon \in(0,1]$, the problem $\operatorname{SVT}(s, \epsilon, \zeta)$ can be solved classically with probability at least $1-1 / \operatorname{poly}(N)$ in $O^{*}\left(s^{2 d} / \epsilon^{2}\right)$ time for any $\zeta \leq \epsilon / 8$.

To prove Theorem4.1, we will use the following lemma.
Lemma 4.2. Let $P \in \mathbb{R}[x]$ be an even polynomial of degree $2 d$. There is a $O^{*}\left(s^{2 d}\right)$-time classical procedure that given

- query-access to an s-sparse matrix $A \in \mathbb{C}^{M \times N}$,
- query-access to a vector $u \in \mathbb{C}^{N}$,
- an index $i \in\{1, \ldots, N\}$,
outputs the $i$-th entry of $P\left(\sqrt{A^{+} A}\right) u$.
Proof. Let us write

$$
P(x)=a_{0}+a_{2} x^{2}+a_{4} x^{4}+\cdots+a_{2 d} x^{2 d}
$$

where $2 d$ is the degree of $P$ and $a_{0}, \ldots, a_{2 d} \in \mathbb{R}$. We have

$$
P\left(\sqrt{A^{\dagger} A}\right)=a_{0} I_{2^{n}}+a_{2} A^{\dagger} A+a_{4}\left(A^{\dagger} A\right)^{2}+\cdots+a_{2 d}\left(A^{\dagger} A\right)^{d}
$$

and thus

$$
P\left(\sqrt{A^{\dagger} A}\right) u=a_{0} u+a_{2} A^{\dagger} A u+a_{4}\left(A^{\dagger} A\right)^{2} u+\cdots+a_{2 d}\left(A^{\dagger} A\right)^{d} u .
$$

We show below how to compute the $i$-th entry of each term $\left(A^{\dagger} A\right)^{r} u$, for $r \in\{1, \ldots, d\}$.

The procedure. Since $\left(A^{\dagger} A\right)^{r} u$ involves products of $r$ matrices, we give a more general recursive procedure that receives as input

- query-access to $s$-sparse matrices $B^{[1]}, \ldots, B^{[r]}$, where $B^{[j]} \in \mathbb{C}^{M_{j} \times N_{j}}$ for each $j \in\{1, \ldots, r\}$, which satisfy the following conditions: $N_{j}=M_{j+1}$ for each $j \in\{1, \ldots, r-1\}$, and $N_{r}=$ $N$;
- query-access to a vector $u \in \mathbb{C}^{N}$;
- an index $i \in\{1, \ldots, N\}$,
and outputs the $i$-th entry of the vector $B^{[1]} \cdots B^{[r]} u$.
The procedure repeatedly applies the following idea to recursively decompose the product $B^{[1]} \ldots B^{[r]}$ : to obtain the $i$ th entry of $B^{[1]} \ldots B^{[r]} u$, it suffices to know only the $s$ non-zero entries of row $i$ of $B^{[1]}$ (which can be queried directly), together with the corresponding $s$ entries in column vector $B^{[2]} \ldots B^{[r]} u$ (which are computed recursively).

Formally, the procedure first queries all the entries of the $i$-th row of $B^{[1]}$, and gets at most $s$ non-zero entries and their positions, which we write $j_{1}, \ldots, j_{\ell}$, with $\ell \leq s$. The procedure then computes recursively the $j_{q}$-th entry of the product $B^{[2]} \ldots B^{[r]} u$, for each $q \in\{1, \ldots, \ell\}$. Let $y_{q}$ denote the value obtained, for each $q \in\{1, \ldots, \ell\}$. The procedure finally outputs the value

$$
B_{i, j_{1}}^{[1]} y_{1}+\cdots+B_{i, j_{\ell}}^{[1]} y_{\ell}
$$

which is equal to the $i$-th entry of the vector $B^{[1]} \cdots B^{[r]} u$.

Runtime. Let $\mathcal{T}(r)$ denote the running time of this procedure. We have

$$
\mathcal{T}(r) \leq s \cdot O(\operatorname{poly}(\log (M N)))+s \mathcal{T}(r-1)+s
$$

and thus $\mathcal{T}(r)=O\left(s^{r} \cdot \operatorname{poly}(\log (M N))\right)$.
The complexity of the whole procedure is

$$
O\left(\left(s^{2}+\cdots+s^{2 d}\right) \cdot \operatorname{poly}(\log (M N))\right)=O^{*}\left(s^{2 d}\right)
$$

as claimed.
Proof of Theorem 4.1. To compute an estimate of $v^{\dagger} P\left(\sqrt{A^{\dagger} A}\right) u$ we will use the following approach (similar to Proposition 4.2 in [Tan19]). Remember that we have sampling-access to $v \in \mathbb{C}^{2^{n}}$. Let $\left(\mathcal{Q}_{v}, \mathcal{S} \mathcal{Q}_{v}, m\right)$ denote the triple from Definition 2.4. Let us write $w=P\left(\sqrt{A^{\dagger} A}\right) u$. Consider the following procedure:

1. Using Algorithm $\mathcal{S} \mathcal{Q}_{v}$, sample one index $j \in\{1, \ldots, N\}$;
2. Using Algorithm $\mathcal{Q}_{v}$, query the corresponding entry $v_{j}$;
3. Compute the value $w_{j}$ using the algorithm of Lemma 4.2,
4. Output the value $\frac{w_{j} m^{2}}{v_{j}}$.

Let $X$ denote the complex random variable corresponding to the output of this procedure. Let real $(X)$ and $\operatorname{imag}(X)$ denote its real part and complex part, respectively. We have

$$
\mathbb{E}[X]=\sum_{i=1}^{N} \frac{w_{i} m^{2}}{v_{i}} p(i)
$$

where $m \in[(1-\zeta)\|v\|,(1+\zeta)\|v\|]$ and

$$
p(i) \in\left[(1-\zeta) \frac{\left|v_{i}\right|^{2}}{\|v\|^{2}},(1+\zeta) \frac{\left|v_{i}\right|^{2}}{\|v\|^{2}}\right]
$$

for each $i \in\{1, \ldots, N\}$. Note that the inequalities

$$
\begin{aligned}
& m^{2} p(i) \leq(1+\zeta)^{3}\left|v_{i}\right|^{2} \leq(1+7 \zeta)\left|v_{i}\right|^{2}, \\
& m^{2} p(i) \geq(1-\zeta)^{3}\left|v_{i}\right|^{2} \geq(1-4 \zeta)\left|v_{i}\right|^{2},
\end{aligned}
$$

and thus also

$$
\left|m^{2} p(i)-\left|v_{i}\right|^{2}\right| \leq 7 \zeta\left|v_{i}\right|^{2}
$$

hold for each $i \in\{1, \ldots, N\}$. We thus have

$$
\begin{aligned}
\left|\mathbb{E}[X]-v^{\dagger} P\left(\sqrt{A^{\dagger} A}\right) u\right|=\left|\sum_{i=1}^{N} \frac{w_{i} m^{2}}{v_{i}} p(i)-\sum_{i=1}^{N} v_{i}^{*} w_{i}\right| & =\left|\sum_{i=1}^{N} \frac{w_{i}}{v_{i}}\left(m^{2} p(i)-\left|v_{i}\right|^{2}\right)\right| \\
& \leq \sum_{i=1}^{N}\left|\frac{w_{i}}{v_{i}}\right| \cdot\left|m^{2} p(i)-\left|v_{i}\right|^{2}\right| \\
& \leq 7 \zeta \sum_{i=1}^{N}\left|w_{i}\right|\left|v_{i}\right| \\
& \leq 7 \zeta\|w\|\|v\| \\
& \leq 7 \zeta\left\|P\left(\sqrt{A^{\dagger} A}\right)\right\|\|u\|\|v\| \\
& \leq 7 \zeta .
\end{aligned}
$$

We have

$$
\begin{aligned}
\operatorname{var}[\operatorname{real}(X)] & =\mathbb{E}\left[\text { real }(X)^{2}\right]-\mathbb{E}[\operatorname{real}(X)]^{2} \\
& \leq \mathbb{E}\left[\operatorname{real}(X)^{2}\right] \\
& \leq \sum_{i=1}^{N}\left(\frac{\left|w_{i}\right| m^{2}}{\left|v_{i}\right|}\right)^{2} p(i) \\
& \leq(1+7 \zeta) \sum_{i=1}^{N}\left|w_{i}\right|^{2} m^{2} \\
& \leq(1+7 \zeta)(1+\zeta)^{2}\|w\|^{2}\|v\|^{2} \\
& \leq(1+7 \zeta)^{2}
\end{aligned}
$$

and similarly $\operatorname{var}[\operatorname{imag}(X)] \leq(1+7 \zeta)^{2}$. This procedure can be implemented in

$$
O(\operatorname{poly}(\log N))+O^{*}\left(s^{2 d}\right)=O^{*}\left(s^{2 d}\right)
$$

time.
Apply the above procedure $r$ times, each time getting a complex number $X_{1}, \ldots, X_{r}$, and output the mean. Let $Z=\frac{X_{1}+\cdots+X_{r}}{r}$ denote the corresponding complex random variable, and let real $(Z)$ and $\operatorname{imag}(Z)$ denote its real part and complex part, respectively. Since the variables $X_{1}, \ldots, X_{r}$ are independent, we have

$$
\begin{aligned}
\mathbb{E}[\operatorname{real}(Z)] & =\mathbb{E}[\operatorname{real}(X)], \\
\operatorname{var}[\operatorname{real}(Z)] & =\frac{1}{r^{2}}\left(\operatorname{var}\left[\operatorname{real}\left(X_{1}\right)\right]+\cdots+\operatorname{var}\left[\operatorname{real}\left(X_{r}\right)\right]\right)=\frac{\operatorname{var}[\operatorname{real}(X)]}{r} \leq(1+7 \zeta)^{2} / r,
\end{aligned}
$$

and similarly $E[\operatorname{imag}(Z)]=\mathbb{E}[i \operatorname{mag}(X)]$ and $\operatorname{var}[\operatorname{imag}(Z)] \leq(1+7 \zeta)^{2} / r$. By Chebyshev's inequality we obtain

$$
\begin{aligned}
\operatorname{Pr}\left[|\operatorname{real}(Z)-\mathbb{E}[\operatorname{real}(X)]| \geq \frac{\epsilon-7 \zeta}{\sqrt{2}}\right] & \leq \frac{2 \cdot \operatorname{var}[\operatorname{real}(Z)]}{(\epsilon-7 \zeta)^{2}} \leq \frac{2(1+7 \zeta)^{2}}{r(\epsilon-7 \zeta)^{2}}, \\
\operatorname{Pr}\left[|\operatorname{imag}(Z)-\mathbb{E}[\operatorname{imag}(X)]| \geq \frac{\epsilon-7 \zeta}{\sqrt{2}}\right] & \leq \frac{2 \cdot \operatorname{var}[\operatorname{imag}(Z)]}{(\epsilon-7 \zeta)^{2}} \leq \frac{2(1+7 \zeta)^{2}}{r(\epsilon-7 \zeta)^{2}},
\end{aligned}
$$

and thus

$$
\begin{aligned}
\operatorname{Pr}\left[\left|Z-v^{\dagger} P\left(\sqrt{A^{\dagger} A}\right) u\right| \leq \epsilon\right] \geq & \operatorname{Pr}[|Z-\mathbb{E}[X]| \leq \epsilon-7 \zeta] \\
\geq & 1-\operatorname{Pr}\left[|\operatorname{real}(Z)-\mathbb{E}[\operatorname{real}(X)]| \geq \frac{\epsilon-7 \zeta}{\sqrt{2}}\right] \\
& -\operatorname{Pr}\left[|\operatorname{imag}(Z)-\mathbb{E}[\operatorname{imag}(X)]| \geq \frac{\epsilon-7 \zeta}{\sqrt{2}}\right] \\
\geq & 1-\frac{4(1+7 \zeta)^{2}}{r(\epsilon-7 \zeta)^{2}} .
\end{aligned}
$$

For $\zeta \leq \epsilon / 8$, taking $r=\Theta\left(1 / \epsilon^{2}\right)$ guarantees that the estimate is an $\epsilon$-additive-approximation of $v^{\dagger} P\left(\sqrt{A^{\dagger} A}\right) u$ with probability at least $3 / 4$. Using Lemma 2.1 for each the real part and the imaginary part of the estimate, the success probability can then be improved to $1-1 / \operatorname{poly}(N)$ by repeating this process $O(\log N)$ times.

### 4.2 Singular Value Estimation by dequantizing the QSVT

In this subsection we consider the following problem, which we denote $\operatorname{SVE}\left(s, t_{1}, t_{2}, \theta_{1}, \theta_{2}, \delta, \zeta\right)$, for an integer $s \geq 2$, real numbers $t_{1}, t_{2}, \theta_{1}, \theta_{2} \in[0,1]$ satisfying the inequalities $\theta_{1} \leq t_{1}<t_{2} \leq$ $1-\theta_{2}$, and parameters $\delta \in(0,1]$ and $\zeta \in[0,1)$. This problem asks to decide if there exist singular values in a given interval, given sampling-access to a state that has some overlap with the corresponding subspace.
$\operatorname{SVE}\left(s, t_{1}, t_{2}, \theta_{1}, \theta_{2}, \delta, \zeta\right) \quad$ (singular value estimation)
Input: query-access to an $s$-sparse matrix $A \in \mathbb{C}^{M \times N}$ with $\|A\| \leq 1$
$\zeta$-sampling-access to a vector $u \in \mathbb{C}^{N}$ such that $\|u\| \leq 1$
Promise: (i) $A$ has a singular value in the interval $\left[t_{1}, t_{2}\right]$, and $\left\|\Pi_{A}^{\left[t_{1}, t_{2}\right]} u\right\| \geq \delta$ holds, or (ii) $A$ has no singular value in the interval $\left(t_{1}-\theta_{1}, t_{2}+\theta_{2}\right)$

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

Note that the condition about $u$ is only imposed in the case (i): for the case (ii) we do not require $u$ to have any overlap with a specified subspace.

Our main result is the following theorem.
Theorem 4.3 (Formal version of Theorem 1.4). For any $t_{1}, t_{2}, \theta_{1}, \theta_{2} \in[0,1]$ such that $\theta_{1} \leq t_{1}<$ $t_{2} \leq 1-\theta_{2}$, any $\delta \in(0,1]$ and any $s \geq 2$, the problem $\operatorname{SVE}\left(s, t_{1}, t_{2}, \theta_{1}, \theta_{2}, \delta, \zeta\right)$ can be solved classically
with probability at least $1-1 / \operatorname{poly}(N)$ in

$$
O^{*}\left(\frac{s^{c \cdot\left(1 / \theta_{1}+1 / \theta_{2}\right) \log (\sqrt{3} / \delta)}}{\delta^{4}}\right)
$$

time, for some universal constant $c$, for any $\zeta \leq \delta^{2} / 56$.
In order to prove Theorem4.3, we will need the following lemma, which can be considered as a variant of Lemma 29 in [GSLW19].

Lemma 4.4. For any $t_{1}, t_{2}, \theta_{1}, \theta_{2} \in[0,1]$ such that $\theta_{1} \leq t_{1}<t_{2} \leq 1-\theta_{2}$ and any $\chi \in(0,1)$, there exists an even polynomial $P \in \mathbb{R}[x]$ of degree

$$
O\left(\left(\frac{1}{\theta_{1}}+\frac{1}{\theta_{2}}\right) \log (1 / \chi)\right)
$$

such that $|P(x)| \leq 1$ for all $x \in[-1,1]$ and:

$$
\begin{cases}P(x) \in[1-\chi, 1] & \text { if } x \in\left[t_{1}, t_{2}\right], \\ P(x) \in[0, \chi] & \text { if } x \in\left[0, t_{1}-\theta_{1}\right] \cup\left[t_{2}+\theta_{2}, 1\right] .\end{cases}
$$

Proof. Low and Chuang [LC17] (see also Lemma 25 in [GSLW19]) showed that for all $\eta>$ $0, \xi \in(0,1 / 2)$, there exists an efficiently computable odd polynomial $P^{\prime} \in \mathbb{R}[x]$ of degree $O(\log (1 / \xi) / \eta)$ such that:

$$
\begin{cases}P^{\prime}(x) \in[-1,1] & \text { for all } x \in[-2,2] \\ P^{\prime}(x) \in[-1,-1+\xi] & \text { if } x \in[-2,-\eta] \\ P^{\prime}(x) \in[1-\xi, 1] & \text { if } x \in[\eta, 2]\end{cases}
$$

Let $P_{1}^{\prime}$ be the polynomial obtained by this construction with $\eta=\theta_{1} / 2$ and $\xi=2 \chi / 5$, and $P_{2}^{\prime}$ be the polynomial obtained by this construction with $\eta=\theta_{2} / 2$ and $\xi=2 \chi / 5$. Define

$$
Q(x)=(1-\xi) \frac{P_{1}^{\prime}\left(x-t_{1}+\theta_{1} / 2\right)+P_{2}^{\prime}\left(-x+t_{2}+\theta_{2} / 2\right)}{2}+\xi .
$$

This polynomial satisfies the following properties:

- for all $x \in\left[t_{1}, t_{2}\right]$,

$$
Q(x) \in\left[(1-\xi) \frac{1-\xi+1-\xi}{2}+\xi,(1-\xi) \frac{1+1}{2}+\xi\right] \subseteq[1-\xi, 1] ;
$$

- for all $x \in\left[0, t_{1}-\theta_{1}\right] \cup\left[t_{2}+\theta_{2}, 1\right]$,

$$
Q(x) \in\left[(1-\xi) \frac{1-\xi-1}{2}+\xi,(1-\xi) \frac{1+(-1+\xi)}{2}+\xi\right] \subseteq[0,3 \xi / 2] .
$$

Additionally, using similar calculations, it can be checked that $Q(x) \in[0,1]$ for all $x \in[0,1]$, and $Q(x) \in[0, \xi]$ for all $x \in[-1,0]$.

Finally, set

$$
P(x)=\frac{Q(x)+Q(-x)}{1+\xi} .
$$

This is an even polynomial of degree $O\left(\left(1 / \theta_{1}+1 / \theta_{2}\right) \log (1 / \chi)\right)$, which satisfies the conditions

- $P(x) \in[0,1]$ for all $x \in[-1,1]$;
- for all $x \in\left[t_{1}, t_{2}\right]$,

$$
P(x) \in\left[\frac{1-\xi+0}{1+\xi}, \frac{1+\xi}{1+\xi}\right] \subseteq[1-2 \xi, 1] \subseteq[1-\chi, 1] ;
$$

- for all $x \in\left[0, t_{1}-\theta_{1}\right] \cup\left[t_{2}+\theta_{2}, 1\right]$,

$$
P(x) \in\left[\frac{0+0}{1+\xi^{\prime}}, \frac{3 \xi / 2+\xi}{1+\xi}\right] \subseteq[0,5 \xi / 2]=[0, \chi] ;
$$

as required.
Proof of Theorem 4.3. Let $P$ denote the polynomial from Lemma 4.4 for the value $\chi=\delta^{2} / 3$. Observe that the degree of $P$ is $O\left(\left(1 / \theta_{1}+1 / \theta_{2}\right) \log (\sqrt{3} / \delta)\right)$.

Let

$$
A=\sum_{i=1}^{N} \sigma_{i} u_{i} v_{i}^{\dagger}
$$

be a singular value decomposition of $A$. Let $\Lambda_{\left[t_{1}, t_{2}\right]}$ be the set of indices $i \in\{1, \ldots, N\}$ such that $\sigma_{i} \in\left[t_{1}, t_{2}\right]$.

Let us decompose $u$ in the basis $\left\{v_{i}\right\}$ :

$$
u=\sum_{i=1}^{N} \alpha_{i} v_{i},
$$

where $\alpha_{i} \in \mathbb{C}$ for each $i \in\{1, \ldots, N\}$. Remember that by assumption we have $\sum_{i=1}^{N}\left|\alpha_{i}\right|^{2} \leq 1$. We have

$$
\begin{aligned}
P\left(\sqrt{A^{\dagger} A}\right) u & =\left(\sum_{i=1}^{N} P\left(\sigma_{i}\right) v_{i} v_{i}^{\dagger}\right)\left(\sum_{i=1}^{N} \alpha_{i} v_{i}\right) \\
& =\sum_{i=1}^{N} P\left(\sigma_{i}\right) \alpha_{i} v_{i} .
\end{aligned}
$$

Observe that $u^{\dagger} P\left(\sqrt{A^{\dagger} A}\right) u=\sum_{i=1}^{N} P\left(\sigma_{i}\right)\left|\alpha_{i}\right|^{2}$ is a real number.
If $A$ has a singular value in the interval $\left[t_{1}, t_{2}\right]$, then the promise guarantees that the inequality

$$
\sum_{i \in \Lambda_{\left[t_{1}, t_{2}\right]}}\left|\alpha_{i}\right|^{2} \geq \delta^{2}
$$

holds. We thus have

$$
u^{+} P\left(\sqrt{A^{+} A}\right) u=\sum_{i=1}^{N} P\left(\sigma_{i}\right)\left|\alpha_{i}\right|^{2} \geq \sum_{i \in \Lambda_{\left[t_{1}, t_{2}\right]}} P\left(\sigma_{i}\right)\left|\alpha_{i}\right|^{2} \geq \sum_{i \in \Lambda_{\left[\left[_{1}, t_{2}\right]\right.}}(1-\chi)\left|\alpha_{i}\right|^{2} \geq(1-\chi) \delta^{2} \geq 2 \delta^{2} / 3
$$

If $A$ has no singular value in the interval $\left(t_{1}-\theta_{1}, t_{2}+\theta_{2}\right)$, then we have

$$
u^{\dagger} P\left(\sqrt{A^{\dagger} A}\right) u=\sum_{i=1}^{N} P\left(\sigma_{i}\right)\left|\alpha_{i}\right|^{2} \leq \sum_{i=1}^{N} \chi\left|\alpha_{i}\right|^{2} \leq \chi \leq \delta^{2} / 3 .
$$

Using the algorithm $\operatorname{SVT}(s, \epsilon, \zeta)$ on input $(A, u, u, P)$ with $\epsilon=\delta^{2} / 7$, Theorem4.1 guarantees that for $\zeta \leq \epsilon / 8=\delta^{2} / 56$, we can distinguish with probability at least $1-1 / \operatorname{poly}(N)$ between the two cases in

$$
O^{*}\left(\frac{s^{O\left(\left(1 / \theta_{1}+1 / \theta_{2}\right) \log (\sqrt{3} / \delta)\right)}}{\delta^{4}}\right)
$$

time.

### 4.3 Classical estimation of the ground state energy of local Hamiltonians

In this subsection we show how the results from Section 4.2 can be used to prove Theorem 1.1
First, we introduce a "gapped decision version" of the problem GLH ${ }^{\text {est }}$, which we write $\operatorname{GLH}(k, a, b, \delta)$, for an integer $k \geq 2$, two real numbers $a, b \in[-1,1]$ such that $a<b$, and $\delta \in(0,1]$. This problem is defined as follows.

```
\(\operatorname{GLH}(k, a, b, \delta) \quad\) (guided local Hamiltonian, gapped decision version)
    Input: a \(k\)-local Hamiltonian \(H\) acting on \(n\) qubits such that \(\|H\| \leq 1\)
        sampling-access to a vector \(u \in \mathbb{C}^{2^{n}}\) such that \(\|u\| \leq 1\)
Promises: (i) \(\left\|\Pi_{H} u\right\| \geq \delta\)
(ii) either \(\lambda_{H} \leq a\) or \(\lambda_{H} \geq b\) holds
Goal: decide which of \(\lambda_{H} \leq a\) or \(\lambda_{H} \geq b\) holds
```

The following result follows easily from Theorem 4.3,
Proposition 4.5. For any constants $a, b \in[-1,1]$ such that $a<b$, any constant $\delta \in(0,1]$ and any $k=$ $O(\log n)$, the problem $\operatorname{GLH}(k, a, b, \delta)$ can be solved classically with probability at least $1-1 / \exp (n)$ in poly ( $n$ ) time.

Proof. Let $(H, u)$ denote the input of the problem $\operatorname{GLH}(k, a, b, \delta)$. Since $H$ is a $k$-local Hamiltonian, i.e., $H$ is a sum of $m=\operatorname{poly}(n)$ terms, where each term is a $k$-local Hermitian matrix, $H$ is $m 2^{k}$-sparse. Since $\|H\| \leq 1$, all the eigenvalues of $H$ are in the interval $[-1,1]$. The Hermitian matrix $\frac{H+3 I}{4}$ is thus $\left(m 2^{k}+1\right)$-sparse, and has all its eigenvalues in the interval $[1 / 2,1]$. In order to solve $\operatorname{GLH}(k, a, b, \delta)$, we can simply use the algorithm of Theorem4.3 for $\operatorname{SVE}\left(m 2^{k}+1,1 / 2,(3+a) / 4,1 / 2,(b-a) / 4, \delta, 0\right)$ on input $\left(\frac{H+3 I}{4}, u\right)$.

We then show how Theorem 1.1 follows from Proposition 4.5 ,
Proof of Theorem 1.1. Let $(H, u)$ denote the input of the problem $\mathrm{GLH}^{\text {est }}(k, \epsilon, \delta)$.
For some integer $r=O(1)$ that will be fixed later, decompose the interval $[-1,1]$ into $2 r$ intervals, each of length $1 / r$. The $i$-th interval is $\left[\frac{i-r-1}{r}, \frac{i-r}{r}\right]$, for each $i \in\{1, \ldots, 2 r\}$.

For each $i \in\{1, \ldots, 2 r\}$, we apply the algorithm of Proposition4.5 (which we denote below by Algorithm $\mathcal{C}$ ) to solve the problem $\operatorname{GLH}\left(k, a_{i}, b_{i}, \delta\right)$ on input $(H, u)$ with $a_{i}=\frac{i-r-1}{r}$ and $b_{i}=$ $\frac{i-r}{r}$. With probability at least $1-1 / \exp (n)$, Algorithm $\mathcal{C}$ never errs during the $2 r$ iterations. We assume below that this is the case. Under this assumption, for each $i$ we have:

- if $\lambda_{H} \leq a_{i}$ then $\mathcal{C}$ outputs " $\lambda_{H} \leq a_{i}{ }^{\prime}$,
- if $\lambda_{H} \geq b_{i}$ then $\mathcal{C}$ outputs " $\lambda_{H} \geq b_{i}$ ",
- if $\lambda_{H} \in\left(a_{i}, b_{i}\right)$ then the output of $\mathcal{C}$ is arbitrary (i.e., the output is either " $\lambda_{H} \leq a_{i}$ " or " $\lambda_{H} \geq b_{i}$ ").

One of the following three cases necessarily occurs:
(a) Algorithm $\mathcal{C}$ outputs " $\lambda_{H} \leq a_{i}$ " for all iterations $i \in\{1, \ldots, 2 r\}$,
(b) Algorithm $\mathcal{C}$ outputs " $\lambda_{H} \geq b_{i}$ " for all iterations $i \in\{1, \ldots, 2 r\}$,
(c) there is one value $i_{0} \in\{1, \ldots, 2 r\}$ such that Algorithm $\mathcal{C}$ outputs " $\lambda_{H} \geq b_{i}$ " for all iterations $i=1, \ldots, i_{0}$ and outputs " $\lambda_{H} \leq a_{i}$ " for all iterations $i=\left(i_{0}+1\right), \ldots, 2 r$.

In Case (a) we can conclude that $\lambda_{H} \in[-1,-1+1 / r]$. In Case (b) we can conclude that $\lambda_{H} \in$ $[1-1 / r, 1]$. In Case (c) we can conclude that $\lambda_{H} \in\left[\frac{i_{0}-r-1}{r}, \frac{i_{0}-r+1}{r}\right] \cap[-1,1]$. Taking $r=2 / \epsilon$ thus guarantees that we get an $\epsilon$-additive-approximation of $\lambda_{H}$.

Remark 4.6. While the problem $\mathrm{GLH}^{\text {est }}$ stated in the introduction and the problem GLH stated above assume (for simplicity) perfect sampling-access to the vector u, i.e., 0 -sampling-access, Theorem 1.1 also holds for the version of the problem where $u$ is given via $\zeta$-sampling-access with $\zeta \leq \delta^{2} / 56$, since the algorithm of Theorem 4.3 works for such sampling-access as well.

## 5 Implications to the Quantum PCP Conjecture

In this section we discuss the implications of our dequantization results to the quantum PCP conjecture.

We first define a notion of quantum states that have a short classical description from which efficient sampling-access to the state can be simulated. More precisely, we consider families of complex vectors that can be represented by short binary strings that allow to efficiently generate the triple ( $\left.\mathcal{Q}_{u}, \mathcal{S} \mathcal{Q}_{u}, m\right)$ of Definition 2.4.

Definition 5.1. For a parameter $\zeta \in[0,1]$, we say that a family of complex vectors $\mathcal{F}=\left\{u^{(\ell)}\right\}_{\ell \in \mathbb{N}}$ has a succinct representation allowing $\zeta$-sampling-access if there exist an injective function rep: $\mathcal{F} \rightarrow$ $\{0,1\}^{*}$, a polynomial $q$, and a polynomial-time algorithm $\mathcal{A}$ satisfying the following conditions.

- For all $u \in \mathcal{F}$, the length of the binary string $\operatorname{rep}(u)$ is at most $q(\log (\operatorname{dim} u))$, where $\operatorname{dim} u$ denotes the dimension of $u$.
- Algorithm $\mathcal{A}$ receives as input a binary string. If the string does not correspond to the string $\operatorname{rep}(u)$ for some $u \in \mathcal{F}$, then Algorithm $\mathcal{A}$ outputs an error message. Otherwise Algorithm $\mathcal{A}$ on input rep $(u)$ outputs the classical description ${ }^{25}$ of Algorithms $\mathcal{Q}_{u}$ and $\mathcal{S Q}_{u}$ and a real number $m$ satisfying the specifications in Definition 2.4

If $\zeta=0$, we say that the family has a succinct representation allowing perfect-sampling-access.
Let $\zeta \geq 0$ be a constant. Consider a family of complex vectors $\mathcal{F}$ that has a succinct representation allowing $\zeta$-sampling-access. For any $n \geq 1$, we define

$$
\mathcal{F}_{n}=\left\{u \in \mathcal{F} \mid \operatorname{dim} u=2^{n}\right\},
$$

[^14]the set of vectors of $\mathcal{F}$ that are $2^{n}$-dimensional. We now define a new version of the $k$-local Hamiltonian problem in which we are guaranteed that there exists a vector in $\mathcal{F}$ that has some overlap with the ground state of the Hamiltonian. We denote this problem by $\mathrm{LH}_{\mathcal{F}}(k, a, b, \delta)$, for an integer $k \geq 2$, two real numbers $a, b \in[-1,1]$ such that $a<b$, and $\delta \in(0,1]$. Note that while $\zeta$ does not appear explicitly in the notation, the problem $\mathrm{LH}_{\mathcal{F}}(k, a, b, \delta)$ does depend on this parameter since $\mathcal{F}$ depends on $\zeta$.
$\mathrm{LH}_{\mathcal{F}}(k, a, b, \delta) \quad$ ( $k$-local Hamiltonian problem with samplable approximate ground state)
Input: a $k$-local Hamiltonian $H$ acting on $n$ qubits such that $\|H\| \leq 1$
Promise: (i) there exists a vector $u \in \mathcal{F}_{n}$ such that $\left\|\Pi_{H} u\right\| \geq \delta$
(ii) $\lambda_{H} \leq a$, or $\lambda_{H} \geq b$ holds

Goal: decide which of $\lambda_{H} \leq a$ or $\lambda_{H} \geq b$ holds
A consequence of our dequantization result is the following result.
Theorem 5.2 (Formal version of Theorem [1.5). For any $k=O(\log n)$, any $a, b \in[-1,1]$ such that $b-a=\Omega(1)$, and any constant $\delta>0$ such that $\delta \geq \sqrt{56 \zeta}$, the problem $\mathrm{LH}_{\mathcal{F}}(k, a, b, \delta)$ is in MA.

Note that when $\zeta=0$, i.e., when the vectors in the family $\mathcal{F}$ have a succinct representation allowing perfect sampling-access, the only condition on $\delta$ is that $\delta$ should be a positive constant.

Proof of Theorem 5.2 We show below that the problem can be solved with high probability by a poly $(n)$-time classical verifier that receives a classical witness of length poly $(n)$. Note that since $H$ is a $k$-local Hamiltonian, i.e., $H$ is a sum of $m=\operatorname{poly}(n)$ terms where each term is a $k$-local Hermitian matrix, $H$ is $m 2^{k}$-sparse.

Consider the case where $\lambda_{H} \leq a$. The witness in this case is simply the classical representation rep $(u)$ of a vector $u \in \mathcal{F}_{n}$ such that $\left\|\Pi_{H} u\right\| \geq \delta$. The verifier checks that the representation is valid using Algorithm $\mathcal{A}$ from Definition5.1. If the representation is not valid, then it rejects. Otherwise, the verifier applies the polynomial-time algorithm of Theorem4.3 (which we denote below by Algorithm $\mathcal{B})$ to solve the problem $\operatorname{SVE}\left(m 2^{k}+1,1 / 2,(3+a) / 4,1 / 2,(b-a) / 4, \delta, \zeta\right)$ on input $\left(\frac{H+3 I}{4}, u\right)$ using the triple $\left(\mathcal{Q}_{u}, \mathcal{S} \mathcal{Q}_{u}, m\right)$ obtained by Algorithm $\mathcal{A}$. The verifier accepts if Algorithm $\mathcal{B}$ outputs that $\frac{H+3 I}{4}$ has a singular value in the interval $[1 / 2,(3+a) / 4]$, and rejects otherwise. Theorem 4.3 says the verifier will accept with probability $1-1 / \exp (n)$.

Consider the case where $\lambda_{H} \geq b$. For any binary string received as witness, the verifier will reject unless the witness is of the form $\operatorname{rep}(u)$ for some vector $u \in \mathcal{F}_{n}$. If the witness is of this form, Theorem 4.3 then guarantees that Algorithm $\mathcal{B}$ outputs with probability at least $1-1 / \exp (n)$ that $\frac{H+3 I}{4}$ has no singular value in the interval $(0,(3+b) / 4)$, in which case the verifier rejects.

Finally, using Definition 5.1, we give the formal version of the conjecture introduced in Section 1.4

Conjecture 4 (NLSS conjecture - formal version of Conjecture 3). There exist a family of $O(1)$ local Hamiltonians $\left\{H_{n}\right\}_{n \in \mathbb{N}}$, where each $H_{n}$ acts on $n$ qubits, and a constant $\epsilon>0$ such that for any family of complex vectors $\mathcal{F}$ that has a succinct representation allowing perfect-sampling-access, we have for any sufficiently large $n$ :

$$
u^{\dagger} H_{n} u>\lambda_{H_{n}}+\epsilon
$$

for any unit-norm vector $u \in \mathcal{F}_{n}$.

## Acknowledgments

The authors are grateful to Ryan Babbush, Dominic Berry and András Gilyén for helpful comments about the manuscript, and to anonymous referees for their feedback. FLG was supported by JSPS KAKENHI grants JP19H04066, JP20H05966, JP20H00579, JP20H04139, JP21H04879 and MEXT Quantum Leap Flagship Program (MEXT Q-LEAP) grants JPMXS0118067394, JPMXS0120319794. SG was supported by DFG grants 450041824 and 432788384.

## References

[AA11] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing (STOC 2011), pages 333-342, 2011.
[AAB $\left.{ }^{+} 19\right]$ Frank Arute, Kunal Arya, Ryan Babbush, et al. Quantum supremacy using a programmable superconducting processor. Nature, 574:505-510, 2019.
[Aar09] Scott Aaronson. Why quantum chemistry is hard. Nature Physics, 5:707-708, 2009.
[AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. Guest column: the quantum PCP conjecture. SIGACT News, 44(2):47-79, 2013. ArXiv: 1309.7495.
[ABN23] Anurag Anshu, Nikolas P. Breuckmann, and Chinmay Nirkhe. NLTS Hamiltonians from good quantum codes. In Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023), pages 1090-1096, 2023.
[AGDLHG05] Alán Aspuru-Guzik, Anthony D. Dutoi, Peter J. Love, and Martin HeadGordon. Simulated quantum computation of molecular energies. Science, 309(5741):1704-1707, 2005.
[AL99] Daniel S. Abrams and Seth Lloyd. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. Physical Review Letters, 83:5162-5165, 1999.
[Amb14] Andris Ambainis. On physical problems that are slightly more difficult than QMA. In Proceedings of the 29th IEEE Conference on Computational Complexity (CCC 2014), pages 32-43, 2014.
[ATS03] Dorit Aharonov and Amnon Ta-Shma. Adiabatic quantum state generation and statistical zero knowledge. In Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing (STOC 2003), page 20-29, 2003.
[BBMC20] 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.
[BCC $\left.{ }^{+} 14\right]$ Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Exponential improvement in precision for simulating sparse hamiltonians. In Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing (STOC 2014), page 283-292, 2014.
[Ber23] 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), volume 261 of LIPIcs, pages 20:1-20:18, 2023.
[BFNV19] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. On the complexity and verification of quantum random circuit sampling. Nature Physics, 15:159-163, 2019.
[BH13] Fernando G.S.L. Brandao and Aram W. Harrow. Product-state approximations to quantum ground states. In Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing (STOC 2013), page 871-880, 2013.
[BHMT02] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. Contemporary Mathematics, 305:53-74, 2002.
[BK02] Sergey B. Bravyi and Alexei Yu. Kitaev. Fermionic quantum computation. Annals of Physics, 298(1):210-226, 2002.
[Bra15] Sergey Bravyi. Monte Carlo simulation of stoquastic Hamiltonians. Quantum Information and Computation, 15(13\&14):1122-1140, 2015.
[CAB $\left.{ }^{+} 21\right]$ Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. Nature Reviews Physics, 3:625-644, 2021.
[CGJ19] Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. The power of blockencoded matrix powers: Improved regression techniques via faster hamiltonian simulation. In Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019), volume 132 of LIPIcs, pages 33:1-33:14, 2019.
[CGL ${ }^{+}$20a] Nai-Hui Chia, András Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. In Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2020), pages 387-400, 2020.
[CGL ${ }^{+} 20 \mathrm{~b}$ ] 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), volume 181 of LIPIcs, pages 47:1-47:17, 2020.
[Coo72] Stephen Cook. The complexity of theorem proving procedures. In Proceedings of the 3rd ACM Symposium on Theory of Computing (STOC 1972), pages 151-158, 1972.
[DKBC21] Charles Derby, Joel Klassen, Johannes Bausch, and Toby Cubitt. Compact fermion to qubit mappings. Physical Review B, 104:035118, 2021.
[EA07] Pablo Echenique and José Luis Alonso. A mathematical and computational review of Hartree-Fock SCF methods in quantum chemistry. Molecular Physics, 105(23-24):3057-3098, 2007.
[EH17] Lior Eldar and Aram W. Harrow. Local hamiltonians whose ground states are hard to approximate. In Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS 2017), pages 427-438, 2017.
[FH14] Michael H. Freedman and Matthew B. Hastings. Quantum systems on non-$k$-hyperfinite complexes: a generalization of classical statistical mechanics on expander graphs. Quantum Information and Computation, 14(1-2):144-180, 2014.
[FKV04] Alan M. Frieze, Ravi Kannan, and Santosh S. Vempala. Fast Monte-Carlo algorithms for finding low-rank approximations. Journal of the ACM, 51(6):10251041, 2004.
[GHLS15] 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.
[GK12a] Sevag Gharibian and Julia Kempe. Approximation algorithms for qmacomplete problems. SIAM Journal on Computing, 41(4):1028-1050, 2012.
[GK12b] Sevag Gharibian and Julia Kempe. Hardness of approximation for quantum problems. arXiv:quant-ph/1209.1055, 2012.
[GK12c] Sevag Gharibian and Julia Kempe. Hardness of approximation for quantum problems. In Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP 2012), pages 387-398, 2012.
[GKS16] Alex B. Grilo, Iordanis Kerenidis, and Jamie Sikora. QMA with subset state witnesses. Chicago Journal of Theoretical Computer Science, 2016:4, 2016.
[GL22] Sevag Gharibian and François Le Gall. Dequantizing the quantum singular value transformation: hardness and applications to quantum chemistry and the quantum PCP conjecture. In Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, pages 19-32, 2022. arXiv:2111.09079.
[Gri18] Alex Bredariol Grilo. Quantum proofs, the local Hamiltonian problem and applications. (Preuves quantiques, le problème des Hamiltoniens locaux et applications). PhD thesis, Sorbonne Paris Cité, France, 2018.
[Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In Proceedings of the 28th Annual ACM Symposium on the Theory of Computing, pages 212-219, 1996.
[GSLW19] 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 (STOC 2019), pages 193-204, 2019.
[GST20] András Gilyén, Zhao Song, and Ewin Tang. An improved quantum-inspired algorithm for linear regression. ArXiv:2009.07268, 2020.
[GY19] Sevag Gharibian and Justin Yirka. The complexity of simulating local measurements on quantum systems. Quantum, 3:189, 2019.
[HHL09] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. Physical Review Letters, 103:150502, 2009. Full version available as ArXiv: 0811.3171.
[JGS20] 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), volume 170 of LIPIcs, pages 53:1-53:14, 2020.
[JMBN19] Zhang Jiang, Jarrod McClean, Ryan Babbush, and Hartmut Neven. Majorana loop stabilizer codes for error mitigation in fermionic quantum simulations. Physical Review Applied, 12:064041, 2019.
[JVV86] 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.
[JW28] Pascual Jordan and Eugene Wigner. Über das Paulische Äquivalenzverbot. Zeitschrift für Physik, 47:631-651, 1928.
[Kit95] Alexei Yu. Kitaev. Quantum measurements and the abelian stabilizer problem. arXiv:quant-ph/9511026, 1995.
[KSV02] Alexei Yu. Kitaev, Alexander H. Shen, and Mikhail N. Vyalyi. Classical and Quantum Computation. American Mathematical Society, 2002.
[LBG $\left.{ }^{+} 21\right]$ 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.
[LC17] Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by uniform spectral amplification. ArXiv: 1707.05391, 2017.
[LC19] Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by qubitization. Quantum, 3:163, 2019.
[Lev73] Leonid Levin. Universal search problems. Problems of Information Transmission, 9(3):265-266, 1973.
[Liu21] 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.
[Llo96] Seth Lloyd. Universal quantum simulators. Science, 273(5278):1073-1078, 1996.
[LT20] Lin Lin and Yu Tong. Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems. Quantum, 4:361, 2020.
[MBLAG14] Jarrod R. McClean, Ryan Babbush, Peter J. Love, and Alán Aspuru-Guzik. Exploiting locality in quantum computation for quantum chemistry. The Journal of Physical Chemistry Letters, 5(24):4368-4380, 2014. PMID: 26273989.
[MNRS11] Frédéric Magniez, Ashwin Nayak, Jérémie Roland, and Miklos Santha. Search via quantum walk. SIAM Journal on Computing, 40(1):142-164, 2011.
[Mon15] Ashley Montanaro. Quantum speedup of Monte Carlo methods. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 471(2181):20150301, 2015.
[MRTC21] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. A grand unification of quantum algorithms. arXiv: 2105.02859, 2021.
[Osb12] Tobias J. Osborne. Hamiltonian complexity. Reports on Progress in Physics, 75(2):022001, 2012.
[ $\mathrm{RWC}^{+}$20] Alessandro Rudi, Leonard Wossnig, Carlo Ciliberto, Andrea Rocchetto, Massimiliano Pontil, and Simone Severini. Approximating Hamiltonian dynamics with the Nyström method. Quantum, 4:234, 2020.
[RWS ${ }^{+}$17] 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.
[SBMW19] Kanav Setia, Sergey Bravyi, Antonio Mezzacapo, and James D. Whitfield. Superfast encodings for fermionic quantum simulation. Physical Review Research, 1:033033, 2019.
[SdN13] Martin Schwarz and Maarten Van den Nest. Simulating quantum circuits with sparse output distributions. arXiv: 1310.6749, 2013.
[SV09] Norbert Schuch and Frank Verstraete. Computational complexity of interacting electrons and fundamental limitations of density functional theory. Nature Physics, 5:732-735, 2009.
[SW19] Mark Steudtner and Stephanie Wehner. Quantum codes for quantum simulation of fermions on a square lattice of qubits. Physical Review A, 99:022308, 2019.
[Sze04] Mario Szegedy. Quantum speed-up of Markov chain based algorithms. In Proceedings of the 45 th Symposium on Foundations of Computer Science (FOCS 2004), pages 32-41, 2004.
[Tan19] 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.
[VC05] Frank Verstraete and Juan Ignacio Cirac. Mapping local Hamiltonians of fermions to local Hamiltonians of spins. Journal of Statistical Mechanics: Theory and Experiment, 2005(09):P09012-P09012, 2005.
[VDN11] Maarten Van Den Nest. Simulating quantum computers with probabilistic methods. Quantum Information and Computation, 11(9-10):784-812, 2011.
[WHT16] James D. Whitfield, Vojtěch Havlíček, and Matthias Troyer. Local spin operators for fermion simulations. Physical Review A, 94:030301, 2016.
[WLAG13] James D. Whitfield, Peter J. Love, and Alán Aspuru-Guzik. Computational complexity in electronic structure. Physical Chemistry Chemical Physics, 15:397-411, 2013.
[WZ06] Pawel Wocjan and Shengyu Zhang. Several natural BQP-complete problems. ArXiv: 0606179, 2006.


[^0]:    ${ }^{1}$ More accurately, the runtimes of [CGL ${ }^{+}$20a, JGS20] depend polynomially on the Frobenius norm of $A$, which is small when the rank is small.

[^1]:    ${ }^{2}$ Most, if not all, known quantum many-body systems in nature evolve in time according to local Hamiltonians, which are $2^{n} \times 2^{n}$ Hermitian matrices specified succinctly with poly $(n)$ bits via local constraints, analogous to local clauses for 3-SAT. The spectrum of said Hamiltonians is particularly important, as each distinct eigenvalue encodes an energy level of the system, and each corresponding eigenspace the state of the system at that energy. The smallest eigenvalue and corresponding eigenvector are called the ground state energy and ground state, respectively.
    ${ }^{3}$ The name "phase" estimation arises since all eigenvalues of unitary $U$ are of form $e^{i \theta}$ for $\theta \in \mathbb{R}$, where $\theta$ is called the "phase".
    ${ }^{4}$ We stress that this is a strong assumption. First, from a complexity-theory perspective, Hamiltonians occurring in chemistry can be QMA-complete [SV09]. Thus, a uniformly generated efficient quantum circuit approximating the ground state would yield BQP = QMA, which is highly unlikely. Indeed, even a non-uniformly generated efficient quantum circuit would yield QCMA $=$ QMA, for QCMA defined as QMA but with a classical proof. From a practical perspective, it is also not expected that one can always produce an initial state with at least inversepolynomial overlap with the ground state. In fact, it is possible to contrive instances where the initial state produced by the Hartree-Fock method (described later) has support on the ground state that vanishes exponentially in the number of particles. A classic example is the "Van Vleck catastrophe" (see MBLAG14]).
    ${ }^{5}$ This statement requires clarification. Roughly, the Hartree-Fock method may be viewed as "optimization over a product state ansatz" in the second quantization formalism. More formally, one wishes to compute the Fock state with minimal energy against a given fermionic Hamiltonian. For such approaches, however, it should be clarified that even if a Fock state achieves a good approximation of the ground state energy, said state does not necessarily achieve good overlap with the ground state. This phenomenon appears for qubit-systems as well -

[^2]:    polynomially in the sparsity and Frobenius norm of the Hamiltonian, whereas our runtime is independent of the Frobenius norm (and thus of the rank).
    ${ }^{8}$ This problem has similarities with the "Guided Stoquastic Local Hamiltonian" problem introduced by Bravyi [Bra15]. In the Guided Stoquastic Local Hamiltonian problem the main differences are: (1) only stoquastic Hamiltonians are considered; (2) in the YES case, an "easy guiding state" exists, but in the NO case, no assumption on guiding states is made; (3) the guiding state is not given as an input. The definition of "guiding state" in Bra15 also differs from ours, requiring non-negative amplitudes and "point-wise overlap" with the entries of the ground state, as opposed to a global overlap here.

[^3]:    ${ }^{9}$ For clarity, there is no connection between our terminology "semi-classical" and (e.g.) semiclassical physics.

[^4]:    ${ }^{10}$ Formally, $|A|:=\sqrt{A^{\dagger} A}$ generalizes the absolute value function to matrix $A$, and ensures $|A|$ is positivesemidefinite (which implies $|A|$ is diagonalizable with real, non-negative eigenvalues). The polynomial $P$ is then applied as an "operator function" to $|A|$, i.e., applied to the eigenvalues of $|A|$.

[^5]:    ${ }^{11}$ In contrast, "idling" usually refers to the same trick applied at the end of $U$, so as to increase the amplitude on time step $m$ of the history state.
    ${ }^{12} \mathrm{P}^{\mathrm{QMA}}[\log ]$ is the set of decision problems decidable by a poly-time deterministic Turing machine making logarithmically many queries to a QMA oracle.

[^6]:    ${ }^{13}$ Additionally, as explained in Section [1.2, when a quantum state representing the guiding state can be constructed efficiently (e.g., when the guiding state is a semi-classical state), this problem can be solved in quantum polynomial-time.

[^7]:    ${ }^{14}$ While the statement of Theorem 1.1 given in Section 1.2 is for $k$-local Hamiltonians with $k=O(\log n)$, the same classical simulation result holds more generally for $s$-sparse Hamiltonians with $s=O(\operatorname{poly}(n))$, since the dequantization results of Section 4 apply to such sparse matrices as well.
    ${ }^{15}$ We note that a related approach has been used by Bravyi [Bra15] and Liu [Liu21]. The latter introduced the notion of an "easy witness" for stoquastic local Hamiltonian, and showed that the Stoquastic Local Hamiltonian instance with an easy witness in the YES case can be verified in MA. The approach from Liu21] considers query access and sample access to the witness. This approach actually works for inverse-polynomial precision, whereas ours (Theorem 1.5 works only for constant precision.

[^8]:    ${ }^{16}$ We remark that since the short version of the article was published [GL22], Anshu, Breuckmann, and Nirkhe have proved the NLTS conjecture ABN23]. For clarity, our NLSS conjecture remains open.
    ${ }^{17}$ We thank Matthew Coudron for pointing this out to us.

[^9]:    ${ }^{18}$ All our upper bounds can easily be converted into upper bounds for the bit-complexity of the problem as well if the notions of query-access and sampling-access to the input are redefined appropriately to take in consideration the bit-complexity of the input vectors and matrices.
    ${ }^{19}$ Lemma 6.1 in [JVV86] is actually stated in terms of multiplicative approximation (approximation within an $1+\epsilon$ ratio), rather than additive approximation. The proof, however, only uses Chernoff's bound and thus applies to additive approximation as well.

[^10]:    ${ }^{20}$ In contrast to [KSV02], which studied QMA, here we do not require a separate proof register, as BQP does not take in a proof.
    ${ }^{21}$ Alternatively, see Lemma 23 of the arXiv version GK12b.

[^11]:    ${ }^{22}$ In other words, $H$ is, up to permutation, block diagonal with respect to the standard basis on $D$, and $\lambda_{H}$ lies in the 1-block of this decomposition.

[^12]:    ${ }^{23}$ That $\mathcal{N}$ is spanned by history states is well-known from Kitaev's original analysis [KSV02]; the 1-dimensionality arises since there is no proof register in our setting, i.e., we are reducing from BQP rather than QMA.

[^13]:    ${ }^{24}$ The $\delta^{2}$ drops to $\delta$ here since we want a lower bound on $\| \Pi_{H}|u\rangle \|$, not $\| \Pi_{H}|u\rangle \|^{2}$.

[^14]:    ${ }^{25}$ By a classical representation of $\mathcal{Q}_{u}$ and $\mathcal{S} \mathcal{Q}_{u}$ we mean the description of poly $(\log \operatorname{dim} u)$-size circuits that implement these two algorithms.

