

Improvements in Quantum SDP-Solving with Applications

Joran van Apeldoorn* András Gilyén†

Abstract

Following the first paper on quantum algorithms for SDP-solving by Brandão and Svore [BS17] in 2016, rapid developments has been made on quantum optimization algorithms. In 2017 van Apeldoorn et al. [AGGW17] improved the quantum algorithm introduced by [BS17] and gave stronger lower bounds as well. Recently Brandão et al. [BKL⁺18] improved the quantum SDP-solver in the so-called quantum state input model, where the input matrices of the SDP are given as purified mixed states. They also gave the first non-trivial application of quantum SDP-solving by obtaining a more efficient algorithm for the problem of shadow tomography [Aar17].

In this paper we improve on all previous quantum SDP-solvers. Mainly we construct better Gibbs-samplers for both input models, which directly gives better bounds for SDP-solving. We also combine the Fast Quantum OR lemma of Brandão et al. [BKL⁺18] and the Gentle Quantum Search Lemma of Aaronson [Aar17], and use the techniques from van Apeldoorn et al. [AGGW17] to give an improved general quantum SDP-solving framework. For an SDP with m constraints involving $n \times n$ matrices, our improvements yield an $\tilde{O}((\sqrt{m} + \sqrt{n}\gamma)s\gamma^4)$ upper bound on SDP-solving in the sparse matrix input model and an $\tilde{O}((\sqrt{m} + B^{2.5}\gamma^{3.5})B\gamma^4)$ upper bound in the quantum state input model. Here $\gamma = Rr/\varepsilon$ is the additive error ε scaled down with bounds R and r on the size of optimal solutions, s is the row-sparsity of the input matrices in the sparse matrix input model and B is a normalization factor for the input states in the quantum state model. We also introduce the quantum operator input model, which generalizes both other input models. In this more general model we give an $\tilde{O}((\sqrt{m} + \sqrt{n}\gamma)\alpha\gamma^4)$ -query algorithm, where α is a normalization factor of the input operators.

We then apply these results to the problem of shadow tomography to simultaneously improve the best known upper bounds on sample complexity [Aar17] and complexity [BKL⁺18]. Furthermore, we apply our quantum SDP-solvers to the problems of quantum state discrimination and E-optimal design. In both cases we beat the classical lower bound in terms of some parameters, at the expense of heavy dependence on some other parameters.

Finally we prove two lower bounds for solving SDPs using quantum algorithms: (1) $\tilde{\Omega}(\sqrt{m}B/\varepsilon)$ in the quantum state input model, and (2) $\tilde{\Omega}(\sqrt{m}\alpha/\varepsilon)$ in the quantum operator input model. These lower bounds show that the \sqrt{m} factor and the polynomial dependence on the parameters B , α , and $1/\varepsilon$ are necessary.

*QuSoft, CWI, the Netherlands. Supported by the Netherlands Organization for Scientific Research, grant number 617.001.351. apeldoor@cwi.nl

†QuSoft, CWI, the Netherlands. Supported by ERC Consolidator Grant 615307-QPROGRESS. gilyen@cwi.nl

1 Introduction

1.1 Semidefinite programs

In this paper we consider *Semidefinite programs* (SDPs). SDPs have many applications in optimization, notable examples include approximation of NP-hard problems like MAXCUT [GW95] and polynomial optimization through the Sum-Of-Squares hierarchy [Las01, Par00]. SDPs have also found applications in quantum information theory. Examples include POVM measurement design [Eld03] and finding the winning probability of non-local games [CHTW04].

We consider the basic (primal) form of an SDP as follows:

$$\begin{aligned} \text{OPT} = \max \quad & \text{Tr}(CX) \\ \text{s.t.} \quad & \text{Tr}(A_j X) \leq b_j \quad \text{for all } j \in [m], \\ & X \succeq 0, \end{aligned} \tag{1}$$

where $[m] := \{1, \dots, m\}$. The input to the problem consists of $n \times n$ Hermitian constraint matrices A_1, \dots, A_m , an objective matrix C and reals b_1, \dots, b_m . For normalization purposes we assume $\|C\|, \|A_j\| \leq 1$. The number of constraints is m (we do not count the standard $X \succeq 0$ constraint for this). The variable X of this SDP is an $n \times n$ positive semidefinite (psd) matrix. We assume that $A_1 = I$ and $b_1 = R$, giving a known bound on the trace of a solution: $\text{Tr}(X) \leq R$. A primal SDP also has a *dual*. For a primal SDP of the above form (1) the dual SDP is

$$\begin{aligned} \text{OPT} = \min \quad & b^T y \\ \text{s.t.} \quad & \sum_{j=1}^m y_j A_j - C \succeq 0, \\ & y \geq 0. \end{aligned} \tag{2}$$

We assume that the dual optimum is attained and that an explicit $r \geq 1$ is known such that at least one optimal dual solution y exists $\|y\|_1 \leq r$. These assumptions imply that strong duality holds, justifying the use of OPT for both optimal values. Linear programs (LPs) correspond to the case where all constraint matrices are diagonal.

In this paper we build on the observation that a normalized psd matrix can be naturally represented as a quantum state. Since operations on quantum states can sometimes be cheaper to perform on a quantum computer than operations on classical descriptions of matrices, this can give rise to faster algorithms for solving SDPs on a quantum computer [BS17].

We say an algorithm is an ε -approximate *quantum SDP-solver* if for all input numbers $g \in \mathbb{R}$ and $\zeta \in (0, 1)$, with success probability $1 - \zeta$, all of the following hold:

- The algorithm determines whether $\text{OPT} \leq g - \varepsilon$ or $\text{OPT} \geq g + \varepsilon$. If $\text{OPT} \in [g - \varepsilon, g + \varepsilon]$ then it may output either.
- The algorithm finds a $y \in \mathbb{R}^{m+1}$ that is an ε -feasible solution to the dual problem with objective value at most $g + \varepsilon$, i.e.,

$$\sum_{j=1}^m y_j A_j - C \succeq -\varepsilon I,$$

and $\langle y, b \rangle \leq g + \varepsilon$, or it concludes that no such y exists even if we would set $\varepsilon = 0$.

- The algorithm finds a vector $y' \in \mathbb{R}^{m+1}$ and a real number z such that for

$$\rho := \frac{e^{-\sum_{j=1}^m y'_j A_j + y'_0 C}}{\text{Tr}\left(e^{-\sum_{j=1}^m y'_j A_j + y'_0 C}\right)} \quad (3)$$

we have that $z\rho$ is an ε -feasible primal solution with objective value at least $g - \varepsilon$, i.e.,

$$\forall j \in [m]: \text{Tr}(z\rho A_j) \leq b_j + \varepsilon,$$

and $\text{Tr}(z\rho C) \geq g - \varepsilon$, or concludes that no such z and y' exist even if we would set $\varepsilon = 0$.

Notice that we can easily find an approximation of OPT using binary search on g if we have an ε -approximate SDP-solver. An algorithm that only satisfies the last of the three points will be called an ε -approximate SDP primal oracle. Due to the form of the objective value constraint in this last point, and to simplify statements like (3), we write $A_0 := -C$ and $b_0 := -g$.

In Sections 3.2 and 3.3 we will work with *subnormalized density operators*:

Definition 1 (Subnormalized density operators & Purification). A subnormalized density operator ϱ is a psd matrix of trace at most 1.

A purification of a subnormalized density operator ϱ is a pure state consisting of 3 registers such that tracing out the third register¹ and projecting on the subspace where the second register is $|0\rangle$ yields ϱ .

We write “ ϱ ” and “ ζ ” for subnormalized density operators to distinguish them from normalized density operators, for which we write “ ρ ” and “ σ ”.

Notation. We use the following definition for $\tilde{\mathcal{O}}$:

$$\tilde{\mathcal{O}}_{d,e}(f(a,b,c)) := \mathcal{O}(f(a,b,c) \cdot \text{polylog}(f(a,b,c), d, e)).$$

We define $\tilde{\Omega}$ in a similar way and $\tilde{\Theta}$ as the intersection of the two. We write δ_{ij} for the Kronecker delta function and e_j for the j th basis vector in the standard basis when the dimension of the space is clear from context. For a Hermitian matrix H we write $\text{Spec}(H)$ for its spectrum (set of eigenvalues). For a function $f: \mathbb{R} \rightarrow \mathbb{R}$ we write $f(H)$ for the matrix we get by applying f to the eigenvalues of H , i.e.,

$$f(H) = U \begin{bmatrix} f(\lambda_1) & & \\ & \ddots & \\ & & f(\lambda_n) \end{bmatrix} U^{-1} \text{ where } H = U \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} U^{-1}.$$

1.2 Input models & Subroutines

We will consider three input models: the *sparse matrix model*, the *quantum state model*, and the *quantum operator model*. In all models we assume quantum oracle access to the numbers b_j via the input oracle O_b satisfying² for all $j \in [m]$:

$$O_b|j\rangle|0\rangle = |j\rangle|b_j\rangle.$$

For all input oracles we assume we can apply both the oracle and its inverse³ in a controlled fashion.

¹For simplicity we assume that for a d -dimensional density operator a purification has at most $\text{polylog}(d)$ qubits.

²For simplicity we assume the bitstring representation has at most $\mathcal{O}(\log(nmRr/\varepsilon))$ bits.

³When we talk about samples, e.g. in Section 4.1, then we do not assume we can apply the inverse operation.

Sparse matrix model. In the *sparse matrix model* the input matrices are assumed to be s -row sparse for a known bound $s \in [n]$, meaning that there are at most s non-zero elements per row. Access to the A_j matrices is provided by two oracles, similar to previous work on Hamiltonian simulation in [BCK15]. The first of the two oracles is a unitary O_{sparse} , which serves the purpose of sparse access. This oracle calculates the **index** : $[m] \times [n] \times [s] \rightarrow [n]$ function, which for input (j, k, ℓ) gives the column index of the ℓ th non-zero element in the k th row of A_j . We assume this oracle computes the index “in place”:

$$O_{\text{sparse}}|j, k, \ell\rangle = |j, k, \mathbf{index}(j, k, \ell)\rangle. \quad (4)$$

(In the degenerate case where the k th row has fewer than ℓ non-zero entries, $\mathbf{index}(j, k, \ell)$ is defined to be ℓ together with some special symbol indicating this case.)

We also need another oracle O_A , returning a bitstring² representation of $(A_j)_{ki}$ for every $j \in [m]$ and $k, i \in [n]$:

$$O_A|j, k, i, z\rangle = |j, k, i, z \oplus (A_j)_{ki}\rangle. \quad (5)$$

This model corresponds directly to a classical way of accessing sparse matrices.

Quantum state model. In contrast to the sparse matrix model, the *quantum state model* is inherently quantum and has no classical counterpart for SDPs.⁴ In this model we assume that each A_j has a fixed decomposition of the form

$$A_j = \mu_j^+ \varrho_j^+ - \mu_j^- \varrho_j^- + \mu_j^I I$$

for (subnormalized) density operators ϱ_j^\pm , non-negative reals μ_j^\pm and real number $\mu_j^I \in \mathbb{R}$. We assume access to an oracle O_μ that takes as input an index j and outputs binary representations² of μ_j^+, μ_j^- and μ_j^I .

Furthermore we assume access to a state-preparing oracle $O_{|\cdot\rangle}$ that prepares purifications¹ $|\psi_j^\pm\rangle$ of ϱ_j^\pm :

$$O_{|\cdot\rangle}|j\rangle|\pm\rangle|0\rangle = |j\rangle|\pm\rangle|\psi_j^\pm\rangle.$$

Finally we assume that a bound $B \in \mathbb{R}_+$ is known such that

$$\forall j : \mu_j^+ + \mu_j^- + |\mu_j^I| \leq B.$$

Note that a tight upper bound B can easily be found using $\mathcal{O}(\sqrt{m})$ quantum queries to O_μ by means of maximum finding [DH96].

Quantum operator model. We propose a new input model that we call the *quantum operator model*. In this model the input matrices are given by a unitary that implements a block-encoding:

Definition 2 (Block encoding). *Suppose that A is a w -qubit operator, $\alpha, \varepsilon \in \mathbb{R}_+$ and $k \in \mathbb{N}$, then we say that the $(a + w)$ -qubit unitary U is an (α, a, ε) -block-encoding of A , if*

$$\|A - \alpha(|0\rangle^{\otimes a} \otimes I)U(|0\rangle^{\otimes a} \otimes I)\| \leq \varepsilon.$$

⁴However, there is a natural classical analogue in the case of LPs, when the constraints a_j are given by random variables that outputs k with probability proportional to a_{jk} . In this case a classical algorithm with complexity $\tilde{\mathcal{O}}_n(m \text{poly}(B, \gamma))$ is possible. In a similar manner, the classical input model can be sped up for classical LP solvers as well using techniques similar to those presented in this paper, which would lead to a $\tilde{\mathcal{O}}((n + m) \text{poly}(\gamma))$ algorithm.

Roughly speaking this means that A is represented by a unitary

$$U \approx \begin{pmatrix} A/\alpha & \cdot \\ \cdot & \cdot \end{pmatrix}.$$

In the quantum operator model we assume access to an oracle O_U that acts as follows:

$$O_U|j\rangle|\psi\rangle = |j\rangle(U_j|\psi\rangle).$$

Where U_j is an $(\alpha, a, 0)$ -block-encoding⁵ of A_j , for some fixed⁶ $\alpha \in \mathbb{R}$ and $a = \mathcal{O}(\log(nmRr/\varepsilon))$.

In Section 3.2 we will show that the sparse input model can be reduced to the quantum operator model with $\alpha = s$ and that the quantum state model can be reduced to it with $\alpha = B$. We will also argue that if we can perform a measurement corresponding to $A_j \succeq 0$ using a ancilla qubits, i.e., accept a state ρ with probability $\text{Tr}(A_j\rho)$, then we can implement a $(1, a + 1, 0)$ -block-encoding of A_j .

Computational cost. We will analyze the query complexity of algorithms and subroutines, i.e., the number of queries to controlled versions of the input oracles and their inverses. We will denote the optimal quantum query complexity of an ε -approximate *quantum SDP-solver* with success probability $2/3$ by $T_{SDP}(\varepsilon)$. We only consider success probability $2/3$ to simplify the notation and proofs. However in all cases an ε -approximate SDP-solver with success probability $1 - \zeta$ can easily be constructed using $\mathcal{O}(\log(1/\zeta)T_{SDP}(\varepsilon))$ queries.

In our algorithms we will assume access to a quantum-read/classical-write RAM (QCRAM), and assume one read/write operation has a constant gate complexity⁷; the size of the QCRAM will typically be $\tilde{\mathcal{O}}_{n,m}\left(\left(\frac{Rr}{\varepsilon}\right)^2\right)$ bits. Most often in our results the number of non-query elementary operations, i.e., two-qubit gates and QCRAM calls, matches the query complexity up to polylog factors. In particular, if not otherwise stated, in our results a T -query quantum algorithm uses at most $\tilde{\mathcal{O}}_{n,m}(T)$ elementary operations.

Subroutines. We will work with two major subroutines which need to be implemented according to the specific input model. First, the algorithm will require an implementation of a Gibbs-sampler.

Definition 3 (Gibbs-sampler). *A θ -precise Gibbs-sampler is a unitary that takes as input a data structure storing a vector $y \in \mathbb{R}_{\geq 0}^{m+1}$ and creates as output a purification of a θ -approximation in trace distance of the Gibbs state $e^{-\sum_{j=0}^m y_j A_j} / \text{Tr}\left(e^{-\sum_{j=0}^m y_j A_j}\right)$. If $\|y\|_1 \leq K$ and the support of y has size at most d , then we write $T_{\text{Gibbs}}(K, d, 4\theta)$ for the cost of this unitary.*

For technical reasons we also allow Gibbs-samplers that require a random classical input seed $S \in \{0, 1\}^a$ for some $a = \mathcal{O}(\log(1/\theta))$. In this case the output should be a θ -approximation of the Gibbs state with high probability ($\geq 4/5$) over a uniformly random input seed S .

⁵If n is not a power of 2, then we simply define A_j to be zero on the additional $2^w - n$ dimensions.

⁶Having a single normalization parameter α is not a serious restriction as it is easy to make a block-encoding more subnormalized so that every A_j gets the same normalization, cf. Lemma 14.

⁷Note that read/write operations of a QRAM or QCRAM of size S can be implemented using $\tilde{\mathcal{O}}(S)$ two-qubit gates, so this assumption could hide a factor in the gate complexity which is at most $\tilde{\mathcal{O}}(S)$.

We will use the approximate Gibbs states in order to compute the quantity $\text{Tr}(A_j\rho)$ using a trace estimator.

Definition 4 (Trace estimator). *A (θ, σ) -trace estimator is a unitary that as input takes a state ρ and index j . It outputs a sample from a random variable $X_j \in \mathbb{R}$ such that X_j is a trace estimator that is at most $\theta/4$ biased:*

$$|\text{Tr}(A_j\rho) - \mathbb{E}[X_j]| \leq \theta/4,$$

and the standard deviation of X_j is at most σ . We write $T_{Tr}^\sigma(\theta)$ for the cost of such a unitary.

1.3 Previous work

Classical SDP-solvers roughly fall into two categories: those with logarithmic dependence on R , r and $1/\varepsilon$, and those with polynomial dependence on these parameters but better dependence on m and n . In the first category the best known algorithm [LSW15] at the time of writing has complexity

$$\tilde{O}_{Rr/\varepsilon}(m(m^2 + n^\omega + mns)).$$

where $\omega \in [2, 2.38]$ is the yet unknown exponent of matrix multiplication.

In the second category Arora and Kale [AK16] gave an alternative framework for solving SDPs, using a matrix version of the “multiplicative weights update” method. Their framework can be tuned for specific types of SDPs, allowing for near linear-time algorithms in the case of for example the Goemans-Williamson SDP for the approximation of the maximum cut in a graph [GW95].

In 2016 Brandão and Svore [BS17] used the Arora-Kale framework to implement a general quantum SDP-solver in the sparse matrix model. They observed that the matrix

$$\rho := \frac{e^{-\sum_{j=0}^m y_j A_j}}{\text{Tr}\left(e^{-\sum_{j=0}^m y_j A_j}\right)},$$

that is used for calculations in the Arora-Kale framework is in fact a $\log(n)$ -qubit Gibbs state and can be efficiently prepared as a quantum state on a quantum computer. Using this they achieved a quantum speedup in terms of n . Combining this with a Grover-like speedup allowed for a speedup in terms of m as well, leading to an ε -approximate quantum SDP solver with complexity

$$\tilde{O}\left(\sqrt{mns}^2 \left(\frac{Rr}{\varepsilon}\right)^{32}\right).$$

They also showed an $\Omega(\sqrt{m} + \sqrt{n})$ quantum query lower bound for solving SDPs when all other parameters are constant. This left as open question whether a better lower bound, matching the \sqrt{mn} upper bound, could be found. The upper bound for the sparse input model was subsequently improved by van Apeldoorn et al. [AGGW17] to

$$\tilde{O}\left(\sqrt{mns}^2 \left(\frac{Rr}{\varepsilon}\right)^8\right).$$

van Apeldoorn et al. also gave an $\Omega(\sqrt{\max(n, m)} \min(n, m)^{3/2})$ lower bound, albeit for non-constant parameters R and r . This bound implies that there is no general quantum SDP-solver that has a

$o(nm)$ dependence on n and m and logarithmic dependence on R , r and $1/\varepsilon$. They also showed that every SDP-solver whose efficiency relies on outputting sparse dual solutions (including their algorithm and that of Brandão and Svore [BS17]) is limited, since problems with a lot of symmetry (like maxflow-mincut) in general require non-sparse dual solutions. Furthermore, they showed that for many combinatorial problems (like MAXCUT) R and r increase linearly with n and m .

Very recently Brandão et al. [BKL⁺17] gave an improved SDP-solver for the quantum state input model⁸ that has a complexity bound with logarithmic dependence on n :

$$T_{SDP}(\varepsilon) = \tilde{O}_n \left(\sqrt{m} \operatorname{poly} \left(\frac{Rr}{\varepsilon}, B, \max_{j \in \{0, \dots, m\}} [\operatorname{rank}(A_j)] \right) \right).$$

Brandão et al. also applied their algorithm to the problem of *shadow tomography*, giving the first non-trivial application of a quantum SDP-solver.

Subsequently these results were further improved by the introduction of the Fast Quantum OR lemma by the same authors [BKL⁺18]. Approaches prior to [BKL⁺18] searched for a violated constraint in the SDP using Grover-like techniques, resulting in a multiplicative complexity of Gibbs-sampling and searching. The Fast Quantum OR lemma can be used to separate the search phase from the initial Gibbs-state preparation phase. This led to the improved complexity bound [BKL⁺18] of

$$\tilde{O}_n \left(\left(\sqrt{m} + \operatorname{poly} \left(\max_{j \in \{0 \dots m\}} [\operatorname{rank}(A_j)] \right) \right) \operatorname{poly} \left(\frac{Rr}{\varepsilon}, B \right) \right).$$

Using the Fast Quantum OR Lemma the complexity bound on $T_{SDP}(\varepsilon)$ can be improved in the sparse input model as well, as independently observed by the authors of [BKL⁺18] and by us. We thank the authors of [BKL⁺18] for sending us an early draft of [BKL⁺18] introducing the Fast Quantum OR Lemma, which enabled us to work on these improvements. During the correspondence the application of the OR lemma to the sparse matrix model was independently suggested by Brandão et al. [Wu17] and by us.

1.4 Our results

In this paper we build on the Arora-Kale framework for SDP-solving in a similar fashion as [AGGW17, BS17] and also use results from [BKL⁺17, LRS15] to construct a primal oracle. We improve on the previous results about quantum SDP-solving in three different ways:

- We give a computationally more efficient version of the Gentle Quantum Search Lemma [Aar17] using the Fast Quantum OR Lemma from [BKL⁺18]. We also extend this to minimum finding to get our *Two-Phase Quantum Minimum finding* (Lemma 7). As independently observed by the authors of [BKL⁺18] the Fast Quantum OR Lemma gives a speed-up for SDP primal oracles in general. Moreover, using Two-Phase Quantum Minimum finding, we show how to improve the upper bound on the complexity of general SDP-solving from

$$T_{SDP}(\varepsilon) = \tilde{O}_n \left(\sqrt{m} (T_{Tr}^\sigma(\gamma) T_{Gibbs}(\gamma, \gamma^2, \gamma^{-1})) \gamma^3 \sigma \right) \quad (6)$$

⁸This model was already introduced in the first version of [BS17] together with a similar complexity statement, but there were some unresolved issues in the proof, that were only fixed by the contributions of [BKL⁺17].

as implied in previous work [BS17, AGGW17] to

$$T_{SDP}(\varepsilon) = \tilde{\mathcal{O}}_n((\sqrt{m}T_{Tr}^\sigma(\gamma) + T_{Gibbs}(\gamma, \gamma^2, \gamma^{-1}))\gamma^4\sigma^2), \quad (7)$$

where $\gamma = \Theta(Rr/\varepsilon)$. For the complexity of SDP primal oracles, the same upper bounds holds.

- We introduce the quantum operator input model, and show that it is a simultaneous generalization of the other two input models which were considered earlier. In particular we show that both the sparse model and the quantum state model can be reduced to the quantum operator model with a constant overhead and with the choices of $\alpha = s$ and $\alpha = B$ respectively. Moreover, we show that for $\sigma = \Theta(1)$, we have that

$$T_{Tr}^\sigma(\gamma) = \tilde{\mathcal{O}}_\gamma(\alpha),$$

in the quantum operator model. We also show how to simulate a linear combination of Hamiltonians efficiently using this input model, and prove that

$$T_{Gibbs}(K, d, \theta) = \tilde{\mathcal{O}}_{\theta, d}(\sqrt{n}K\alpha).$$

This result is based on the idea of gradually building up an efficient data structure for state preparation, following ideas of [KP17]. This significantly improves the complexity of Gibbs-sampling compared to [AGGW17], which presented a Gibbs sampler subroutine in the sparse matrix input model with complexity

$$T_{Gibbs}(K, d, \theta) = \tilde{\mathcal{O}}_\theta(\sqrt{n}Ks^2d^2).$$

- We develop a new method for Gibbs-sampling in the quantum state model. Our approach, in contrast to the one in [BKL⁺18], does not introduce a dependence on the rank of the input matrices in the complexity. In particular we improve the complexity bound of [BKL⁺18]

$$T_{Gibbs}(K, d, \theta) = \mathcal{O}\left(\text{poly}(K, B, d, 1/\theta, \max_{j \in \{0 \dots m\}} [\text{rank}(A_j)])\right)$$

to

$$T_{Gibbs}(K, d, \theta) = \tilde{\mathcal{O}}_{d, \theta, n}((KB)^{3.5}).$$

An important consequence of this improvement is that in the complexity of SDP solving we do not get a dependence on the rank of the input matrices, unlike Brandão et al. [BKL⁺18]. For some quantum SDPs given in the quantum state input model the A_j matrices could naturally correspond to quantum states. In this case B would be just 1, but the rank could easily be proportional to n , e.g., for a highly mixed state, eliminating the speedup over the sparse input model. Finally note that this Gibbs-sampling method is only beneficial if $\sqrt{n} \leq (KB)^{2.5}$, otherwise the reduction to the quantum operator model with $\alpha = B$ gives a better algorithm.

For the quantum operator input model the above improvements lead to the complexity bound

$$T_{SDP}(\varepsilon) = \tilde{\mathcal{O}}((\sqrt{m} + \sqrt{n}\gamma)\alpha\gamma^4), \quad (8)$$

where $\gamma := \frac{Rr}{\varepsilon}$. Note that the $\Omega(\sqrt{n} + \sqrt{m})$ lower bound of [BS17] also applies to the quantum operator model due to our reductions, matching the above upper bound (8) up to polylog factors in n and m when γ and α are constant. For the quantum state input model our improved Gibbs-sampler yields the complexity bound

$$T_{SDP}(\varepsilon) = \tilde{\mathcal{O}}_n((\sqrt{m} + B^{2.5}\gamma^{3.5})B\gamma^4).$$

In both cases, the same bound holds for an SDP primal oracle but with $\gamma := R/\varepsilon$.

	With OR lemma / Two-Phase Search		Without OR lemma / Two-Phase Search	
	Sparse input	Quantum state input	Sparse input	Quantum state input
Previous Gibbs-sampling	$\tilde{\mathcal{O}}((\sqrt{m} + \sqrt{ns}\gamma^5)s\gamma^4)$ Theorem ⁹ 8 + [AGGW17]	$\tilde{\mathcal{O}}_n((\sqrt{m} + \text{poly}(\text{rk}))\text{poly}(\gamma, B))$ [BKL ⁺ 18]	$\tilde{\mathcal{O}}(\sqrt{mns}^2\gamma^8)$ [AGGW17]	$\tilde{\mathcal{O}}_n(\sqrt{m}\text{poly}(\gamma, B, \text{rk}))$ [BKL ⁺ 17]
Improved Gibbs-sampling	$\tilde{\mathcal{O}}((\sqrt{m} + \sqrt{n}\gamma)s\gamma^4)$ Theorem 17	$\tilde{\mathcal{O}}_n((\sqrt{m} + B^{2.5}\gamma^{3.5})B\gamma^4)$ Theorem 24	$\tilde{\mathcal{O}}(\sqrt{mns}\gamma^4)$ Corollary 18	$\tilde{\mathcal{O}}_n(\sqrt{m}B^{3.5}\gamma^{6.5})$ Corollary 25

Table 1: Summary of our query complexity bounds illustrating the role of our various improvements. Here we present the results for the sparse matrix and quantum state input models for comparison to prior work. However, note that our results presented for sparse input hold more generally for the quantum operator input model; to get the corresponding results one should just replace s by α in the table. Thereby similar bounds hold in the case of the quantum state input model too, after replacing s by B , which can be beneficial when $B^{2.5}\gamma^{2.5} \geq \sqrt{n}$. Notation: $\text{rk} = \max_{j \in \{0, \dots, m\}} \text{rank}(A_j)$ and $\gamma = \frac{Rr}{\varepsilon}$.

In Section 4 we give some applications of quantum SDP-solvers:

- We extend the idea of applying SDP-solving to the problem of shadow tomography: given an unknown, n -dimensional quantum state ρ , find ε -additive approximations of the expectation values $\text{Tr}(E_1\rho), \dots, \text{Tr}(E_m\rho)$ of several binary measurement operators. This problem was introduced by Aaronson in [Aar17], he gave an efficient algorithm in terms of the number of samples from ρ . In particular he proved that $\tilde{\mathcal{O}}(\log^4(m) \log(n)/\varepsilon^5)$ samples suffice. Brandão et al. [BKL⁺18] applied their SDP-solver to get a more efficient algorithm in terms of computation time when the measurements E_i are given in the quantum state model, while keeping the sample complexity as low as $\text{poly}(\log(m), \log(n), 1/\varepsilon, B)$. We simultaneously improve on both results, giving a sample bound of $\tilde{\mathcal{O}}(\log^4(m) \log(n)/\varepsilon^4)$ while also improving the best known time complexity [Aar17, BKL⁺18] of the implementation for all input models. Finally we show that if we can efficiently implement the measurements $\text{Tr}(E_1\rho), \dots, \text{Tr}(E_m\rho)$ on a quantum computer, then we can also efficiently represent E_1, \dots, E_m using the quantum operator input model.
- We apply the SDP-solvers to the problem of Quantum State discrimination: given a set of quantum states, what is the best POVM for discriminating between the states? We consider the case of minimizing the total error in the measurements. In this case we get an algorithm with running time $\tilde{\mathcal{O}}(\sqrt{k} \text{poly}(d, 1/\varepsilon))$ in the sparse input model, where k is the number of

⁹A similar result was independently proved by Brandão et al. [BKL⁺18].

states and d is the dimension of the states. Due to the quantum state model for SDP-solving, we can also solve the problem when the states that need to be discriminated are actually given as quantum states, rather than classical descriptions of density operators.

- We use the SDP-solver for the sparse matrix model to solve the problem of E-optimal design: given a set of k experiments, find the optimal distribution of the experiments that minimizes the variance in our knowledge of a d -dimensional system. Our final bound is $\tilde{O}\left((\sqrt{k} + \sqrt{d})\text{poly}(1/\varepsilon, P)\right)$, where P is a parameter that depends on the standard deviation of the experiments.

We end the paper with proving new lower bounds. Lower bounds on the quantum query complexity of SDP-solving for the sparse input model were presented in previous works [BS17, AGGW17]. We add to this by giving $\Omega(\sqrt{m}B/\varepsilon)$ and $\Omega(\sqrt{m}\alpha/\varepsilon)$ bounds for the quantum state model and quantum operator model respectively. These lower bounds show that the \sqrt{m} factor and the polynomial dependence on the parameters B, α , and $1/\varepsilon$ are necessary.

Compared to problems with a discrete input, proving lower bounds on continuous-input quantum problems gives rise to extra challenges and often requires more involved techniques, see for example the work of Belovs [Bel15] on generalizations of the adversary method. Due to these difficulties, fewer results are known in this regime. Examples of known continuous-input lower-bound results include phase-estimation related problems (cf. Bessen [Bes05]) and the complexity-theoretic version of the no-cloning theorem due to Aaronson [Aar09]. Recently, a new hybrid-method based approach was developed by Gilyén et al. [GAW17] in order to handle continuous-input oracles, which they use for proving a lower bound for gradient computation. We use their techniques to prove our lower bounds, combined with efficient reductions between input models stemming from the smooth-functions of Hamiltonians techniques developed in the work of van Apeldoorn et al. [AGGW17].

2 SDP-solving frameworks

In this section we present two frameworks for SDP-solving. The first is the Arora-Kale framework, which is used to find a good approximation of the optimal value and an almost feasible solution to the dual. Then we present an algorithm to implement a primal oracle. These together implement a full SDP-solver.

2.1 The Arora-Kale framework

Similarly to previous work [BS17, AGGW17] we build our results on the Arora-Kale framework. For a detailed description see the original paper by Arora and Kale [AK16]. For the specific application to general SDP-solvers, see [AGGW17]. For our application, the following broad overview suffices.

We assume that the first constraint is $\text{Tr}(X) \leq R$, i.e., $A_1 = I$ and $b_1 = R$. Remember that we set $A_0 = -C$ and $b_0 = -g$.

1. let $y = 0 \in \mathbb{R}^{m+1}$ and set $\theta = \frac{\varepsilon}{6Rr}$.
2. Repeat $\frac{\ln(n)}{\theta^2}$ times the following:
 - (a) Define $\rho := e^{-\sum_{j=0}^m y_j A_j} / \text{Tr}\left(e^{-\sum_{j=0}^m y_j A_j}\right)$.

(b) Find a \tilde{y} in the polytope

$$\mathcal{P}_\delta(\rho) := \left\{ \tilde{y} \in \mathbb{R}^{m+1} : b^T \tilde{y} \leq 0, \right. \\ \left. \sum_{j=0}^m \tilde{y}_j \text{Tr}(A_j \rho) \geq -\delta, \right. \\ \left. \tilde{y} \geq 0, \tilde{y}_0 = \frac{1}{2r}, \|\tilde{y}\|_1 \leq 1 \right\}.$$

for $\delta = \theta$ or conclude that none exists for $\delta = 0$.

(c) If no such \tilde{y} exists, then conclude that $\text{OPT} > g$ and stop.

(d) If such a \tilde{y} exists, then update $y \leftarrow \theta \tilde{y}$.

3. Conclude $\text{OPT} \leq g + \varepsilon$ and output $\frac{2r\theta}{\ln(n)}y + \frac{\varepsilon}{R}e_1 - e_0$ as a dual solution.

Brandão and Svore [BS17] observed that $\rho := e^{-\sum_j y_j A_j} / \text{Tr}(e^{-\sum_j y_j A_j})$ is a quantum Gibbs state and this state can be prepared efficiently on a quantum computer, allowing fast trace estimation, in particular resulting in a quadratic speedup in n .

A procedure that solves step (b) is called a θ -oracle. In the rest of this paper we will assume that the cost of updating the y vector is lower than the cost of a θ -oracle call.

Van Apeldoorn et al. [AGGW17] gave an oracle implementation that always outputs a 2-sparse \tilde{y} . The oracle is constructed using a geometric argument that boils down to minimizing over m angles, one for each constraint. Each angle is easily computed from the corresponding b_j and $\text{Tr}(A_j \rho)$, where a θ -additive error is allowed in the approximation of $\text{Tr}(A_j \rho)$. In [AGGW17] this minimization was done using Quantum Minimum finding [DH96], allowing for a quadratic speedup in m . Previously Brandão and Svore [BS17] applied other techniques to similarly get a quadratic speedup in m but this introduced a worse dependence on Rr/ε .

2.2 An SDP primal oracle

To construct a primal oracle, we use the same algorithm as Brandão et al. [BKL⁺18] following the proof of Lemma 4.6 of Lee, Raghavendra and Steurer [LRS15]. A few small reductions are required to apply this technique. To be able to work with density operators instead of X , the b_j s in the constraints $1 \dots m$ are scaled down by a factor R , such that every solution X' to the new SDP has trace at most 1. Then, we add one new variable denoted by ω such that

$$\rho := \begin{bmatrix} X' & 0 \\ 0 & \omega \end{bmatrix}.$$

Now $\text{Tr}(\rho) = 1$ and $\rho \succeq 0$ imply that $\text{Tr}(X') \leq 1$, and we get a new SDP that is equivalent to the previous one. It can be shown that in our input models this reduction does not introduce more than a constant factor overhead in the complexity.

The framework for an SDP primal oracle can now be summarized as follows (here we write A_j and b_j for the input after the reductions).

1. Let $y = 0 \in \mathbb{R}^{m+1}$ and $\theta = \frac{\varepsilon}{2R}$.

2. Repeat $\frac{\ln(n)}{\theta^2}$ times the following:
 - (a) Define $\rho := e^{-\sum_{j=0}^m y_j A_j} / \text{Tr}\left(e^{-\sum_{j=0}^m y_j A_j}\right)$.
 - (b) Find an index j such that $\text{Tr}(A_j \rho) \geq b_j$ or conclude correctly that for all j , $\text{Tr}(A_j \rho) \leq b_j + \theta$.
 - (c) If no j is found, then we are done and output y and $z = R\text{Tr}(X')$, where $\text{Tr}(X')$ is the probability¹⁰ of measuring ρ to be in the subspace corresponding to the variable X' .
 - (d) Otherwise update $y \leftarrow y + \theta e_j$.
3. Conclude that there is no solution for $\theta = 0$.

Both frameworks have a very similar structure. The main difference is that the primal oracle framework requires only a simple search, whereas the θ -oracle needed for the Arora-Kale framework is slightly more complex. Our implementation of the θ -oracle is always returning a 3-sparse vector, thus in both cases we will work with a y vector that is non-negative and $\tilde{\mathcal{O}}_n(1/\theta^2)$ sparse.

3 Improvements

3.1 Fast Quantum OR Lemma and Two-Phase Minimum Finding

To speed up the SDP solvers derived from this framework we use a fast version of the Quantum OR Lemma, as in [BKL⁺18]. They prove the following lemma:

Lemma 5 (Fast Quantum OR Lemma [BKL⁺18]). *Let Π_1, \dots, Π_m be projectors and ρ a quantum state. Suppose that either*

1. $\exists j$ s.t. $\text{Tr}(\Pi_j \rho) \geq 1 - \delta_1$, or
2. $\frac{1}{m} \sum_{j=1}^m \text{Tr}(\Pi_j \rho) \leq \delta_2$

for some $0 < \delta_1 \leq 1/2$ and $0 < \delta_2 \leq \frac{(1-\delta_1)^2}{12m}$. Then for all $\xi \in (0, 1)$ there is a procedure that accepts with probability at least $(1 - \delta_1)^2/4 - \xi$ in the first case, and probability at most $3m\delta_2 + \xi$ in the second case, and that uses 1 copy of ρ and $\mathcal{O}(\log(m)\sqrt{m}/\xi)$ applications of a controlled version of the reflection $I - 2\sum_{j=1}^m \Pi_j \otimes |j\rangle\langle j|$.

This lemma is almost the same as the original Quantum OR Lemma [HLM17] but with the addition that the algorithm requires only $\mathcal{O}(\log(m)\sqrt{m}/\xi)$ applications of the controlled reflection. In a recent paper Aaronson [Aar17] proved the *Gentle Quantum Search Lemma* using the Quantum OR Lemma. His proof can easily be extended to use the Fast Quantum OR Lemma. We call the resulting more efficient version *Two-Phase Quantum Search*.

In the setting of the Two-Phase Quantum Search we will have m algorithms for decision problems and we ask whether one of them evaluates to 1, and if so, to find one. We also know that all algorithms start with preparing some state ρ , followed by some procedure U_j that depends on the index of the decision problem $j \in [m]$. In classical deterministic processes it is quite natural that only one preparation of ρ is needed since the result can be stored. For bounded error classical

¹⁰Note that a θ -approximation of $\text{Tr}(X')$ is easy to compute by means of amplitude estimation if ρ can be efficiently prepared as a quantum state – which is the case in our algorithms.

processes $\mathcal{O}(\log(m))$ preparations of ρ suffice to get the error probability of one decision problem below $1/m$. By the classical union bound this is low enough that we can find a marked element with constant probability. However, if ρ is a quantum state and the U_j are quantum algorithms, then such a bound is not so straightforward, since progress made in constructing ρ might be destroyed when running one of the U_j . Nevertheless, using the Fast Quantum OR Lemma it can be shown that $\tilde{\mathcal{O}}(\log^4(m))$ samples from ρ suffice.

Lemma 6 (Two-Phase Quantum Search). *Let $\nu \in (0, 1)$. Let ρ be a quantum state and U_1, \dots, U_m be unitaries with the U_j accessible through a unitary U that acts as $U|j\rangle|\psi\rangle = |j\rangle U_j|\psi\rangle$. Then there is a quantum algorithm that using $\tilde{\mathcal{O}}(\log^4(m) \log(\nu))$ samples of ρ and $\tilde{\mathcal{O}}(\sqrt{m} \log(\nu))$ applications of U and its inverse, outputs with success probability at least $1 - \nu$ either*

- *a j such that $\text{Tr}\left((I \otimes |1\rangle\langle 1|)U_j \rho U_j^\dagger\right) \geq 1/3$, i.e., a j such that U_j outputs 1 with probability at least $1/3$ on input ρ ,*
- *or concludes correctly that $\text{Tr}\left((I \otimes |1\rangle\langle 1|)U_j \rho U_j^\dagger\right) < 2/3$ for all j , i.e., no unitary outputs 1 with probability at least $2/3$ on input ρ .*

Proof. This follows from the proof of Gentle Quantum Search in [Aar17, Lemma 15] using the Fast Quantum OR Lemma [BKL⁺18] instead the normal Quantum OR Lemma. \square

Using the above lemma we construct the *Two-Phase Quantum Minimum Finding* algorithm. It turns out that we need to use this algorithm in a situation where different values have different error-bars, therefore the statement gets slightly complicated. In typical use-cases one can probably just choose each error-margin η equal to say δ resulting in a simpler statement.

Lemma 7 (Two-Phase Quantum Minimum Finding). *Let $\delta, \nu' \in (0, 1)$. Let ρ be a quantum state and U_1, \dots, U_m be unitaries, with the U_j accessible through a unitary U that acts as $U|j\rangle|\psi\rangle = |j\rangle U_j|\psi\rangle$. Let $a_1, \dots, a_m, \eta_1, \dots, \eta_m$ be numbers such that $\min_j |a_j| + |\eta_j| \leq M$. Assume that with probability at least $2/3$, U_j computes a binary representation of a_j up to additive error η_j using one copy of ρ . Then, with probability at least $1 - \nu'$, we can find a j such that $a_j - \eta_j \leq \min_i (a_i + \eta_i) + \delta$ using $\tilde{\mathcal{O}}(\log^4(m) \log(M/\delta) \log(\nu'))$ samples of ρ and $\tilde{\mathcal{O}}(\sqrt{m} \log(\nu') \log(M/\delta))$ applications of U and its inverse.*

Proof. Do a binary search on the value v to precision δ by checking whether there is still an element with $a_i + \eta_i \leq v$ using Lemma 6 in each round with setting $\nu = \Theta(\nu' / \log(M/\delta))$. This binary search will result in a value $v \leq \min_i (a_i + \eta_i) + \delta$ with probability at least $1 - \nu/2$, and it is not hard to see that the last j found by Lemma 6 during the binary search will be such that $a_j - \eta_j \leq v$ with probability at least $1 - \nu/2$. Therefore this j satisfies the required inequality with probability at least $1 - \nu$. \square

This leads to the following general bound on SDP-solving.

Theorem 8. *Assume that updating an entry of $y \in \mathbb{R}^{m+1}$ in the data structure requires at most $\tilde{\mathcal{O}}(T_{\text{Gibbs}}(\gamma, \gamma^2, \gamma^{-1}))$ elementary operation, where $\gamma := 6Rr/\varepsilon$. Then there is a quantum SDP-solver for which*

$$T_{\text{SDP}}(\varepsilon) = \tilde{\mathcal{O}}_n\left(\left(\sqrt{m}T_{\text{Tr}}^\sigma(\gamma^{-1}) + T_{\text{Gibbs}}(\gamma, \gamma^2, \gamma^{-1})\right)\gamma^4\sigma^2\right),$$

similarly there is also a quantum algorithm with the same complexity, but with $\gamma := 6R/\varepsilon$, that implements an SDP primal oracle.

Proof. To construct an SDP-solver use both frameworks in succession, otherwise use only the primal oracle. The frameworks run for $\tilde{\mathcal{O}}_n(\gamma^2)$ iterations. In each iteration we need to update at most three entries of the y vector, which takes at most $\tilde{\mathcal{O}}(T_{Gibbs}(\gamma, \gamma^2, \gamma^{-1}))$ elementary operations by assumption. To search for a violated constraint when using the primal oracle framework, we use the Two-Phase Quantum Search, and we use Two-Phase Minimum Finding to implement the minimum finding needed in the Oracle for the Arora-Kale framework¹¹, following the geometric approach of [AGGW17, Lemma 14] for implementing a γ^{-1} -oracle.

Let $\rho := \rho_S^{\otimes k}$ where $k = 6(4\sigma\gamma)^2$, S is a uniform random seed and $\tilde{\rho}_S$ is a density operator, that is a $\gamma^{-1}/4$ -approximation in the trace distance of the Gibbs state $\rho_{Gibbs(y)}$ corresponding to the current y vector (for at least a $4/5$ fraction of the possible input seeds). Let U_j be the operator that applies a $(\gamma^{-1}/4, \sigma)$ -trace estimator to each copy of $\tilde{\rho}_S$ and takes the average of the outcomes. I.e., it obtains estimates of $\text{Tr}(A_j \tilde{\rho}_S)$ with bias at most $\gamma^{-1}/4$ and standard deviation at most σ independently k -times, taking the average at the end. By Chebyshev’s inequality we can see that this way U_j computes a $2\gamma^{-1}/4$ -precise estimate of $\text{Tr}(A_j \tilde{\rho}_S)$ with probability at least $5/6$. Also with probability at least $4/5$ we have that $\|\tilde{\rho}_S - \rho_{Gibbs(y)}\| \leq \gamma^{-1}/4$ and thus we get a $3\gamma^{-1}/4$ -precise estimate of $\text{Tr}(A_j \rho_{Gibbs(y)})$ with probability at least $4/5 \cdot 5/6 = 2/3$. The preparation of ρ can be performed using $kT_{Gibbs}(\gamma, \gamma^2, \gamma^{-1})$ queries by definition, whereas U_j can be implemented with query complexity $kT_{Tr}^\sigma(\gamma^{-1})$.

By using Two-Phase Quantum Search and Two-Phase Quantum Minimum Finding with $\nu = \tilde{\mathcal{O}}_n(\gamma^{-2})$ we get that it takes $\tilde{\mathcal{O}}_n(((\sqrt{m}T_{Tr}^\sigma(\gamma^{-1}) + T_{Gibbs}(\gamma, \gamma^2, \gamma^{-1}))\gamma^2\sigma^2)$ queries to implement an iteration. The stated final complexity follows considering the number of iterations. \square

In the rest of this paper we will give upper bounds on T_{Gibbs} and T_{Tr}^σ for different input models. In particular, due to the results from the next section $T_{Gibbs}(K, d, \theta)$ will always depend only logarithmically on d and θ , and $T_{Tr}^\sigma(\theta)$ will only depend logarithmically on θ . Nevertheless we left these parameters in Theorem 8 for completeness and to allow for comparison with previous results.

3.2 SDP-solving using the quantum operator input model

In this section we present recent results on Hamiltonian simulation which motivate the quantum operator input model. In particular, we restate results from [LC16] showing that using “block-encodings” of Hamiltonians, efficient Hamiltonian simulation can be implemented, moreover we show that this input model generalizes both other input models. Then we show that efficient Hamiltonian simulation leads to efficient trace estimation. Since we need to prepare Gibbs states for Hamiltonians of the form $\sum_{j=0}^m y_j A_j$ we also show how to obtain a block-encoding of a linear combination of Hamiltonians using an efficient data structure. Combining these techniques with the meta-algorithm of Theorem 8 leads to an efficient quantum SDP-solver, see Theorem 17.

3.2.1 Block-encodings of operators

Inspired by the work of Low and Chuang [LC16] and following the exposition of Gilyén and Wiebe [GW18] we introduce block-encodings of operators, which will be the central concept to

¹¹In the Oracle implementation of [AGGW17, Lemma 14] the minimum finding is not done over the computed traces, but rather the angles calculated using these traces. The trace \rightarrow angle conversion suffices with precision $\delta^{-1} = \mathcal{O}(\text{poly}(\gamma))$, and since the magnitude M of angles is bounded by π , we get $\log(M/\delta) = \mathcal{O}(\log(\gamma))$ in Lemma 7.

the quantum operator approach. Recall that a unitary U is a block-encoding of a matrix A if the top-left block of U is proportional to A :

$$U = \begin{pmatrix} A/\alpha & \cdot \\ \cdot & \cdot \end{pmatrix}$$

Note that $\|U\| = 1$ hence we must have $\|A\| \leq \alpha$.

Notation. Since unitaries are also operators, we will sometimes call a unitary a $(1, a, \varepsilon)$ -block-encoding of another unitary if it uses a ancillary qubits and is an ε -approximation in the operator norm. Note that every unitary is a $(1, 0, 0)$ -block-encoding of itself.

The following Hamiltonian simulation theorem is a corollary of the results of [LC16, Theorem 1], and provides the main motivation for this input model. For a detailed proof see the work of Chakraborty, Gilyén and Jeffrey [CGJ18].

Theorem 9 (Optimal block-Hamiltonian simulation [LC16]). *Suppose that U is an $(\alpha, a, \varepsilon/|2t|)$ -block-encoding of the Hamiltonian H . Then we can implement an ε -precise Hamiltonian simulation unitary V which is an $(1, a + 2, \varepsilon)$ -block-encoding of e^{itH} , with $\mathcal{O}\left(|\alpha t| + \frac{\log(1/\varepsilon)}{\log \log(1/\varepsilon)}\right)$ uses of controlled- U or its inverse and with $\mathcal{O}\left(a|\alpha t| + a\frac{\log(1/\varepsilon)}{\log \log(1/\varepsilon)}\right)$ two-qubit gates.*

Following the approach of [AGGW17], we show in Appendix A that efficient Hamiltonian simulation implies efficient trace estimation:

Corollary 10. *For $\sigma = 6$, we have that $T_{Tr}^\sigma(\theta) = \tilde{\mathcal{O}}_{\frac{\varepsilon}{2}}(\alpha)$ in the quantum operator input model.*

Proof. This is shown by Corollary 40 from Appendix A. □

Now we present a lemma stating the complexity of implementing a block-encoding of an operator accessed via sparse matrix input oracles. As a corollary we get that the quantum operator input model generalizes the sparse-matrix input model, when we set $\alpha = s$.

Lemma 11. *Suppose that A is a w -qubit s -sparse operator given in the sparse matrix input model (Oracles (4) and (5) with respect to a fixed j). Then we can implement an $(s, \tilde{\mathcal{O}}_{\frac{\varepsilon}{2}}(w), \varepsilon)$ -block-encoding of A with $\mathcal{O}(1)$ -queries, and $\tilde{\mathcal{O}}_{\frac{\varepsilon}{2}}(w)$ -other two-qubit quantum gates.*

Proof. See e.g., [BCK15]. □

We present a lemma based on ideas of [LC16, Corollary 9] showing how to implement a block-encoding of a (subnormalized) density operator using purified access to the density operator.

Lemma 12 (Block-encoding of a (subnormalized) density operator [LC16]). *Let G be a $(w + a)$ -qubit unitary which on the input state $|0\rangle^{\otimes w}|0\rangle^{\otimes a}$ prepares a purification $|\varrho\rangle$ of the subnormalized w -qubit density operator ϱ . Then we can implement a $(1, w + a, 0)$ -block-encoding of ϱ , with single use of G and its inverse and with $w + 1$ two-qubit gates.*

Proof. Let us write $|\varrho\rangle = \alpha|\rho_0\rangle + \beta|\rho_1\rangle$, where $\alpha, \beta \in \mathbb{R}$, $(I_{2^w} \otimes |0\rangle\langle 0| \otimes I_{2^{a-1}})|\rho_0\rangle = |\rho_0\rangle$ and $(I_{2^w} \otimes |1\rangle\langle 1| \otimes I_{2^{a-1}})|\rho_1\rangle = |\rho_1\rangle$. Moreover, without loss of generality we can assume that $|\varrho_0\rangle = \sum_{j=1}^D \sqrt{p_j}|\psi_j\rangle|0\rangle|\tilde{\psi}_j\rangle$ such that $\langle \tilde{\psi}_i | \tilde{\psi}_j \rangle = \delta_{ij}$ and $\varrho = \alpha^2 \sum_{j=1}^D p_j |\psi_j\rangle\langle \psi_j|$. Consider the $(2w+a+1)$ -qubit unitary $V = (I_{2^{w+1}} \otimes G^\dagger)(\text{SWAP}_{w+1} \otimes I_{2^a})(I_{2^{w+1}} \otimes G)$, where SWAP_{w+1} denotes the unitary which swaps the first $w+1$ -qubit register with the second $w+1$ -qubit register. Observe that

$$\begin{aligned}
(I_{2^w} \otimes \langle 0|^{\otimes 1+w+a})V(I_{2^w} \otimes |0\rangle^{\otimes 1+w+a}) &= (I_{2^w} \otimes \langle 0|\langle \varrho|)(\text{SWAP}_{w+1} \otimes I_{2^a})(I_{2^w} \otimes |0\rangle|\varrho\rangle) \\
&= \alpha^2 (I_{2^w} \otimes \langle 0|\langle \rho_0|)(\text{SWAP}_{w+1} \otimes I_{2^a})(I_{2^w} \otimes |0\rangle|\rho_0\rangle) \\
&= \alpha^2 \sum_{j=1}^D p_j (I_{2^w} \otimes \langle 0|\langle \psi_j|\langle 0|)(\text{SWAP}_{w+1} \otimes I_{2^a})(I_{2^w} \otimes |0\rangle|\psi_j\rangle|0\rangle) \\
&= \alpha^2 \sum_{j=1}^D p_j (I_{2^w} \otimes \langle 0| \otimes \langle \psi_j|\langle 0|)(|\psi_j\rangle|0\rangle \otimes I_{2^w} \otimes |0\rangle) \\
&= \alpha^2 \sum_{j=1}^D p_j (|\psi_j\rangle \otimes \langle \psi_j|) \\
&= \alpha^2 \sum_{j=1}^D p_j |\psi_j\rangle\langle \psi_j| \\
&= \varrho.
\end{aligned}$$

□

The above corollary essentially shows that the quantum operator input model generalizes the quantum state input model too, by choosing $\alpha = B$ (since $\mu_j^+ - \mu_j^- + |\mu_j^I| \leq B$). What is left is to show how to implement a linear combination of block-encodings, e.g., $A_j = \mu_j^+ \varrho_j^+ - \mu_j^- \varrho_j^- + \mu_j^I I$. We show how to efficiently implement such a block-encoding in the next subsection.

3.2.2 Implementing a linear combination of block-encodings

Following the work of Gilyén and Wiebe [GW18], in this subsection we show how to efficiently implement a linear combination of block-encodings. Together with the optimal block-Hamiltonian simulation theorem from the previous subsection this enables a clean and efficient way to implement Gibbs-sampling using the techniques developed in [AGGW17, Appendix B]. Using linear combinations of block-encodings by-passes the entrywise summation of the input matrices which was a major bottleneck in previous SDP-solvers for the sparse input model [AGGW17].

We define state-preparation unitaries in order to conveniently state our next lemma about implementing a linear combinations of block-encodings.

Definition 13 (State-preparation pair). *Let $y \in \mathbb{C}^m$ and $\beta \geq \|y\|_1$, the pair of unitaries (P_L, P_R) is called a (β, b, ε) -state-preparation pair for y if $P_L|0\rangle^{\otimes b} = \sum_{j=0}^{2^b-1} c_j|j\rangle$ and $P_R|0\rangle^{\otimes b} = \sum_{j=0}^{2^b-1} d_j|j\rangle$ such that $\sum_{j=0}^{m-1} |\beta \cdot (c_j^* d_j) - y_j| \leq \varepsilon$ and for all $j \in m, \dots, 2^b - 1$ we have $c_j^* d_j = 0$. A symmetric state-preparation pair also satisfies $c_j = d_j$ for all $j \in 0 \dots m - 1$.*

Lemma 14 (Linear combination of block encodings). *Let $A = \sum_{j=0}^m y_j A_j$ be a w -qubit operator and $\varepsilon \in \mathbb{R}_+$. Suppose that (P_L, P_R) is a $(\beta, b, \varepsilon_1)$ -state-preparation pair for y , $W = \sum_{j=0}^{m-1} U_j \otimes |j\rangle\langle j| +$*

$(I_{2^{w+a}} \otimes (I_{2^b} - \sum_{j=0}^{m-1} |j\rangle\langle j|))$ is an $w+a+b$ qubit unitary such that for all $j \in 0, \dots, m$ we have that U_j is an $(1, a, \varepsilon_2)$ -block-encoding of A_j . Then we can implement a $(\beta, a+b, \varepsilon_1 + \beta\varepsilon_2)$ -block-encoding of A , with a single use of W , P_R and P_L^\dagger .

Proof. Observe that

$$\widetilde{W} = (I_{2^w} \otimes I_{2^a} \otimes P_L^\dagger)W(I_{2^w} \otimes I_{2^a} \otimes P_R)$$

is $(\beta, a+b, \varepsilon_1 + \beta\varepsilon_2)$ -block-encoding of A . \square

Now we describe how to use a quantum-access classical RAM (QCRAM) to efficiently implement a state-preparation-pair unitary that can be used to construct the linear combinations of the block-encodings. In the SDP-solver we use this data structure for the summation of constraint matrices needed for Gibbs-sampling.

Lemma 15. *There is a data structure that can store an m -dimensional d -sparse vector y with θ -precision using a QCRAM of size $\widetilde{\mathcal{O}}_{\frac{m}{\theta}}(d)$. Furthermore:*

- Given a classical $\mathcal{O}(1)$ -sparse vector, adding¹² it to the stored vector has classical cost $\widetilde{\mathcal{O}}_{\frac{m}{\theta}}(1)$.
- Given that $\beta \geq \|y\|_1$ we can implement a (symmetric) $(\beta, \widetilde{\mathcal{O}}_{\frac{m}{\theta}}(1), \theta)$ -state-preparation pair for y with $\widetilde{\mathcal{O}}_{\frac{m}{\theta}}(1)$ queries to the QCRAM.

Proof. We use the data structure of [KP17, Appendix A]. \square

Corollary 16. *Having access to the above data structure for y , we have $T_{\text{Gibbs}}(K, d, \theta) = \widetilde{\mathcal{O}}_{\theta}(\alpha K \sqrt{n})$ in the quantum operator input model.*

Proof. This can be proven using the Gibbs-sampler of [AGGW17], combined with the above lemma for simulating the linear combination of the operators $\sum_{j=0}^m y_j A_j$. \square

This directly gives the following result for SDP-solving:

Theorem 17. *In the quantum operator input model*

$$T_{SDP}(\varepsilon) = \widetilde{\mathcal{O}}((\sqrt{m} + \sqrt{n}\gamma)\alpha\gamma^4),$$

where $\gamma = Rr/\varepsilon$ is. For a primal oracle the same complexity can be accomplished with $\gamma = R/\varepsilon$. The input oracle of the quantum operator model can be constructed using $\mathcal{O}(1)$ queries and $\widetilde{\mathcal{O}}_{mn\alpha\gamma}(1)$ elementary operations in the sparse matrix model and also in the quantum state model with setting $\alpha = s$ or $\alpha = B$ respectively. Therefore the above bound applies to these input models too.

Proof. The complexity statement follows from Theorem 8 using Corollary 10 and 16. The reductions follow from Lemma 11-12. \square

If we do not apply Two-Phase Minimum Finding but use standard quantum minimum finding [DH96] instead, then we get a result with a slightly better dependence on γ , cf. (6):

Corollary 18. *In the quantum operator input model*

$$T_{SDP}(\varepsilon) = \widetilde{\mathcal{O}}(\sqrt{mn}\alpha\gamma^4),$$

where $\gamma = Rr/\varepsilon$ for a full SDP-solver and $\gamma = R/\varepsilon$ for a SDP primal oracle.

¹²In order to avoid error accumulation from repeated roundings, we assume for simplicity that there can be at most $\mathcal{O}(\text{poly}(m/\theta))$ such calls to the data structure in total.

3.3 Exponentially improved Gibbs-sampling in the quantum state input model

In this subsection we show how to harness the special structure of the quantum state input model, to improve the complexity of Corollary 16. As shown in [BKL⁺18] this allows for an SDP-solver with a polylog dependence on n . We improve on the results of [BKL⁺18] by constructing a Gibbs-sampler with no explicit dependence on the rank of the input matrices. Moreover, we also improve the dependence on precision from polynomial to logarithmic.

We will use the following lemma about projectors.

Lemma 19. *Let $0 < q < 1$, Π be a projector and ϱ a subnormalized density operator. Suppose that $q\Pi \preceq \varrho$, $(I - \Pi)\varrho(I - \Pi) = 0$ and we have access to an a -qubit unitary $U_{\tilde{\varrho}}$ preparing a purification of a subnormalized density operator $\tilde{\varrho}$ such that $\|\varrho - \tilde{\varrho}\|_1 \leq 4\nu$. Then we can prepare a purification of a subnormalized density operator $\tilde{\varrho}_{\text{unif.}}$ such that $\left\| \frac{q}{4}\Pi - \tilde{\varrho}_{\text{unif.}} \right\|_1 \leq \tilde{\mathcal{O}}(\nu/q)$, with $\tilde{\mathcal{O}}_\nu(1/q)$ queries to $U_{\tilde{\varrho}}$ and its inverse and using $\tilde{\mathcal{O}}_\nu(a/q)$ two-qubit gates.*

Proof. First let us assume that we have access to U_ϱ instead of $U_{\tilde{\varrho}}$. Then, using Corollary 38 from Appendix A, we could implement a unitary W which is a $(1, \tilde{\mathcal{O}}_{\nu,q}(a), 0)$ -block-encoding of V such that $\left\| \left(V - \frac{\sqrt{q}}{2\sqrt{\varrho}} \right) \Pi \right\| \leq \nu$. Note that since $q\Pi \preceq \varrho$ and $(I - \Pi)\varrho(I - \Pi) = 0$ we know that ϱ is supported on the image of Π , in particular $\Pi\varrho\Pi = \varrho$. Using Hölder's inequality it is easy to see that

$$\left\| V\varrho V^\dagger - \frac{\sqrt{q}}{2\sqrt{\varrho}}\varrho\frac{\sqrt{q}}{2\sqrt{\varrho}} \right\|_1 \leq 2\nu,$$

which is equivalent to saying

$$\left\| V\varrho V^\dagger - \frac{q}{4}\Pi \right\|_1 \leq 2\nu.$$

Corollary 38 shows that W can be implemented with a single use of a controlled Hamiltonian simulation unitary $e^{it\varrho}$, with maximal simulation time $|t| \leq \tilde{\mathcal{O}}_\nu(1/q)$. This translates to $\tilde{\mathcal{O}}_\nu(1/q)$ uses of U_ϱ and its inverse as shown by Lemma 41.

Considering that $\|\varrho - \tilde{\varrho}\|_\infty \leq \|\varrho - \tilde{\varrho}\|_1 \leq 4\nu$, if in the implementation of the controlled Hamiltonian simulation we replace U_ϱ by $U_{\tilde{\varrho}}$, then we make no bigger error than $\tilde{\mathcal{O}}(\nu/q)$, as shown by Theorem 9. The resulting new unitary \tilde{W} will be therefore an $(1, \tilde{\mathcal{O}}_{\nu,q}(a), \tilde{\mathcal{O}}(\nu/q))$ -block-encoding of V . Therefore we can prepare a purification of the a subnormalized density operator $\tilde{\varrho}_{\text{unif.}}$ such that

$$\left\| \frac{q}{4}\Pi - \tilde{\varrho}_{\text{unif.}} \right\|_1 \leq \tilde{\mathcal{O}}(\nu/q).$$

□

In the proof of the next lemma we will mostly be looking at *Eigenvalue threshold projectors*.

Definition 20 (Eigenvalue threshold projector). *Suppose H is a Hermitian matrix and $q \in \mathbb{R}$. Let $\Pi_{H>q}$ denote the orthogonal projector corresponding to the subspace spanned by the eigenvectors of H that have eigenvalue larger than q . We define $\Pi_{H\leq q} = I - \Pi_{H>q}$ in a similar way.*

We are now ready to prove the main lemma of this section, an improved Gibbs-sampler in the quantum state model. In the proof we will use some specific conventions and notation to simplify the form of the proof. We say that two subnormalized density operator are δ -close when their trace distance is at most δ , and that two unitaries are δ -close when their operator norm distance is at

most δ . We will always work with purifications of subnormalized density operators, so when for example we say that we apply an operator to a subnormalized density operator, we mean that we apply the operator to its purification.

Lemma 21 (Gibbs-sampling of the difference of density operators). *Suppose that we have unitaries U_{ϱ^\pm} preparing a purification of the subnormalized density operators ϱ^\pm using $a = \mathcal{O}(\text{poly log}(n))$ qubits. Let¹³ $H := (\varrho^+ - \varrho^-)/2$, $\beta \in [1, n/2]$ and $\delta, \eta \in (0, 1]$. Assume we are given a point $q \in [2/n, 1/\beta]$ that is η -far from the spectrum of H , i.e.*

$$|\lambda - q| \geq \eta \text{ for all } \lambda \in \text{Spec}(H).$$

Then we can prepare a purification of an approximate Gibbs-state $\tilde{\rho}_{\text{Gibbs}}$ such that

$$\left\| \frac{e^{\beta H}}{\text{Tr}(e^{\beta H})} - \tilde{\rho}_{\text{Gibbs}} \right\|_1 \leq \delta$$

with¹⁴ $\tilde{\mathcal{O}}_{\frac{n}{\delta}}(q^{-1.5}/\eta)$ queries to controlled- U_{ϱ^\pm} and their inverses.¹⁵

Proof. The main idea of the proof is that we prepare (slightly subnormalized) density operators corresponding to $\Pi_{H>q}$ and $\Pi_{H\leq q}$, i.e., uniform distributions over a partition of eigenspaces of H . Utilising these states we prepare subnormalized Gibbs states on the corresponding subspaces, then merge and amplify the states in order to obtain the final Gibbs state. This is beneficial since on the subspace corresponding to $\Pi_{H\leq q}$ the map $e^{\beta H}$ is nicely bounded. However, on the image of $\Pi_{H>q}$ the map $e^{\beta H}$ might behave wildly, and in the extreme case this map might magnify the amplitude of some eigenvectors tremendously. This implies that we need to “find” such magnified elements, as the Gibbs state is concentrated around them. Fortunately the rank of $\Pi_{H>q}$ is at most $1/q$, which makes it easier to “find” the extreme vectors than if we would apply the same procedure to the uniform distribution I/n .

We start with implementing the unitary $\tilde{V}_{H>q}$, that labels eigenstates of H corresponding to which component of $\mathbb{R} \setminus \{q\}$ their eigenvalue lies in. More precisely, we set $\delta' := \tilde{\Theta}(\delta q^2)$, and we want to implement a unitary $\tilde{V}_{H>q}$ which is a $(1, \tilde{\mathcal{O}}_{\frac{n}{\delta'}}(1), \delta')$ -block-encoding of $(\Pi_{H>q} \otimes I + \Pi_{H\leq q} \otimes X)$. Due to the assumption that q lies at least η -far from $\text{Spec}(H)$ we can implement these unitaries using $\Theta(\eta)$ precise phase estimation of the operator e^{iH} , repeated $\mathcal{O}(\log(1/\delta'))$ times. This can be implemented with $\tilde{\mathcal{O}}_{\delta'}(1/\eta)$ queries as show by Lemmas 12 and 14.

Now let us consider Gibbs-sampling on the image of $\Pi_{H>q}$. Let $\varsigma^+ := (\varrho^+ + \varrho^-)/2$ be a subnormalized density operator, which we can prepare in a purified manner using $\mathcal{O}(1)$ queries. Also let

$$\varsigma_{H>q}^+ := \Pi_{H>q} \varsigma^+ \Pi_{H>q},$$

and observe that we can prepare $\tilde{\varsigma}_{H>q}^+$, such that

$$\left\| \tilde{\varsigma}_{H>q}^+ - \varsigma_{H>q}^+ \right\|_1 \leq 2\delta' (\leq q/4), \tag{9}$$

¹³In case n is not a power of 2 we still represent H on $\lceil \log_2(n) \rceil$ qubits, but think about it as an $\mathbb{C}^{n \times n}$ operator.

¹⁴We think that it should be possible to improve the complexity to $\tilde{\mathcal{O}}_{\frac{n}{\delta}}(q^{-1}/\eta)$ using recent results about variable-time amplitude amplification and estimation [CGJ18].

¹⁵If U_{ϱ^\pm} are not controlled, then it is easy to construct a controlled version using $\mathcal{O}(a)$ extra ancilla qubits.

by applying $\tilde{V}_{H>q}$ to ς^+ (the second inequality can be assumed w.l.o.g. since $\delta' = \tilde{\Theta}(\delta q^2)$).

Observe that

$$\begin{aligned}
q\Pi_{H>q} &= \Pi_{H>q}(q\Pi_{H>q})\Pi_{H>q} \\
&\preceq \Pi_{H>q}(H)\Pi_{H>q} \\
&\preceq \Pi_{H>q}(H + 2\varrho^-)\Pi_{H>q} \\
&= \varsigma_{H>q}^+.
\end{aligned} \tag{10}$$

This allows us to apply Lemma 19 to $\varrho := \varsigma_{H>q}^+$, $\tilde{\varrho} := \tilde{\varsigma}_{H>q}^+$ and $\Pi := \Pi_{H>q}$ with $\nu := \delta'$ so we get that we can prepare a state $\tilde{\varrho}_{\text{unif.}}$ such that

$$\left\| \frac{q}{4}\Pi_{H>q} - \tilde{\varrho}_{\text{unif.}} \right\|_1 \leq \tilde{\mathcal{O}}(\delta'/q)$$

using $\tilde{\mathcal{O}}_{\delta'}(q^{-1}/\eta)$ queries.

Now we can check if $\Pi_{H>q} = 0$ or not as follows. If it is not 0 then $\text{Tr}(\Pi_{H>q}) \geq 1$ and hence by (9)-(10) we have $\text{Tr}(\tilde{\varsigma}_{H>q}^+) \geq \text{Tr}(\varsigma_{H>q}^+) - q/4 \geq q - q/4 = 3q/4$. Since $\varsigma_{H>q}^+ = \text{Tr}(\Pi_{H>q})\varsigma_{H>q}^+ + \text{Tr}(\Pi_{H>q})$, when $\Pi_{H>q} = 0$ it similarly follows that $\text{Tr}(\tilde{\varsigma}_{H>q}^+) \leq \text{Tr}(\varsigma_{H>q}^+) + q/4 = q/4$. Thus we can check whether $\Pi_{H>q} = 0$ by checking whether $\text{Tr}(\tilde{\varsigma}_{H>q}^+) \leq q/4$ or $\text{Tr}(\tilde{\varsigma}_{H>q}^+) \geq 3q/4$. This can be done with success probability at least $1 - \delta'$ by using amplitude estimation with $\tilde{\mathcal{O}}_{\delta'}(q^{-0.5})$ calls to the procedure preparing $\tilde{\varsigma}_{H>q}^+$, costing $\tilde{\mathcal{O}}_{\delta'}(q^{-0.5}/\eta)$ queries in total. For the final Gibbs-sampling we will consider the Gibbs state on the image of $\Pi_{H>q}$ and $\Pi_{H \leq q}$ separately. Therefore if $\Pi_{H>q} = 0$ we only need to consider the Gibbs state on the image of $\Pi_{H \leq q}$, which we do later in this proof. For now we assume $\Pi_{H>q} \neq 0$ and consider the Gibbs state on its image.

Now we use binary search in order to find λ_{\max} the maximal eigenvalue of H , with precision $\beta^{-1}/2$ and success probability $1 - \delta'$. By our assumption $\lambda_{\max} \in (q, 1]$. We start each iteration of the binary search by performing phase estimation on $\tilde{\varsigma}_{H>q}^+$ using the unitary e^{iH} with precision $\beta^{-1}/4$ and success probability $1 - q/4$. By (9)-(10) we know that the eigenvector corresponding to λ_{\max} is present with probability at least $3q/4$ in $\tilde{\varsigma}_{H>q}^+$, and the other eigenvalues are present with a probability at most 1 in total. Therefore the probability of obtaining a phase estimate $\tilde{\lambda}$ such that $\tilde{\lambda} \geq \lambda_{\max} - \beta^{-1}/4$ is at least $q/2$, whereas the total probability of obtaining a phase estimate λ' such that $\lambda' \geq \lambda_{\max} + \beta^{-1}/4$ is at most $q/4$. Therefore we can perform each iteration of the binary search with precision $\beta^{-1}/2$ and success probability $1 - \delta'/\log(q^{-1})$ by applying amplitude estimation to the probability of getting an eigenvalue estimate from the current search interval, using $\tilde{\mathcal{O}}_{\beta}(q^{-0.5})$ repetitions of the initial state preparation and phase estimation procedure. Thus each iteration has query complexity $\tilde{\mathcal{O}}_{\beta}(q^{-0.5}(\beta + 1/\eta)) = \tilde{\mathcal{O}}(q^{-1.5} + q^{-0.5}/\eta)$, giving the same total query complexity bound $\tilde{\mathcal{O}}(q^{-1.5} + q^{-0.5}/\eta)$ for the complete binary search.

After finding the minimum up to precision $1/(2\beta)$ we can compute a number $\tilde{\lambda}_{\max}$ such that

$$\tilde{\lambda}_{\max}I \succeq H \text{ but } (\tilde{\lambda}_{\max} - 1/\beta)I \not\succeq H. \tag{11}$$

Using Lemmas 12-14, and the results of [AGGW17, Appendix B] we can implement an $(1/2, \tilde{\mathcal{O}}_{\beta/\delta'}(1), \delta')$ -block encoding of $e^{\beta \frac{H - \tilde{\lambda}_{\max}I}{2}}$ using $\tilde{\mathcal{O}}_{\delta'}(\beta)$ queries. Applying this map to $\tilde{\varrho}_{\text{unif.}}$ gives an $\tilde{\mathcal{O}}(\delta'/q)$ -approximation of the subnormalized density operator $\frac{qe^{-\beta \tilde{\lambda}_{\max}}}{16}\Pi_{H>q}e^{\beta H}$. Observe that since we

assumed $\text{Tr}(\Pi_{H>q}) \geq 1$, by (11) we get that

$$\text{Tr}\left(\frac{qe^{-\beta\tilde{\lambda}_{\max}}}{16}\Pi_{H>q}e^{\beta H}\right) \geq q/(16e). \quad (12)$$

Thus we can prepare a subnormalized $\tilde{\mathcal{O}}(\delta'/q)$ -approximation of the Gibbs state on the image of $\Pi_{H>q}$ having trace at least $q/(16e)$.

Now we consider the Gibbs state on $\Pi_{H\leq q}$. First observe that we can prepare the density operator I/n by preparing a the maximal entangled state $\frac{1}{\sqrt{n}}\sum_{j=1}^n|j\rangle|j\rangle$ using $\tilde{\mathcal{O}}(\log(n))$ two-qubit quantum gates. With a single use of the unitary $\tilde{V}_{H>q}$ we can prepare an $\mathcal{O}(\delta')$ approximation of $\frac{1}{n}\Pi_{H\leq q}$ by simply marking the appropriate eigenstates of I/n , which takes $\tilde{\mathcal{O}}(1/\eta)$ queries. Then we apply the map $e^{\beta H/2}/(2\sqrt{e})$ on the subspace $\Pi_{H\leq q}$ with δ' accuracy. Since $q \leq 1/\beta$ all eigenvalues of $\beta H/2$ that we are concerned with are smaller in absolute value than $1/2$. As shown in [AGGW17, Appendix B] this implies that implementing the map $e^{\beta H/2}/(2\sqrt{e})$ with δ' precision requires Hamiltonian simulation of βH for constant time, which can be done using $\tilde{\mathcal{O}}_{\delta'}(\beta)$ queries. Therefore we can prepare a $\mathcal{O}(\delta')$ approximation of the state $\frac{1}{4en}\Pi_{H\leq q}e^{\beta H}$ with $\tilde{\mathcal{O}}_{\delta'}(\beta + 1/\eta) \leq \tilde{\mathcal{O}}_{\delta'}(q^{-1}/\eta)$ queries.

Like before, we would like to lower bound the trace of the created subnormalized density operator. First note that $\|H\|_1 \leq 1$, and so the number¹⁶ of eigenvalues that are larger than q in absolute value is at most $1/q \leq n/2$, thus $\text{Tr}(\Pi_{|H|\leq q}) \geq n/2$. Also note that $\Pi_{H\leq q}e^{\beta H} \succeq \Pi_{|H|\leq q}e^{\beta H}$, and for an eigenvalue λ such that $|\lambda| \leq q$ we have $e^{\beta\lambda} \geq e^{-\beta q} \geq 1/e$. It follows that

$$\text{Tr}\left(\frac{1}{4en}\Pi_{H\leq q}e^{\beta H}\right) \geq \text{Tr}\left(\frac{1}{4en}\Pi_{|H|\leq q}e^{\beta H}\right) \geq \text{Tr}\left(\frac{1}{4e^2n}\Pi_{|H|\leq q}\right) \geq \frac{1}{8e^2}. \quad (13)$$

As we can now Gibbs-sample on both parts of the spectrum, we are ready to combine the two. Let¹⁷ $\xi := \min\left(\frac{qe^{-\beta\tilde{\lambda}_{\max}}}{16}, \frac{1}{4en}\right)$, then we can prepare a purification of $\tilde{\rho}_G$ which is an $\tilde{\mathcal{O}}(\delta'/q)$ -approximation of

$$\rho_G := \frac{\xi}{2}e^{\beta H} = \underbrace{\left(\frac{\xi}{2}\frac{16}{qe^{-\beta\tilde{\lambda}_{\max}}}\right)}_{\leq 1/2}\frac{qe^{-\beta\tilde{\lambda}_{\max}}}{16}\Pi_{H>q}e^{\beta H} + \underbrace{\left(\frac{\xi}{2}\frac{4en}{1}\right)}_{\leq 1/2}\frac{1}{4en}\Pi_{H\leq q}e^{\beta H},$$

by mixing the two subnormalized Gibbs states on the corresponding subspaces with appropriate ($\leq 1/2$) coefficients. This subnormalized Gibbs state $\tilde{\rho}_G$, can be prepared at the same cost as the two partial Gibbs-state preparation, that is $\tilde{\mathcal{O}}(q^{-1}/\eta)$ queries¹⁸ as we have already shown.

Note that $\text{Tr}(\rho_G) = \Omega(q)$ as shown by (12)-(13), therefore we can use $\tilde{\mathcal{O}}(\sqrt{1/q})$ amplitude amplification steps to prepare an $\tilde{\mathcal{O}}(\delta'/q^2)$ approximation of $\frac{\rho_G}{\text{Tr}(\rho_G)}$, which is clearly ρ_{Gibbs} . In total this yields an $\tilde{\mathcal{O}}(q^{-1.5}/\eta)$ query algorithm. Since $\delta' = \tilde{\Theta}(\delta q^2)$ this concludes the proof. \square

¹⁶We count eigenvalues with algebraic multiplicity.

¹⁷In the special case when $\Pi_{H>q} = 0$ we simply set $\xi := \frac{1}{4en}$.

¹⁸Note that we do the maximum finding to find $\tilde{\lambda}_{\max}$ only once, and we do not count its complexity in the state preparation.

The following corollary expands our new Gibbs-sampling result, it gives an exponential improvement in terms of the precision over the previous approach for this input model by Brandão et al. [BKL⁺18]. The dependence on the success probability is worse, but in our application to SDP-solving we only require success for a constant fraction of random seeds. Furthermore, there is no longer a dependence on the rank of the input matrices.

Theorem 22. *Suppose we have query access to the $a = \text{poly log}(n)$ -qubit unitaries U_{ρ^\pm} preparing a purification of the (subnormalized) density operators $\rho^\pm \in \mathbb{C}^{n \times n}$, such that $H = (\rho^+ - \rho^-)/2$, $\beta \geq 1$, $\theta, \delta \in (0, 1]$. Then there is a quantum algorithm, that using¹⁹ $\tilde{\mathcal{O}}_\theta(\beta^{3.5}/\delta)$ queries to controlled- U_{ρ^\pm} or their inverses, prepares a purification of a quantum state ρ_S such that*

$$\left\| \rho_S - \frac{e^{-\beta H}}{\text{Tr}(e^{-\beta H})} \right\|_1 \leq \theta,$$

where S is an $\mathcal{O}(\log(\beta/\delta))$ -bit random seed, and the above holds for at least $(1 - \delta)$ -fraction of the seeds.

Proof. If $\beta \geq n/2$, then we simply use the Gibbs-sampler from Theorem 17. Otherwise, using the random seed S we generate a uniform random number q_S from the interval $[1/(2\beta), 1/\beta]$. Note that since $\text{Tr}(|H|) \leq 1$ we have that $|\text{Spec}(|H|) \cap [1/(2\beta), 1/\beta]| \leq 2\beta$. Also the length of the interval is $1/(2\beta)$ therefore a random point in the interval falls $\delta/(8\beta^2)$ -close to $\text{Spec}(|H|)$ with probability at most δ . Therefore the random seed can be used in such a way that q_S will be $\eta = \delta/(8\beta^2)$ far from any point of $\text{Spec}(|H|)$ with probability at least $1 - \delta$. If this is the case the procedure of Lemma 21 prepares the sought Gibbs state with the stated complexity. \square

Corollary 23. *Having access to the data structure of Lemma 15 storing the vectors $\nu^\pm \in \mathbb{R}^{m+1}$ such that $\nu_j^\pm = y_j \mu_j^\pm$ for all $j \in 0 \dots m$, we have that $T_{\text{Gibbs}}(K, d, \theta) = \tilde{\mathcal{O}}_\theta((BK)^{3.5})$ using the quantum state input model.*

Proof. To start, let us define

$$H := \sum_{j=0}^m \frac{y_j A_j}{KB} = \sum_{j=0}^m \frac{y_j}{KB} (\mu_j^+ \rho_j^+ - \mu_j^- \rho_j^- + \mu^I I) = \underbrace{\sum_{j=0}^m \frac{y_j \mu_j^+}{KB} \rho_j^+}_{\rho^+ :=} - \underbrace{\sum_{j=0}^m \frac{y_j \mu_j^-}{KB} \rho_j^-}_{\rho^- :=} + I \sum_{j=0}^m \frac{y_j \mu_j^I}{KB}.$$

Notice that we can ignore the identity terms since adding identities in the exponent does not change a Gibbs state, also let

$$\beta := KB \geq \sum_{j=0}^m y_j (\mu_j^+ + \mu_j^-).$$

Using Lemma 15 we can see that a $\mathcal{O}(\theta/\beta)$ approximation of ρ^\pm can be prepared with $\mathcal{O}(1)$ queries and using $\mathcal{O}(\text{poly log}(m\beta/\theta))$ elementary operation. By setting $\delta := 1/5$ and using Theorem 22 the statement follows. \square

This directly gives the following result for SDP-solving in the quantum state model

¹⁹Similarly to Lemma 21 we think that it should be possible to improve the complexity to $\tilde{\mathcal{O}}_\theta(\beta^3/\delta)$ using recent results about variable time amplitude amplification and estimation [CGJ18] (maybe at the expense of worse dependence on the error).

Theorem 24. *In the quantum state input model*

$$T_{SDP}(\varepsilon) = \tilde{O}((\sqrt{m} + B^{2.5}\gamma^{3.5})B\gamma^4),$$

where $\gamma = Rr/\varepsilon$. The same bound holds for a primal oracle with $\gamma = R/\varepsilon$.

Proof. This follows directly from Theorem 8 using Corollaries 10, 16 and 23. \square

Also we can simply not apply Two-Phase minimum finding and use standard quantum minimum finding [DH96] instead to get a slightly better dependence on γ , cf. (6):

Corollary 25. *In the quantum state input model*

$$T_{SDP}(\varepsilon) = \tilde{O}(\sqrt{m}B^{3.5}\gamma^{6.5}),$$

where $\gamma = Rr/\varepsilon$. The same bound holds for a primal oracle with $\gamma = R/\varepsilon$.

4 Applications

In previous works on quantum SDP-solving [AGGW17, BS17] it remained an open question whether any applications could be found in the regime where Rr/ε was small enough to get a speedup over the best classical methods. Later Brandão et al. [BKL⁺18] showed that SDP primal oracles can be used to solve the problem of shadow tomography if the input is given in the quantum state model. Shadow tomography was recently proposed by Aaronson [Aar17], who also gave a sample-efficient algorithm. In Section 4.1 we apply our improved SDP primal oracles to this problem, simultaneously improving the sample complexity and computational complexity compared to the previous works.

We also propose new applications to quantum SDP-solvers, namely the problems of *quantum state discrimination* and *E-optimal design*. In both cases we show a speedup over the best possible classical algorithm in terms of some input parameters, while suffering from a massive dependence on other parameters.

4.1 Improved shadow tomography

We apply the idea from Brandão et al. [BKL⁺18] to use an SDP primal oracle to the problem of *shadow tomography* proposed by Aaronson [Aar17]. In *shadow tomography* we are given the ability to sample from an n -dimensional quantum state τ and we have a description of some measurement operators E_1, \dots, E_m ; the goal is to find ε -approximations of the corresponding expectation values $\text{Tr}(E_j\tau)$ for all $j \in [m]$. Aaronson showed that this can be done with only

$$\tilde{O}\left(\frac{\log^4(m) \log(n)}{\varepsilon^5}\right)$$

samples from τ , but his method has high computational costs.

In [BKL⁺18] Brandão et al. showed that a slightly relaxed problem can be efficiently solved using an SDP primal oracle. The problem they solved is to find a $y \in \mathbb{R}^m$ for which $\sigma :=$

$e^{-\sum_{j=1}^m y_j E_j} / \text{Tr}\left(e^{-\sum_{j=1}^m y_j E_j}\right)$ is such that $|\text{Tr}(E_j(\tau - \sigma))| \leq \varepsilon/2 \forall j \in [m]$, i.e., σ is in

$$\begin{aligned} \mathcal{P}_\varepsilon = \{ & \sigma : \sigma \succeq 0 \\ & \text{Tr}(\sigma) = 1 \\ & \text{Tr}(\sigma E_j) \leq \text{Tr}(\tau E_j) + \varepsilon/2 \quad \forall j \in [m] \\ & \text{Tr}(-\sigma E_j) \leq \text{Tr}(-\tau E_j) - \varepsilon/2 \quad \forall j \in [m] \}. \end{aligned}$$

We call the problem of finding a classical description of τ that suffices to solve the shadow tomography problem without any more samples from τ the *descriptive shadow tomography problem*. In particular if we get a vector y as above, then for a given $j \in [m]$ using $\tilde{\mathcal{O}}_m(1/\varepsilon^2)$ invocations of a Gibbs-sampler for y followed by the measurement E_j suffices to find an ε -approximation of $\text{Tr}(\tau E_j)$ with success probability at least $1 - \mathcal{O}(1/m)$. If we can coherently apply E_j , then using amplitude estimation techniques the number of (coherent) Gibbs-sampler calls can be reduced to $\tilde{\mathcal{O}}_m(1/\varepsilon)$.

Due to the output size of the shadow tomography problem, a trivial $\Omega(m \log(1/\varepsilon))$ lower bound can be given on the computational complexity. However, this limitation does not exist for the descriptive shadow tomography problem. Both problems clearly have the same sample complexity, furthermore the best known lower bound on the sample complexity is $\Omega(\log(m)/\varepsilon^2)$ [Aar17].

Theorem 26. *The descriptive shadow tomography problem can be solved using*

$$\tilde{\mathcal{O}}\left(\frac{\log^4(m) \log(n)}{\varepsilon^4}\right)$$

samples from τ . Furthermore, when the E_j matrices are accessible in the quantum operator model this can be done using

$$\tilde{\mathcal{O}}\left(\left(\sqrt{m} + \frac{\sqrt{n}}{\varepsilon}\right) \frac{\alpha}{\varepsilon^4}\right)$$

queries. It follows that the same bound holds with $\alpha = s$ for the sparse model and with $\alpha = B$ for the quantum state model. When the measurements are given in the quantum state model the query complexity can be also bounded by

$$\tilde{\mathcal{O}}_n\left(\left(\sqrt{m} + \min\left(\frac{\sqrt{n}}{\varepsilon}, \frac{B^{2.5}}{\varepsilon^{3.5}}\right)\right) \frac{B}{\varepsilon^4}\right).$$

Proof. The samples from τ are only used for calculating the values b_j , i.e., $\varepsilon/4$ approximations of $\text{Tr}(\tau E_j)$, when checking the constraints in the SDP primal oracle. Like in [BKL⁺18] we make a small adjustment to our SDP primal oracle: when Gibbs-sampling the Gibbs state ρ , we also sample τ to create the state $\rho \otimes \tau$. Then, when checking the constraint, we measure $E_j \otimes -E_j$ to obtain an approximation of $\text{Tr}(E_j \rho) - \text{Tr}(E_j \tau)$. Notice that our SDP primal oracle uses $\tilde{\mathcal{O}}\left(\frac{\log^4(m) \log(n)}{\varepsilon^4}\right)$ Gibbs states ($\tilde{\mathcal{O}}_{\log(n)}(\log^4(m)/\varepsilon^2)$ in each of the $\mathcal{O}(\log(n)/\varepsilon^2)$ iterations) and hence the modified version uses that many samples from τ too.

The statement about the computational complexity follows directly from Theorem 17 and 24. \square

As a final remark, similarly to Low and Chuang [LC17], we note that if one can perform a POVM measurement on a quantum computer, then one can also implement a block-encoding of

the corresponding measurement operator. First we clarify what we mean by performing a POVM measurement on a quantum computer. For simplicity assume that the POVM is a two-outcome measurement, represented by the operators $M, (I - M)$. Then we assume (without too much loss of generality) that an implementation on a quantum computer is as follows: We get as input a mixed state ρ , and attach a ancilla qubits to it. Then we apply some unitary on the state, and finally perform a measurement in the computational basis, accepting only measurement outcomes where the last qubit is $|0\rangle$. We can summarise the procedure as follows:

$$\rho \rightarrow \rho \otimes |0\rangle\langle 0|^{\otimes a} \rightarrow U(\rho \otimes |0\rangle\langle 0|^{\otimes a})U^\dagger \rightarrow \text{Tr}\left((I \otimes |0\rangle\langle 0|)U(\rho \otimes |0\rangle\langle 0|^{\otimes a})U^\dagger\right).$$

Suppose that the above implementation has at most ε bias, then

$$\forall \rho: \left| \text{Tr}\left((I \otimes |0\rangle\langle 0|)U(\rho \otimes |0\rangle\langle 0|^{\otimes a})U^\dagger\right) - \text{Tr}(M\rho) \right| \leq \varepsilon. \quad (14)$$

Observe that $\rho \otimes |0\rangle\langle 0|^{\otimes a} = (I \otimes |0\rangle\langle 0|^{\otimes a})\rho(I \otimes \langle 0|^{\otimes a})$, and thus

$$\text{Tr}\left((I \otimes |0\rangle\langle 0|)U(\rho \otimes |0\rangle\langle 0|^{\otimes a})U^\dagger\right) = \text{Tr}\left((I \otimes \langle 0|^{\otimes a})U^\dagger(I \otimes |0\rangle\langle 0|)U(I \otimes |0\rangle\langle 0|^{\otimes a})\rho\right).$$

Therefore (14) is equivalent to saying that

$$\begin{aligned} \forall \rho: \quad & \left| \text{Tr}\left(\left[(I \otimes \langle 0|^{\otimes a})U^\dagger(I \otimes |0\rangle\langle 0|)U(I \otimes |0\rangle\langle 0|^{\otimes a}) - M\right]\rho\right) \right| \leq \varepsilon \\ \iff & \left\| (I \otimes \langle 0|^{\otimes a})U^\dagger(I \otimes |0\rangle\langle 0|)U(I \otimes |0\rangle\langle 0|^{\otimes a}) - M \right\| \leq \varepsilon. \end{aligned}$$

Finally let $a' := a + 1$ and $U' := U \otimes I_2$, then it is easy to see that

$$(I \otimes \langle 0|^{\otimes a})U^\dagger(I \otimes |0\rangle\langle 0|)U(I \otimes |0\rangle\langle 0|^{\otimes a}) = (I \otimes \langle 0|^{\otimes a'})U'^\dagger(I \otimes \text{CNOT})U'(I \otimes |0\rangle\langle 0|^{\otimes a'}),$$

therefore $U'^\dagger(I \otimes \text{CNOT})U'$ is a $(1, a', \varepsilon)$ -block-encoding of M .

This shows that if we can implement the measurements E_j in a controlled fashion on a quantum computer, then we can also implement the corresponding block encoding with essentially the same cost. Hence the descriptive shadow tomography problem can be solved with the same cost as $(\sqrt{m} + \sqrt{n}/\varepsilon)/\varepsilon^4$ controlled measurements of E_j , if the measurement is performed on a quantum computer as we described above.

4.2 Quantum state discrimination

In the *Quantum State Discrimination* problem we are given k d -dimensional quantum states $\rho^{(1)}, \dots, \rho^{(k)} \in \mathbb{C}^{d \times d}$, in some oracular access model. Our goal is to find a POVM $M^{(1)}, \dots, M^{(k)}$ that has the “best” probability of discriminating between the states. Here “best” can mean two things:

- The minimal success probability is maximized: $\max_M \min_{i \in [k]} \text{Tr}(M^{(i)}\rho^{(i)})$.
- The total success probability, the sum of all the success probabilities, is maximized.

Both problems can be cast as an SDP [Eld03] but we will only consider the second here since it lends it self better to the Arora-Kale framework.. Our goal will be to get a quantum speedup in k at the expense of a slowdown in terms of d . However, the interesting cases of the problem seem to occur when $d \ll k$. Furthermore, we will use the quantum state model not to get a further speedup over the sparse matrix access oracle, but to show that it is possible to solve the problem even when just given access to unitaries that prepare the quantum states.

Theorem 27. *Given access to the matrix entries of the quantum states $\rho^{(1)}, \dots, \rho^{(k)} \in \mathbb{C}^{d \times d}$ the total error quantum state discrimination problem can be solved up to additive error ε on a quantum computer using*

$$\tilde{O}\left(\frac{\sqrt{k}}{\varepsilon^5} \text{poly}(d)\right)$$

queries to the input.

Given access to a unitary that creates a purified version of the quantum states $\rho^{(1)}, \dots, \rho^{(k)} \in \mathbb{C}^{d \times d}$ the total error quantum state discrimination problem can be solved up to additive error ε on a quantum computer using

$$\tilde{O}\left(\frac{k^{1.5}}{\varepsilon^5} \text{poly}(d)\right)$$

queries.

Proof. To maximize the total success probability, notice that the probability of measuring $\rho^{(i)}$ correctly is $\text{Tr}(M^{(i)}\rho^{(i)})$. Writing the problem as an SDP we get:

$$\begin{aligned} \max \quad & \sum_{i=1}^k \text{Tr}\left(M^{(i)}\rho^{(i)}\right) \\ \text{s.t.} \quad & \sum_{i=1}^k M^{(i)} = I_d \\ & M^{(i)} \succeq 0 \text{ for all } i \in [k]. \end{aligned}$$

This can be written in the standard form (1) as follows:

- $X = \text{diag}(M^{(1)}, \dots, M^{(k)})$.
- $C = \text{diag}(\rho^{(1)}, \dots, \rho^{(k)})$.
- $A_{st} = \bigoplus_{i=1}^k E_{st}$ and $b_{st} = \delta_{st}$, for δ_{st} the Kronecker delta.

Notice that we have strict equalities, as opposed to the inequalities in the standard form. These equalities can be cast into inequality form by adding a separate upper and lower bound, this is however not needed for the analysis. Just note that an equality in the primal corresponds to a variable in the dual without positivity constraint.

To apply our SDP-solvers we need to give bounds on the input parameters. Clearly here $n = kd$ and $m = \mathcal{O}(d^2)$. Furthermore, since the objective matrix is block diagonal with $d \times d$ blocks, the sparsity s is at most d . To bound B , note that C has trace k and is psd, and all other constraints can clearly be decomposed with a constant trace. It remains to give a bound for R and r .

For R , the bound on the trace of a primal solution, notice that

$$\text{Tr}(X) = \sum_{i=1}^k \text{Tr}(M^{(i)}) = \text{Tr}\left(\sum_{i=1}^k M^{(i)}\right) = \text{Tr}(I_d) = d.$$

For r we need to write out the dual, doing so directly gives:

$$\begin{aligned} \min \quad & \sum_{s,t}^d y_{st} \delta_{st} = \sum_s^d y_{ss} \\ \text{s.t.} \quad & \sum_{j=1}^m y_{st} \left(\oplus_{i=1}^k E_{st}\right) \succeq \oplus_{i=1}^k \rho^{(i)}. \end{aligned}$$

Notice that we do not have a $y \geq 0$ since we have strict equalities in the primal. We could have replaced the equalities by inequalities and then we would get a y^+ and y^- vector, both non-negative, such that $y = y^+ - y^-$. However, since r is a bound on the sum of the values in *one* optimal solution, it is enough to bound the absolute value of the y variables.

To do so, simply rewrite the dual in block form and reorganize the y variables in a matrix Y :

$$\begin{aligned} \min \quad & \text{Tr}(Y) \\ \text{s.t.} \quad & Y \succeq \rho^{(i)} \text{ for all } i \in [k]. \end{aligned}$$

Clearly I_d is feasible for this problem so for an optimal Y we have $\text{OPT} = \text{Tr}(Y) \leq d$. This gives the bound, for $S \in \{-1, 1\}^{d \times d}$

$$\sum_{s,t}^d |Y_{st}| = \max_{S \in \{-1, 1\}^{d \times d}} \text{Tr}(SY) \leq \text{Tr}(Y) \max_{S \in \{-1, 1\}^{d \times d}} \|S\| \leq d \text{Tr}(Y) \leq d^2, \quad (15)$$

so $r = d^2$ suffices.²⁰

Applying our results about the complexity of SDP-solving gives the claimed bounds. \square

Note that the output of the algorithm is a classical description of a dual solution Y and a concise classical description (Y', z) of a primal solution $M^{(i)} \propto e^{Y' - z\rho^i}$, such that the $M^{(i)}$ s form a close to optimal POVM. Note that this representation gives an interesting way of compressing a POVM, since the Y' matrix is only $d \times d$, and encodes k POVM operators with the help of the $\rho^{(i)}$ matrices. The dual solution Y could be of independent interest too, solving the following problem: for a set of density operators, find the matrix with the smallest trace that is psd bigger than all given density operators.

A lower bound. To find a lower bound, fix $d = 2$, i.e., consider a single qubit. Now let $z \in \{0, 1\}^k$ be the input for a search problem, we want to distinguish the cases $|z| = 0$ and $|z| = 1$ under the promise that we are in one of these cases. This is known to take $\Omega(\sqrt{k})$ quantum queries or $\Omega(k)$ classical queries. Now let $\rho^{(j)} = |z_j\rangle\langle z_j|$. Given query access to z it is easy to construct the input

²⁰This also proves that for k states of dimension d the total success probability of discrimination is always at most d , so the average will be at most d/k . Thus the error parameter should scale with $1/k$ if we would consider the average probability. This is why we choose to look at the total success probability instead.

oracles for any of the three input models. Clearly if $z = 0^k$ then all states are equal thus the total success probability is always 1. However, if $z_k = 1$ then by setting $M^{(k)} := |1\rangle\langle 1|$ and choosing the other measurement operators arbitrarily, we clearly get a total success probability of 2. Hence a $1/3$ -approximation to the optimal value of the SDP given above will solve the search problem and hence takes at least $\Omega(\sqrt{k})$ quantum queries or at least $\Omega(k)$ classical queries.

4.3 Optimal design

In the optimal design setting we want to learn a hidden state $\theta \in \mathbb{R}^d$ through experiments. There is a set of k possible experiments, represented by unit vectors $u^{(1)}, \dots, u^{(k)} \in \mathbb{R}^d$, and when we execute the i th experiment we learn $\langle \theta, u^{(i)} \rangle$ with some noise. In particular we get a sample from $\mathcal{N}(\langle \theta, u^{(i)} \rangle, \sigma_i)$. Precise estimation of θ requires a lot of experiments, and the problem in optimal design is to decide which distribution to use when choosing the experiments in order to “minimize” the covariance matrix of the maximum likelihood estimator of θ . Since the variance of the maximum likelihood estimator is hard to express analytically, we instead look at the *Fischer information matrix*, which is a good approximation for the inverse of the covariance matrix, and has a nice closed form:

$$F_p = \sum_{i=1}^k p_i u^{(i)} u^{(i)T} / \sigma_i^2,$$

where p_i is the probability of doing experiment $u^{(i)}$. Now, to get the covariance matrix “small” we would like to get the Fischer information matrix “large”. For a more detailed explanation, see for example [Sil80].

The precise meaning of “small” and “large” can be defined in several sensible ways. The most common criteria are called *A-optimal*, *D-optimal* and *E-optimal design*. In A-optimal design we want to minimize the sum of the eigenvalues of the covariance matrix, or as an approximation the sum of the eigenvalues of the inverse of the Fischer information matrix. Unfortunately the SDP formulation of this problem has parameters r, R that make our methods inefficient. In D-optimal design we want to minimize the determinant of the covariance matrix, this can be approximated with a convex program, but sadly this problem does not naturally correspond to an SDP.

We will consider E-optimal design. In this setting we would like to minimize the operator norm of the covariance matrix. Since this is hard to do, we will try to maximize the smallest eigenvalue of the Fischer information matrix. Let $P := \frac{1}{d \max_i \sigma_i^2}$ be an input parameter dependent on the precision of the experiments averaged over the coordinates.

Theorem 28. *The E-optimal design problem, that is, finding a distribution p such that the smallest eigenvalue of F_p is maximized, can be solved up to additive error ε using sparse access to the s -sparse experiment (unit) vectors $u^{(1)}, \dots, u^{(k)} \in \mathbb{R}^d$ and oracle access to the σ_i values with*

$$\tilde{O}\left(\left(\sqrt{k} + \sqrt{d} \frac{P^2}{\varepsilon}\right) s \frac{P^8}{\varepsilon^4}\right)$$

*queries on a quantum computer.*²¹

²¹The dependence on s can be reduced to \sqrt{s} by using state preparation and the quantum operator model.

Proof. We consider the following SDP:

$$\begin{aligned}
& \max && t \\
& \text{s.t.} && \sum_{i=1}^k p_i u^{(i)} u^{(i)T} / \sigma_i^2 \succeq t I_d \\
& && \sum_{i=1}^k p_i \leq 1 \\
& && p_i \geq 0 \text{ for all } i \in [k]
\end{aligned}$$

Clearly this SDP would maximize the minimal eigenvalue of F_p . We can rewrite this in standard dual form, flipping the sign of the optimal value:

$$\begin{aligned}
& \min && -t \\
& \text{s.t.} && \sum_{i=1}^k p_i \begin{bmatrix} -1 & \\ & u^{(i)} u^{(i)T} / \sigma_i^2 \end{bmatrix} + t \begin{bmatrix} 0 & \\ & -I_d \end{bmatrix} \succeq \begin{bmatrix} -1 & \\ & 0 \end{bmatrix} \\
& && p_i \geq 0 \text{ for all } i \in [k] \\
& && t \geq 0
\end{aligned}$$

The corresponding primal problem is then easy to write down:

$$\begin{aligned}
& \max && -z \\
& \text{s.t.} && -z + \text{Tr}\left(X u^{(i)} u^{(i)T}\right) / \sigma_i^2 \leq 0 \text{ for all } i \in [k] \\
& && \text{Tr}(X) \geq 1 \\
& && z \geq 0, X \succeq 0
\end{aligned}$$

From the size of the input it follows that $n = 1 + d$ and $m = 1 + k$ for this SDP. Furthermore, the row sparsity of the constraint matrices is equal to the vector sparsity of the $u^{(i)}$, which justifies the use of s for the sparsity of the vectors $u^{(i)}$. It remains to give a bound on r and R . Note that the trace constraint on X will be tight for an optimal X and hence $R = 1 + |\text{OPT}|$, where OPT is the optimal value of one of these SDPs. Similar for the sum constraint on p , we get $r = 1 + |\text{OPT}|$. To give a bound on $|\text{OPT}|$ we rewrite the primal again, flip the sign of the optimum and flipping the sign of z :

$$\begin{aligned}
& \min && z \\
& \text{s.t.} && \text{Tr}\left(X u^{(i)} u^{(i)T}\right) / \sigma_i^2 \leq z \text{ for all } i \in [k] \\
& && \text{Tr}(X) \geq 1 \\
& && z \geq 0, X \succeq 0
\end{aligned}$$

Now, let us construct a feasible point, since we have a minimization SDP, this will give an upper bound on $|\text{OPT}|$. Let $X = I_d/d$, then $\text{Tr}(X u^{(i)} u^{(i)T}) = 1/d$, so picking $z = \max_i \frac{1}{d\sigma_i^2}$ will give a feasible point with objective value z . We conclude that $r = R = \mathcal{O}\left(\frac{1}{\max_i d\sigma_i^2}\right)$ suffices. The stated complexity follows using our complexity bounds on quantum SDP-solving. \square

5 Lower bounds for the new input models

Previous work [AGGW17] showed an $\Omega\left(\sqrt{\max\{n, m\}} \min\{n, m\}^{3/2}\right)$ lower bound for the quantum query complexity of SDP-solving in the sparse input model. For this bound $s = 1$, $\varepsilon = 1/3$ and $R = r = \min\{n, m\}^2$. By letting either n or m be constant, the $\Omega(\sqrt{n} + \sqrt{m})$ lower bound from [BS17] can be recovered. The improved upper bounds of this paper show that the dependence on n and m is tight up to logarithmic factors. It remains an open question whether a lower bound with an interesting dependence on s and Rr/ε can be proven.

In this section we prove lower bounds for the new input models: the quantum state model and the quantum operator model. To do so, we first prove a lower bound in the Hamiltonian input model, where we can time-evolve under the matrices A_j , see Definition 29. In all cases the goal is to show that the term \sqrt{m}/ε times the relevant normalization parameter (for example B in the quantum state model) is necessary.

Definition 29 (Hamiltonian input model). *In the Hamiltonian input model for SDPs, we have access to two oracles for the A_j matrices. The first oracle, O_t , gives a classical description of a real vector $t \in \mathbb{R}^j$ in the usual way*

$$O_t|j\rangle|0\rangle = |j\rangle|t_j\rangle.$$

The second oracle, O_H , performs the Hamiltonian simulation with A_j for time $1/t_j$:

$$O_H|j\rangle|\psi\rangle = |j\rangle e^{iA_j/t_j}|\psi\rangle$$

Alongside the oracles we also require an upper bound $\tau \geq \max_j t_j$ as part of the input for an SDP. As in the other input models, we assume that we can also apply the inverse of the oracles.

Now we invoke a result of Gilyén et al. [GAW17, Theorem 2], which gives a lower bound on the number of queries needed for distinguishing different phase oracles. In the spirit of Definition 29, we will view this as the task of distinguish a set of diagonal Hamiltonians.

Theorem 30 (Hybrid method for arbitrary phase oracles). *Let G be a (finite) set of labels and let $\mathcal{H} := \text{Span}(|x\rangle : x \in G)$ be a Hilbert space. For a function $\tilde{f} : G \rightarrow \mathcal{R}$ let $O_{\tilde{f}}$ be the phase oracle acting on \mathcal{H} such that*

$$O_{\tilde{f}} : |x\rangle \rightarrow e^{i\tilde{f}(x)}|x\rangle \quad \text{for every } x \in G.$$

Suppose that \mathcal{F} is a finite set of functions $G \rightarrow \mathcal{R}$, and the function $f_ : G \rightarrow \mathcal{R}$ is not in \mathcal{F} . If a quantum algorithm makes T queries to a (controlled) phase oracle $O_{\tilde{f}}$ (or its inverse) and for all $f \in \mathcal{F}$ can distinguish with probability at least $2/3$ the case $\tilde{f} = f$ from the case $\tilde{f} = f_*$, then*

$$T \geq \frac{\sqrt{|\mathcal{F}|}}{3} \bigg/ \sqrt{\max_{x \in G} \sum_{f \in \mathcal{F}} |f(x) - f_*(x)|^2}.$$

Now the following lower bound follows naturally by reducing the above ‘‘Hamiltonian discrimination problem’’ to solving an SDP in the Hamiltonian input model.

Lemma 31. *Let $\varepsilon \in (0, 1/2]$, $2 \leq m$ and $1 \leq \tau$. Then there is an LP (and hence an SDP) (with $R, r = \mathcal{O}(1)$) for which an ε -approximation of the optimal value requires $\Omega(\sqrt{m} \frac{\tau}{\varepsilon})$ queries to O_H in the Hamiltonian input model.*

Proof. Let $H_1, \dots, H_m \in \mathbb{R}^{2 \times 2}$ be such that

(a) either all H_j are $I/(2\tau)$,

(b) or all but one matrices are $I/(2\tau)$, and there is one H_j that is equal to

$$\begin{bmatrix} 1/(2\tau) + \varepsilon/\tau & 0 \\ 0 & 1/(2\tau) - \varepsilon/\tau \end{bmatrix} = \frac{1}{2\tau}I + \frac{\varepsilon}{\tau}Z.$$

Let us assume that we have access to the phase oracle $O : |j\rangle|b\rangle \rightarrow e^{i(H_j)_{bb}}|j\rangle|b\rangle$. In case (b) there are m possible different choices for this oracle. It is easy to see by Theorem 30 that distinguishing case (a) from (b) requires $\Omega(\sqrt{m}\frac{\tau}{\varepsilon})$ queries to O .

Now we show that using the above phase oracles we can define an SDP in the Hamiltonian input model, solving which to ε -precision distinguishes case (a) from (b), proving the sought lower bound.

Let us define $A_j := \tau H_j$ (and $t = (\tau, \dots, \tau)$), so all the A_j 's are either $I/2$ or $I/2 + \varepsilon Z$, furthermore let

$$C = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix},$$

and $b = (1, \dots, 1)^T$ the all-one vector. Note that since at least one of the A_j matrices is $I/2$, we know that $R = 2$ suffices as an upper bound on the trace. Furthermore, it is easy to verify from the dual that $r = 2$ suffices as well.

Now we analyze the optimal value. If all A_j matrices are $I/2$ then all constraints are the same:

$$X_{11}/2 + X_{22}/2 \leq 1$$

from which it clearly follows that $\text{OPT} = 2$.

If one A_j matrix is not $I/2$, then the constraint

$$(1/2 + \varepsilon)X_{11} + (1/2 - \varepsilon)X_{22} \leq 1$$

is present. It follows that

$$(1/2 + \varepsilon)X_{11} \leq 1 \Rightarrow X_{11} \leq \frac{1}{1/2 + \varepsilon}$$

which will clearly be tight in the optimum. Using that

$$2 - 4\varepsilon \leq \frac{1}{1/2 + \varepsilon} \leq 2 - 2\varepsilon$$

we conclude that $2 - 4\varepsilon \leq \text{OPT} \leq 2 - 2\varepsilon$.

Hence solving this SDP up to precision ε will distinguish case (a) from (b) and hence requires $\Omega(\sqrt{m}\frac{\tau}{\varepsilon})$ queries. \square

To prove the lower bounds for the quantum state model and the quantum operator model we reduce the Hamiltonian input model to them.

Lemma 32. *Let $\varepsilon \in (0, 1]$. Given an SDP in the Hamiltonian input model with parameter $\tau \geq 2$ (for technical reasons also assume that $t_j \geq 2$ for all j), an ε -approximate oracle call in the quantum operator model with $\alpha = 2\tau$ can be simulated using $\tilde{O}_\varepsilon(1)$ queries.*

Proof. For simplicity let us drop the index j . Let $H' := H/t$, and apply Theorem 37 to H' with $f(x) = x$, setting $x_0 = 0$, $r = 1$, $\delta = \pi/2 - 1$, and $K = 2$, providing a $(2, \tilde{\mathcal{O}}_\varepsilon(1), \varepsilon)$ -block encoding of H' via controlled $(\mathcal{O}(\log(1/\varepsilon)), 1)$ -Hamiltonian simulation, which can be easily implemented by $\mathcal{O}(\log(1/\varepsilon))$ controlled oracle calls. \square

Corollary 33. *Let $\varepsilon \in (0, 1/2]$, $2 \leq m$ and $2 \leq \alpha$. Then there is an LP (and hence an SDP) (with $R, r = \mathcal{O}(1)$) for which an ε -approximation of the optimal value requires $\Omega(\sqrt{m} \frac{\alpha}{\varepsilon})$ queries to O_H in the quantum operator model.*

For the quantum state input model we only give a reduction for LPs, i.e., the case where all input matrices are diagonal.

Lemma 34. *Let $\varepsilon \in (0, 1]$. Given an LP in the Hamiltonian input model with parameter $\tau \geq 4$ (for technical reasons also assume that $t_j \geq 4$ for all j), then an ε -approximate oracle call in the quantum state model with $B = n\tau$ can be simulated using $\tilde{\mathcal{O}}_\varepsilon(1)$ queries.*

Proof. For simplicity let us drop the index j . Let $H' := I \pm H/t$, and apply Theorem 37 to H' with $f(x) = \sqrt{x/2}$, setting $x_0 = 1$, $r = 1/2$, $\delta = \pi/6 - 1/2$, and $K = 1$, providing a $(1, \tilde{\mathcal{O}}_\varepsilon(1), \mathcal{O}(\varepsilon))$ -block encoding of $\sqrt{(I + H')/2}$, via controlled $(\mathcal{O}(\log(1/\varepsilon)), 3)$ -Hamiltonian simulation, which can be easily implemented by $\mathcal{O}(\log(1/\varepsilon))$ controlled oracle calls. The validity of the $K = 1$ bound follows from the observation that

$$\sqrt{(1+x)/2} = \frac{1}{\sqrt{2}} \sqrt{1+x} = \frac{1}{\sqrt{2}} \sum_{k=0}^{\infty} \binom{1/2}{k} x^k,$$

thus

$$\frac{1}{\sqrt{2}} \sum_{k=0}^{\infty} \left| \binom{1/2}{k} \right| (r + \delta)^k = \frac{1}{\sqrt{2}} \sum_{k=0}^{\infty} \left| \binom{1/2}{k} \right| (\pi/6)^k = \sqrt{(1 + \pi/6)/2} \leq 1 =: K.$$

The state input oracles can be implemented as follows: controlled on the state $|\pm\rangle$ we apply $\sqrt{(I \pm H/t)/2}$ to the first half of the state $\sum_{i=1}^n |i\rangle|i\rangle/\sqrt{n}$, resulting in subnormalized density operators $\varrho_{\pm} = (I \pm H/t)/(2n)$, so that $\varrho_+ - \varrho_- = H/(nt)$. If needed one can further subnormalize the ϱ -s in order to get a uniform normalization factor $n\tau$ (instead of nt_j). \square

Corollary 35. *Let $\varepsilon \in (0, 1/2]$, $2 \leq m$ and $1 \leq B$. Then there is an LP (and hence an SDP) (with $R, r = \mathcal{O}(1)$) for which an ε -approximation of the optimal value requires $\Omega(\sqrt{m} \frac{B}{\varepsilon})$ queries to O_H in the quantum state model.*

Acknowledgments. We thank the authors of [BKL⁺18] for sending work-in-progress versions of their paper, and Fernando Brãndao, Tongyang Li and Xiaodi Wu for personal communication. A.G. thanks Robin Kothari and Nathan Wiebe for useful discussions. We thank Jamie Sikora for useful discussions about applications and for suggesting the state discrimination problem. We are grateful to Ronald de Wolf and Sander Gribling for useful discussions, and advice about the manuscript.

References

- [Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC 2009, Paris, France, 15-18 July 2009*, pages 229–242, 2009. 10
- [Aar17] Scott Aaronson. Shadow Tomography of Quantum States, 2017, arXiv: 1711.01053. To appear in STOC’18. 1, 7, 9, 12, 13, 23, 24
- [AGGW17] J. van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf. Quantum SDP-solvers: Better upper and lower bounds. In *Proceedings of 58th IEEE FOCS*, pages 403–414, 2017, arXiv: 1705.01843. 1, 6, 7, 8, 9, 10, 11, 14, 15, 16, 17, 20, 21, 23, 30, 35
- [AK16] S. Arora and S. Kale. A combinatorial, primal-dual approach to semidefinite programs. *Journal of the ACM*, 63(2):12, 2016. Earlier version in STOC’07. 6, 10
- [BCK15] D. W. Berry, A. M. Childs, and R. Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. *Proceedings of 56th IEEE FOCS*, pages 792–809, 2015, arXiv: 1501.01715. 4, 15, 35
- [Bel15] Aleksandrs Belovs. Variations on quantum adversary. 2015, arXiv: 1504.06943. 10
- [Bes05] Arvid J. Bessen. Lower bound for quantum phase estimation. *Physical Review A*, 71(4):042313, 2005. 10
- [BKL⁺17] Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Exponential quantum speed-ups for semidefinite programming with applications to quantum learning, 2017, arXiv: 1710.02581v1. First arXiv version. 7, 9
- [BKL⁺18] Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning, 2018, arXiv: 1710.02581v2. 1, 7, 8, 9, 11, 12, 13, 18, 22, 23, 24, 32
- [BS17] Fernando G. S. L. Brandão and Krysta M. Svore. Quantum speed-ups for semidefinite programming. In *Proceedings of 58th IEEE FOCS*, pages 415–426, 2017, arXiv: 1609.05537. 1, 2, 6, 7, 8, 9, 10, 11, 23, 30
- [CGJ18] Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. The power of block-encoded matrix powers: improved regression techniques via faster Hamiltonian simulation. 2018, arXiv: 1804.01973. 15, 19, 22, 35, 36
- [CHTW04] R. Cleve, P. Høyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *19th IEEE Annual Conference on Computational Complexity, 2004. Proceedings*, pages 236–249, 2004. 2

- [DH96] C. Dürr and P. Høyer. A quantum algorithm for finding the minimum. 18 Jul 1996, arXiv: [quant-ph/9607014](https://arxiv.org/abs/quant-ph/9607014). 4, 11, 17, 23
- [Eld03] Y. C. Eldar. A semidefinite programming approach to optimal unambiguous discrimination of quantum states. *IEEE Transactions on Information Theory*, 49:446–456, 2003, arXiv: [quant-ph/0206093](https://arxiv.org/abs/quant-ph/0206093). 2, 26
- [GAW17] András Gilyén, Srinivasan Arunachalam, and Nathan Wiebe. Optimizing quantum optimization algorithms via faster quantum gradient computation, 2017, arXiv: [1711.00465v3](https://arxiv.org/abs/1711.00465v3). 10, 30
- [GW95] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. Earlier version in STOC’94. 2, 6
- [GW18] András Gilyén and Nathan Wiebe. Unpublished manuscript. Work in progress, 2018. 14, 16
- [HLM17] Aram W. Harrow, Cedric Yen-Yu Lin, and Ashley Montanaro. Sequential measurements, disturbance and property testing. In *Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 1598–1611, 2017, arXiv: [1607.03236](https://arxiv.org/abs/1607.03236). 12
- [KP17] Iordanis Kerenidis and Anupam Prakash. Quantum recommendation systems. In *Innovations in Theoretical Computer Science*, 2017, arXiv: [1603.08675](https://arxiv.org/abs/1603.08675). 8, 17
- [Las01] Jean Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001. 2
- [LC16] Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by qubitization. 2016, arXiv: [1610.06546](https://arxiv.org/abs/1610.06546). 14, 15
- [LC17] Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by uniform spectral amplification. 2017, arXiv: [1707.05391](https://arxiv.org/abs/1707.05391). 24
- [LRS15] J. R. Lee, P. Raghavendra, and D. Steurer. Lower bounds on the size of semidefinite programming relaxations. In *Proceedings of 47th ACM STOC*, pages 567–576, 2015. 7, 11
- [LSW15] Y. T. Lee, A. Sidford, and S. C. Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. In *Proceedings of 56th IEEE FOCS*, pages 1049–1065, 2015, arXiv: [1508.04874](https://arxiv.org/abs/1508.04874). 6
- [Par00] Pablo Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, 2000. 2
- [Sil80] S. D. Silvey. *Optimal Design*. Chapman and Hall Ltd, 1980. 28
- [Wu17] Xiaodi Wu. Personal communication. Email, November 2017. 7

A Implementing smooth functions of Block-Hamiltonians

In [AGGW17, Appendix B] techniques were developed that make it possible to implement smooth-functions of a Hamiltonian H , based on Fourier series decompositions and using the Linear Combinations of Unitaries (LCU) Lemma [BCK15]. The techniques developed in [AGGW17, Appendix B] access H only through controlled-Hamiltonian simulation, which is defined in the following way:

Definition 36. Let $M = 2^J$ for some $J \in \mathbb{N}$, $\tau \in \mathbb{R}$ and $\epsilon \geq 0$. We say that the unitary

$$W := \sum_{m=-M}^{M-1} |m\rangle\langle m| \otimes e^{im\tau H}$$

implements controlled (M, τ) -simulation of the Hamiltonian H , where $|m\rangle$ denotes a (signed) bit-string $|b_J b_{J-1} \dots b_0\rangle$ such that $m = -b_J 2^J + \sum_{j=0}^{J-1} b_j 2^j$.

The main theorem of [AGGW17, Appendix B] states the following:

Theorem 37 (Implementing a smooth function of a Hamiltonian [AGGW17, Theorem 40]). Let $x_0 \in \mathbb{R}$ and $r > 0$ be such that $f(x_0+x) = \sum_{\ell=0}^{\infty} a_{\ell} x^{\ell}$ for all $x \in [-r, r]$. Suppose $K > 0$ and $\delta \in (0, r]$ are such that $\sum_{\ell=0}^{\infty} (r+\delta)^{\ell} |a_{\ell}| \leq K$. If $\|H - x_0 I\| \leq r$ and $\epsilon' \in (0, \frac{1}{2}]$, then we can implement a unitary \tilde{U} that is a $(K, a + \mathcal{O}(\log(r \log(1/\epsilon')/\delta)), K\epsilon')$ -block-encoding of $f(H)$, with a single use of a circuit V which is a $(1, a, \epsilon'/2)$ -block-encoding of controlled $(\mathcal{O}(r \log(1/\epsilon')/\delta), \frac{\pi}{2(r+\delta)})$ -simulation of H , and using $\mathcal{O}(r/\delta \log(r/(\delta\epsilon')) \log(1/\epsilon'))$ two-qubit gates.

This theorem can in particular be applied to the function $f(x) = x^{-c}$, a power function of a Hamiltonian with negative exponent:

Corollary 38 ([CGJ18]). Let $\kappa \geq 2$, $c \in (0, \infty)$ and H be a w -qubit Hamiltonian such that $I/\kappa \preceq H \preceq I$.

Then we can implement a unitary \tilde{U} that is a $(2\kappa^c, a + \mathcal{O}(\log(\kappa^c \max(1, c) \log(\kappa^c/\epsilon))), \epsilon)$ -block-encoding of H^{-c} , with a single use of a circuit V which is a $(1, a, \epsilon/(4\kappa^c))$ -block-encoding of controlled $(\mathcal{O}(\kappa \max(1, c) \log(\kappa^c/\epsilon)), \mathcal{O}(1))$ -simulation of H , and using $\mathcal{O}(\kappa \max(1, c) \log^2(\kappa^{1+c} \max(1, c)/\epsilon))$ other two-qubit gates.

We do not proof this here, but a full proof can be found in [CGJ18]. Similarly they get a result about implementing power functions of positive exponents.

Corollary 39 ([CGJ18]). Let $\kappa \geq 2$, $c \in (0, 1]$ and H be an s -qubit Hamiltonian such that $I/\kappa \preceq H \preceq I$.

Then we can implement a unitary \tilde{U} that is a $(2, a + \mathcal{O}(\log \log(1/\epsilon)), \epsilon)$ -block-encoding of H^c , with a single use of a circuit V which is a $(1, a, \epsilon/4)$ -block-encoding of controlled $(\mathcal{O}(\kappa \log(1/\epsilon)), \mathcal{O}(1))$ -simulation of H , and using $\mathcal{O}(\kappa \log(\kappa/\epsilon) \log(1/\epsilon))$ other two-qubit gates.

We can use these results to show that for a quantum state ρ it is possible to sample from a random variable with expectation value θ -close to $\text{Tr}(A\rho)$ and with standard deviation $\sigma \leq 6$, using $(\tilde{\mathcal{O}}_{\theta}(1), \mathcal{O}(1))$ -controlled Hamiltonian simulation of A . For more details see Definition 4.

Corollary 40. Suppose $-I \preceq A \preceq I$, $0 < \theta < 1$. We can implement a trace estimator for A with standard deviation $\sigma \leq 6$ and bias $\leq \theta$ with a single use of an $(\tilde{\mathcal{O}}_{\theta}(1), \mathcal{O}(1))$ -controlled Hamiltonian simulation circuit for A , and with $\tilde{\mathcal{O}}_{\theta}(1)$ other two-qubit gates.

Proof. We apply Corollary 39 to the square root function with the operator $I/2 + A/4$ (so that $\kappa \geq 4$), giving a $(2, \tilde{\mathcal{O}}(1), \Theta(\theta))$ -block-encoding \tilde{U} of $\sqrt{I/2 + A/4}$ using an $(\tilde{\mathcal{O}}_\theta(1), \mathcal{O}(1))$ -controlled Hamiltonian simulation circuit for $I/2 + A/4$. The probability of finding the ancilla qubit to be $|0\rangle$ upon measurement is

$$\begin{aligned} \text{Tr}\left((I \otimes \langle 0|)\tilde{U}^\dagger(\rho \otimes |0\rangle\langle 0|)\tilde{U}(I \otimes |0\rangle)\right) &= \text{Tr}\left(\underbrace{(I \otimes \langle 0|)\tilde{U}(I \otimes |0\rangle)}_{\approx \frac{\sqrt{I/2+A/4}}{2}}\underbrace{(I \otimes \langle 0|)\tilde{U}^\dagger(I \otimes |0\rangle)\rho}_{\approx \frac{\sqrt{I/2+A/4}}{2}}\right) \\ &= \frac{1}{8} + \frac{\text{Tr}(A\rho)}{16} + \mathcal{O}(\theta). \end{aligned}$$

Upon measuring the ancilla qubit and getting outcome $|0\rangle$ we output $16 - 2 = 14$. In case of any other measurement outcome we output -2 . By choosing the right constants so that \tilde{U} is a precise enough block-encoding we can ensure that the bias is less than $\theta/2$, and the standard deviation $\sigma \leq 6$.

Finally observe that a controlled Hamiltonian simulation circuit for $I/2 + A/4$ can be implemented as a product of controlled Hamiltonian simulation circuits for $I/2$ and $A/4$. \square

The following lemma from [CGJ18] gives a way to implement the controlled Hamiltonian simulation used in this appendix, provided a block-encoding of H , as used in this paper. In particular, it shows how to make Theorem 9 controlled by a time parameter.

Lemma 41 ([CGJ18]). *Let $M = 2^J$ for some $J \in \mathbb{N}$, $\tau \in \mathbb{R}$ and $\epsilon \geq 0$. Suppose that U is an $(\alpha, a, \epsilon/|2(J+1)^2 M\tau|)$ -block-encoding of the Hamiltonian H . Then we can implement a $(1, a+2, \epsilon)$ -block-encoding of a controlled (M, τ) -simulation of the Hamiltonian H , with $\mathcal{O}\left(|\alpha M\tau| + J \frac{\log(J/\epsilon)}{\log \log(J/\epsilon)}\right)$ uses of controlled- U or its inverse and with $\mathcal{O}\left(a|\alpha M\tau| + aJ \frac{\log(J/\epsilon)}{\log \log(J/\epsilon)}\right)$ three-qubit gates.*