

Quantum Speed-ups for Semidefinite Programming

FERNANDO G.S.L. BRANDÃO¹ AND KRISTA M. SVORE²

¹*Institute of Quantum Information and Matter, California Institute of Technology, Pasadena, CA*

²*Station Q Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA*

Abstract

We give a quantum algorithm for solving semidefinite programs (SDPs). It has worst-case running time $n^{\frac{1}{2}}m^{\frac{1}{2}}s^2 \text{poly}(\log(n), \log(m), R, r, 1/\delta)$, with n and s the dimension and row-sparsity of the input matrices, respectively, m the number of constraints, δ the accuracy of the solution, and R, r a upper bounds on the size of the optimal primal and dual solutions. This gives a square-root unconditional speed-up over any classical method for solving SDPs both in n and m . We prove the algorithm cannot be substantially improved (in terms of n and m) giving a $\Omega(n^{\frac{1}{2}} + m^{\frac{1}{2}})$ quantum lower bound for solving semidefinite programs with constant s, R, r and δ .

We then argue that in some instances the algorithm offers even exponential speed-ups. This is the case whenever the quantum Gibbs states of Hamiltonians given by linear combinations of the input matrices of the SDP can be prepared efficiently on a quantum computer. An example are SDPs in which the input matrices have low-rank: for SDPs with the maximum rank of any input matrix bounded by rank , we show the quantum algorithm runs in time $\text{poly}(\text{rank}, \log(n), \log(m), \text{rank}, r, R, 1/\delta)m^{\frac{1}{2}}$.

The quantum algorithm is constructed by a combination of quantum Gibbs sampling and the multiplicative weight method. In particular it is based on a classical algorithm of Arora and Kale for approximately solving SDPs. We present a modification of their algorithm to eliminate the need for solving an inner linear program which may be of independent interest.

1 Introduction

Quantum computers harness the unique features of quantum mechanics to compute in novel ways that outperform classical methods. In the past 20 years a variety of quantum algorithms offering speed-ups over classical computations have been found, including Shor's polynomial-time quantum algorithm for factoring [1] and Grover's quantum algorithm for searching a database in time square-root its size [2]. A central challenge in quantum computing is to identify more quantum algorithms beating classical computing, especially for practically relevant problems.

Semidefinite programming is one of the most successful algorithmic frameworks of the past few decades [3]. It has applications ranging from designing efficient algorithms for approximating combinatorial optimization problems [4] to operations research and beyond [5]. The power of semidefinite programs (SDPs) resides in their generality, together with the fact that there are efficient methods for solving them [5]. However, given the steadily increasing sizes of SDPs found in practice, it is an important problem to find even more efficient algorithms.

In this paper we give a quantum algorithm for solving semidefinite programming achieving a quadratic speed-up over any classical method, both in the dimension of the matrices and in the number of constraints of the program (we note however that the algorithm run-time also depends on a parameter measuring the size of the solution of the SDP, discussed below). Our work also gives the first quantum speed-up for linear programming, which is an important subclass of SDPs. Although we show that such quadratic speed-up is not far from the best possible in general, we argue the algorithm sometime offers even exponential speed-ups. Concretely we show this is the case for SDPs with low-rank input matrices (if the input is given in a proper way; see next section). We believe the results of this paper make a compelling case that solving SDPs quickly has the potential to be a relevant application of future quantum computers.

1.1 Semidefinite Programs

A general semidefinite program (SDP) is given by

$$\begin{aligned} & \max \operatorname{tr}(CX) \\ \forall j \in [m], & \operatorname{tr}(A_j X) \leq b_j \\ & X \geq 0, \end{aligned} \tag{1}$$

where the $n \times n$ Hermitian matrices (C, A_1, \dots, A_m) and the real numbers (b_1, \dots, b_m) are the inputs of the problem. The optimization is taken over positive semidefinite $n \times n$ matrices X . The dual program of the primal SDP given by Eq. (1) is the following:

$$\begin{aligned} & \min b \cdot y \\ & \sum_{j=1}^m y_j A_j \geq C \\ & y \geq 0, \end{aligned} \tag{2}$$

where the minimization is taken over vectors $y := (y_1, \dots, y_m)$. Under mild conditions the primal and dual problems have the same optimal value [5]. We can assume that

$$\|A_i\| \leq 1 \quad \forall i \in [m], \quad \text{and} \quad \|C\| \leq 1, \tag{3}$$

with $\|*\|$ the operator norm. This is without loss of generality by normalizing b_i and the optimal solution appropriately.

There are several different classical polynomial-time algorithms for solving semidefinite programs. One is the class of interior point methods. The state-of-the-art algorithm (in terms of rigorous worst-case bounds) has running time $\tilde{O}(m(m^2 + n^\omega + mns) \log(1/\delta))$ [7], with ω the exponent of matrix multiplication, s the row-sparsity of the matrices (C, A_1, \dots, A_m) (i.e., the maximum number of non-zero entries in each of the rows of the matrices), and δ the accuracy of the solution.

If one is willing to tolerate a worse scaling with error, faster algorithms can be sometimes obtained using the multiplicative weight method [8, 9, 10]. In particular Arora and Kale gave an algorithm for solving the SDP given by Eq. (1) in time $\tilde{O}(\omega^2 R^2 mns / \delta^2)$ [9], where ω is a parameter of the problem called *width* of the SDP, and R is an upper bound to the trace of the optimal X . Both R and ω can be shown to be small for many SDPs of interest [9].

It is an important open problem to find even more efficient algorithms for solving SDPs. Nonetheless, as we show in Section 2.3, a limit for any such improvement is the lower bound of $\Omega(n + m)$ for SDPs with $s, \omega, R = O(1)$.

1.2 Problem Statement

The problem we want to solve is the following: Given a list of $n \times n$ matrices $(A_1, \dots, A_m, A_{m+1})$ (with $A_{m+1} := C$), approximate the optimal value of Eqs. (1) and (2) and output optimal primal and/or dual solutions. To formulate the problem precisely we must specify in which form the inputs and outputs are given. We consider two different models of accessing the input:

Input Model 1: We assume there is an oracle \mathcal{P}_A that given the indices $j \in [m + 1]$, $k \in [n]$ and $l \in [s]$, computes a bit string representation of the l 'th non-zero element of the k -th row of A_j ,¹ i.e. the oracle performs the following map:

$$|j, k, l, z\rangle \rightarrow |j, k, l, z \oplus (A_j)_{k f_{jk}(l)}\rangle. \quad (4)$$

with $f_{jk} : [r] \rightarrow [N]$ a function (parametrized by the matrix index j and the row index k) which given $l \in [s]$ computes the column index of the l -th non-zero entry.

Input Model 2: This model is more unusual, but will be relevant to argue that sometimes large speed-ups are possible. We assume there is a quantum oracle which prepares copies of the eigenstates of the input matrices, and oracles for the eigenvalues of the input matrices and the numbers b_i . In more details, for $i \in \{0, \dots, n\}$, let

$$A_i = \sum_{l=1}^{r_i} \kappa_l^i |\eta_l^i\rangle \langle \eta_l^i|, \quad (5)$$

be the spectral decomposition of A_i , with $r_i \leq r$. Then we assume there is an oracle which given i, l , outputs an approximation (to accuracy $\nu > 0$) of a copy of the quantum state $|\eta_l^i\rangle$ and the eigenvalue κ_l^i (for an arbitrary ordering of the eigenvalues of the matrix). One example is when there are polylog(n)-sized quantum circuits which create the states $|\eta_l^i\rangle$. Then the oracle can give the description of the circuits.

Output: One way to specify the output is to require a list with the entries of X or y . However it is clear that in such a case at least $n(n - 1)/2$ (or m) time is required even to write down a primal (or dual) solution. Therefore it is necessary to relax the format of the output in order to obtain faster algorithms.

We require the quantum algorithm provides the following:

- An estimate of the optimal objective value.
- An estimate of $\|y\|_1$ and/or $\text{tr}(X)$.
- *Samples* from the distribution $p := y/\|y\|_1$ and/or from the quantum state $\rho := X/\text{tr}(X)$.²

¹We assume all elements of the input matrices can be represented exactly by a bit string of size polylog(n, m). If not we can truncate the matrices, which will only incur error $\exp(-\text{polylog}(n, m))$ that can be neglected.

²In this paper we present a quantum algorithm producing an estimate of $\|y\|_1$ and samples from the probability distribution $p := y/\|y\|_1$. The algorithm can be modified to also generate an estimate of $\text{tr}(X)$ and samples from $\rho := X/\text{tr}(X)$. However, we leave the details of this improvement to a future version of the paper.

As we show in Section 2.3, $\Omega(n + m)$ calls to the oracle are required even classically to output a solution as above.

2 Results

In this paper we give the first quantum algorithm for solving SDPs offering a speed-up over classical methods. Below we state the main contributions on a high level, describe the algorithm, and present a few open questions related to it.

2.1 Main Ideas

The first contribution of the paper is to notice that classical algorithms for solving SDPs based on the multiplicative weight method [9] imply that in order to solve the SDP of Eq. (1), it is enough to prepare quantum thermal (Gibbs) states of Hamiltonians given by linear combinations of the input matrices A_1, \dots, A_m, C of the program. Therefore in cases where such Gibbs states can be prepared efficiently (in time polynomial in $\log(n)$), quantum computers can give exponential speed-ups. This already suggests that our method might be an interesting heuristic to run on quantum computers (e.g., by using quantum Metropolis sampling [11, 12] to prepare the Gibbs states).

The second contribution is to combine the first observation with amplitude amplification [17] in the preparation of the Gibbs state to achieve a generic quadratic speed-up in terms of n , the dimension of the input matrices of the program. Here we can apply known results [13, 14] on using amplitude amplification to prepare Gibbs states on a quantum computer in time given roughly by the square root of the dimension of the system.

The third contribution is to show one can achieve a quadratic speed-up also in m , the number of constraints of the program. Establishing this fact requires more work. We modify the Arora-Kale algorithm and replace the inner linear program they use by the preparation of a Gibbs state of a classical Hamiltonian, whose entries are estimated from the expectation value (with each of the input matrices) of the Gibbs state prepared in the main thread of the algorithm. We show this replacement is possible using Jaynes' principle of maximum entropy [18], or more specifically a recent approximate version of it [19] with better control of parameters. Then we apply amplitude amplification also to the preparation of this Gibbs state. This modification alone is not enough to give a speed-up, since the sparsity of the Hamiltonians for which we must prepare the associated Gibbs states also depends on m (and the quantum algorithms for preparing Gibbs states we consider have linear dependence on sparsity). We then show that we can sparsify the Hamiltonians under consideration by random sampling, without changing the functioning of the algorithm, so that their sparsity only depends on the original sparsity s of the input matrices (up to polylogarithmic factors in n, m).

The final contribution is to analyse the algorithm assuming the second model of accessing the data discussed in Section 1.2 (in which one is given a quantum oracle for the eigenstates of the input matrices, together with oracles for the eigenvalues and the b_i 's). We show that for SDPs in which all input matrices have rank at most r , we can solve the problem with a quantum computer in time $\text{poly}(\log(n), \log(m), r, R, \alpha, 1/\delta)m^{\frac{1}{2}}$, which is an exponential speed-up in terms of n for small r, R, α . The idea is to show that the Gibbs sampling of the associated Hamiltonians can be done in $\text{poly}(\log(n), \log(m)r)$ time on a quantum computer.

2.2 The Algorithm

Reduction to Feasibility: Using binary search we can reduce the optimization problem to a feasibility one. Let α be a guess for the optimal solution (which will be varied by binary search). We are then concerned with the problem of either sampling from a probability distribution $p := y / \|y\|_1$, with y a dual feasible vector whose value is at most $\alpha(1 + \delta)$ for a small $\delta > 0$, or finding out that the optimal value is larger than $\alpha(1 - \delta)$.³

Gibbs Samplers: A subroutine of the main algorithm is the following:

Definition 1 (Gibbs Sampler). *Let H be a Hamiltonian and $O[H]$ an oracle for its entries.⁴ Then $\text{GibbsSampler}(O[H], \nu)$ is a quantum operation that given access to $O[H]$, outputs a state ρ such that $\|\rho - e^H / \text{tr}(e^H)\|_1 \leq \nu$.*

Several different Gibbs samplers have been proposed in the literature [11, 12, 13, 14, 15, 16], and any of them could be used in the main quantum algorithm.

Oracle by Estimation: Given a Hamiltonian H , in order to run a Gibbs sampler algorithm, we need an oracle for its entries. We also need the notion of a probabilistic oracle. This is an oracle that with high probability outputs the right entry of the Hamiltonian, but with small probability might output a wrong value.

Consider a quantum state ρ and two real numbers λ and μ . We define the Hamiltonian $h(\rho, \lambda, \mu) := \sum_{i=1}^m r_i |i\rangle \langle i|$, with

$$r_i := \lambda \text{tr}(A_i \rho) + \mu b_i, \quad (6)$$

and its truncated version

$$\bar{h}(\rho, \lambda, \mu) := \sum_{i=1}^m \bar{r}_i |i\rangle \langle i|, \quad (7)$$

with \bar{r}_i the rounding of r_i to precision $h_{\text{precision}}$. Throughout the paper we set

$$h_{\text{precision}} = \frac{\delta}{56R^2}. \quad (8)$$

The quantum algorithm will make use of calls to a probabilistic oracle for \bar{h} , which we show how to construct in Section 4. A subtlety is that ρ will not be given explicitly as a density matrix, but only as a quantum state. Therefore in order to implement the oracle for \bar{h} , we need to first estimate (some of) the values $\{\text{tr}(A_i \rho)\}_{i=1}^m$.

The Size Parameter R : Apart from the dimension of the matrices n , the number of constraints m , the sparsity of the input matrices s , and the error δ , the algorithm will depend on another parameter of the SDP. For many problems of interest, this is a constant independent of n and m .⁵

Following [9], we assume $A_1 = I$ and let $b_1 = R$. Thus we have the constraint

$$\text{tr}(X) \leq R. \quad (9)$$

³Alternatively we could also sample from a quantum state $\rho := X / \text{tr}(X)$, with X a primal feasible solution with objective value at least $\alpha(1 - \delta)$; however we do not consider this task in the current version.

⁴In analogy with the input model 1, $O[H]$ is an oracle that given the indices $k \in [n]$ and $l \in [s]$, computes a bit string representation of the l 'th non-zero element of the k -th row of H . Here n is the dimension of H and s its sparsity.

⁵We remind the reader we are also assuming that $\|C\|, \|A_i\| \leq 1$ for all $i \in [m]$ (which is w.l.o.g. by changing the values of the numbers b_j, α appropriately).

We can always add this constraint without changing the optimal solution by choosing R sufficiently large. The parameter R is a measure of the size of the optimal solution of the SDP. Note we have the upper bound $\alpha \leq R$.⁶ We also assume R is chosen sufficiently large such that $\max_i |b_i| = R$.

Reduction to $b_i \geq 1$ and $\alpha \geq 1$ and the dual size parameter r : Our algorithm will only work for SDPs for which $b_i \geq 1$ for all $i \in [m]$. However in Appendix A we prove Lemma 2 below, which shows that if we can solve SDPs with the b_i 's larger than one, we can solve arbitrary SDPs in roughly the same time. We say a feasible solution is δ -optimal if its objective value is within additive error δ to the optimal. Given the SDP of Eq. (2) with an optimal solution (y_1, \dots, y_m) , we assume we are given an upper bound r to the sum $\sum_{i=1}^m y_i$ (if there is more than one solution, any of them can be chosen). We have:

Lemma 2. *One can sample from a δ -optimal solution of the SDP given by Eq. (2) (with dimension n , m variables, size parameter R and upper bound r on optimal solution vector) given the ability to sample from a (δ/r) -optimal solution of the SDP given by Eq. (2) (with dimension $n + 1$, $m + 1$ variables and size parameter $2R + 1$) in which $b_i \geq 1$ for all $i \in [m]$.*

To apply Lemma 2 we need an upper bound r on the ℓ_1 -norm of the optimal solution vector. For example, if all b_i 's are positive, we can take $r = \alpha / \min_i b_i$.

We will also assume that $\alpha \geq 1$. We can ensure this is the case as follows: Given the SDP of Eq. (2) with all b_i 's larger than one, it is clear that we can take $\alpha > 0$ (as all the y_i 's are non-negative). Suppose $\alpha < 1$. Then we consider a new SDP with the b_i 's rescaled by $1/\alpha$. As an effect we must solve the scaled SDP to accuracy $\delta\alpha$. Note that the complexity of solving the SDP (which depends on the accuracy) increases when α approaches zero. It is an open question this drawback can be avoided.

The quantum algorithm for $\alpha \geq 1$, $b_i \geq 1$ discussed below is given in terms of multiplicative error, while the reduction above works with additive error. It is easy to convert the multiplicative approximation to additive approximation by redefining the error $\delta \rightarrow \delta/\alpha$. This results in an increased complexity of the algorithm in terms of α .

Probability of success: The probability of success of our algorithm is greater than

$$1 - O(\exp(-\log^{\zeta}(nm))), \quad (10)$$

where $\zeta > 0$ is a free parameter.

The Algorithm: Let

$$\gamma := \left\lceil \frac{8}{\varepsilon^2} \log(m) R^2 \right\rceil \quad (11)$$

for an $\varepsilon > 0$ defined below. For an integer $k \leq \gamma$, we define

$$\bar{h}(\rho, k) := \bar{h} \left(\rho, -\frac{\varepsilon}{8R^2} k, -\frac{\varepsilon}{8R^2} (\gamma - k) \right), \quad (12)$$

with \bar{h} the Hamiltonian given by Eq. (7). Let $e_1 := (1, 0, \dots, 0)$ be the first computational basis state of \mathbb{R}^m . Let

$$G_{\bar{h}}(\rho) := \max_{k \in [\gamma]} \left(\text{number of calls to the oracle } O[\bar{h}(\rho, k)] \text{ in GibbsSampler} \left(O[\bar{h}(\rho, k)], \frac{\varepsilon}{4} \right) \right). \quad (13)$$

⁶It follows from $\text{tr}(CX) \leq \|X\|_1 \|C\| \leq \text{tr}(X) \leq R$.

The main algorithm is the following:

Algorithm 3 (Quantum Algorithm for SDPs).

Input: Oracles for $\{A_1, \dots, A_m, C\}$, with $\|C\|, \|A_i\| \leq 1$ and $\{b_1, \dots, b_m\}$ with $b_i \geq 1$. Parameters $R, \alpha, \delta, \xi > 0$.

Output: Either a sample from distribution p and a real number number L such $y := Lp$ is dual feasible with objective value less than $(1 + \delta)\alpha$, or the label Larger indicating the optimal objective value is larger than $(1 - \delta)\alpha$.

Set $\rho^{(1)} = I/n$. Let $\varepsilon = \frac{\delta}{28R^2}$, $\varepsilon' = -\ln(1 - \varepsilon)$, $M = 80 \log^{1+\xi}(8R^2nm/\varepsilon)/\varepsilon^2$, $L = 80 \log^{1+\xi}(nm)/\varepsilon^2$ and $Q = 10^6 R^6 \ln^{2+\xi}(nm)/\delta^4$. Let $T = \frac{500R^3 \ln(n)}{\delta^2}$. For $t = 1, \dots, T$:

1. Set $y^{(t)} = (0, \dots, 0)$.
2. For $k = 1, \dots, \gamma = \lceil \frac{8}{\varepsilon^2} \log(m)R^2 \rceil$, $N = 1, \dots, \lceil \frac{\alpha}{\varepsilon} \rceil$,
 - Create M copies of $q \leftarrow \text{GibbsSampler}(O[\bar{h}(\rho^{(t)}, k)], \varepsilon/4)$.
 - Sample i_1, \dots, i_M independently from the distribution q .
 - Compute estimates $\{e_{i_1}, \dots, e_{i_M}, f\}$ of $\{\text{tr}(A_{i_1}\rho^{(t)}), \dots, \text{tr}(A_{i_M}\rho^{(t)}), \text{tr}(C\rho^{(t)})\}$ to accuracy $\varepsilon/2$ using $(M + 1)L$ samples from $\rho^{(t)}$.
 - If $1/M \sum_{j=1}^M e_{i_j} \geq f/(\varepsilon N) - \varepsilon$ and $1/M \sum_{j=1}^M b_{i_j} \leq \alpha/(\varepsilon N) + R\varepsilon$, set $k_t = k$, $N_t = N$, $q^{(t)} = q$ and $y^{(t)} = \varepsilon N q^{(t)}$.
3. If $y^{(t)} = (0, \dots, 0)$, stop and output Larger.
4. Create $Q + 1$ copies of $q^{(t)} \leftarrow \text{GibbsSampler}(O[\bar{h}(\rho^{(t)}, k_t)], \varepsilon/4)$
5. Sample i_1, \dots, i_Q independently from the distribution $q^{(t)}$.
6. Let $M^{(t)} = (\varepsilon N_t Q^{-1} \sum_{j=1}^Q A_{i_j} - C + 2\alpha I) / 4\alpha$.
7. Let $C_t := \frac{10 \log(m)}{\varepsilon^2} (\frac{\gamma \alpha}{\varepsilon} M + Q) G_{\bar{h}}(\rho^{(t)}) + \frac{2\gamma \alpha}{\varepsilon} ML$.
Create C_t copies of the state $\rho^{(t+1)} \leftarrow \text{GibbsSampler}(-\varepsilon'(\sum_{\tau=1}^t M^{(\tau)}), \varepsilon/4)$.

Output $\|\bar{y}\|_1$ and a sample from $\bar{y}/\|\bar{y}\|_1$ with $\bar{y} = \frac{\delta \alpha}{2R} e_1 + \frac{1}{T} \sum_{t=1}^T y^{(t)}$.

Let

$$G_M := \max_{t \leq T} \left(\text{number of calls to the oracle in } \text{GibbsSampler} \left(O \left[-\varepsilon' \left(\sum_{\tau=1}^t M^{(\tau)} \right) \right], \frac{\varepsilon}{4} \right) \right), \quad (14)$$

and

$$G_{\bar{h}} := \max_{t \leq T} G_{\bar{h}}(\rho^{(t)}). \quad (15)$$

Finally, let T_{Meas} be the maximum time needed to estimate one of $\text{tr}(A_i\rho)$, for $i \in [m]$, and $\text{tr}(C\rho)$ (for an arbitrary ρ) within additive error $\varepsilon/2$ (in Lemma 13 we show that for s -sparse matrices, $T_{\text{Meas}} \leq \tilde{O}(s/\varepsilon^2)$).

We prove in Section 5 the following:

Theorem 4. *Algorithm 3 runs in time*

$$\tilde{O}\left(\frac{R^{21}}{\delta^{11}}G_{\bar{h}}G_M\right) + \tilde{O}\left(\frac{R^{13}}{\delta^5}T_{\text{Meas}}\right). \quad (16)$$

The algorithm fails with probability at most

$$O((R/\delta)^{18}(nm)^{10}\exp(-\log^{\tilde{c}}(nm))). \quad (17)$$

Assuming it does not fail, if it outputs Larger, the optimal objective value is larger than $(1 - \delta)\alpha$. Otherwise it outputs a sample of a probability distribution p and a real number L such that $y = Lp$ is dual feasible and $\sum_i y_i b_i \leq (1 + \delta)\alpha$.

We note the algorithm is very costly in terms of the size parameter R and the error δ . We believe it is possible to significantly reduce the complexity in terms of these two parameters, but we leave this possibility as an open question to future work.

One interesting Gibbs sampler to consider is quantum Metropolis [11, 12]. Although it is difficult to obtain rigorous estimates on its running time, it is expected that it is polylogarithmic in many cases. Whenever this is the case for the Hamiltonians involved in the algorithm (given by linear combinations of the A_i 's and C), one would achieve exponential speed-ups.

In Section 5 we show:

Corollary 5. *Using the Gibbs Sampler from Ref. [13], Algorithm 3 runs in time $\tilde{O}(n^{\frac{1}{2}}m^{\frac{1}{2}}s^2R^{32}/\delta^{18})$.*

As we show in the next section, this represents an unconditional polynomial speed-up (in terms of m and n) over any classical method for solving semidefinite programming.

One particular case of interest is when R is a constant independent of all other parameters. In this case the running time only depends on the parameters n, m, s, δ . Moreover the SDP has a clear quantum interpretation: we want to optimize the expectation value of an observable C/R on a quantum state ρ subject to the constraints that the expectation value of A_i on ρ is bounded by b_i/R , for all $i \in [m]$.

2.3 Lower Bounds

We give a lower bound on the complexity of solving SDPs which shows the n, m dependence of Algorithm 3 cannot be substantially improved. Consider the following two instances of the primal problem given by Eq. (1), with $R = 1$, $A_1 = I$, $b_j = 1$ for $j \in [m]$ and either

1. For a random $i \in [n]$, set $C_{ii} = 1$. All other elements of C are set to zero. Choose at random $j \in [m]$ and set $(A_j)_{ii} = 2$. All other elements of the matrices $\{A_j\}_{i=2}^m$ are set to zero.
2. For a random $i \in [n]$, set $C_{ii} = 1$. All other elements of C are set to zero. All elements of the matrices $\{A_j\}_{i=2}^m$ are set to zero.

We claim that to decide which of the two cases we are given requires at least $\Omega(n + m)$ calls to the oracle classically and $\Omega(\sqrt{n} + \sqrt{m})$ calls quantum-mechanically. This follows from an elementary reduction to the search problem.

It is easy to see that the optimal solution of the primal and dual problems in the first case are $X = |i\rangle\langle i|/2$ and $y = |j\rangle\langle j|$, with objective value $1/2$. In the second case, in turn, $X = |i\rangle\langle i|$ and $y = |1\rangle\langle 1|$, with objective value 1 . Therefore we can decide which of the two we are given and find the marked (i, j) (in the first case) given samples of the optimal y and X and a constant-error approximation to $\text{tr}(X)$ or $\|y\|_1$. This is equivalent to solving two search problems, one in a list of n elements and another in a list of m elements.

2.4 Exponential Speed-ups for SDPs with Low-Rank Inputs

Let us consider the second model of getting the input data. Let us assume $\text{rank}(A_i) \leq r$ for all $i \in [n]$, for small r (e.g., $r = \text{polylog}(n)$). To keep with the earlier set-up of always having the constraint $\text{tr}(X) \leq R$, we allow one of the input matrices to be the identity (and therefore not be low rank).

In Section 5.1 we show that the Gibbs sampling of $\varepsilon'(\sum_{\tau=1}^t M^{(\tau)})$ can be done in time

$$\text{poly}(\log(n), \log(m), r) \quad (18)$$

(i.e., $G_M \leq \text{poly}(\log(n), \log(m), r)$). Then combining it with the same analysis behind Theorem 4 we find:

Theorem 6. *Consider Input Model 2 and SDPs with input matrices with rank less than r . Then Algorithm 3 runs in time $\text{poly}(\log(n), \log(m), r, R, \alpha, 1/\delta)m^{\frac{1}{2}}$.*

To compare with classical algorithms, it is also useful to consider the variant of Input Model 2 in which the classical input of the problem is the description of quantum circuits creating the eigenstates of the input matrices. We are not aware of any classical method that does not require exponentially more time ($O(\text{poly}(n))$) to solve the problem in this setting. Indeed it is easy to see that with the encoding we are considering the problem is as hard as any other in BQP.⁷

2.5 Discussion and Open Questions

The core quantum part of the algorithm is the preparation of quantum Gibbs states. Classically there are several interesting applications of the Monte Carlo method and the Metropolis algorithm to problems not related to simulating thermal properties of physical systems [24]. One could expect the same will be the case for quantum Metropolis. We might have to wait until there

⁷Consider the following BQP-complete problem: given a quantum circuit U , decide if $|\langle 0|U|0\rangle|^2 \geq 2/3$ or $|\langle 0|U|0\rangle|^2 \leq 1/3$. We can solve it by a SDP with low-rank matrices. Consider the SDP of Eq. (1) with $C = U|0\rangle\langle 0|U^\dagger$, $A_1 = I$, $R = 1$, $A_2 = -|0\rangle\langle 0|$ and $b_2 = -1$. It is equivalent to

$$\begin{aligned} \max_{\rho} \text{tr}(\rho U|0\rangle\langle 0|U^\dagger) \\ \text{tr}(|0\rangle\langle 0|\rho) &\geq 1, \\ \text{tr}(\rho) &\leq 1, \\ \rho &\geq 0, \end{aligned} \quad (19)$$

whose solution is readily seen to be $|\langle 0|U|0\rangle|^2$.

are working quantum computers to fully explore the usefulness of quantum Metropolis, since in analogy to the classical case, many times heuristic methods based on it might work well in practice even though it is hard to get theoretical guarantees. Nevertheless, as far as we know the results of this paper give the first example of a problem of interest outside the simulation of physical systems in which "quantum Monte Carlo" methods (i.e., sampling from quantum Gibbs states) play an important role. The algorithm can also be seen as a new application of quantum annealing. One difference is that in this case the annealing is used to prepare a finite temperature state, instead of a groundstate as is usually considered in quantum adiabatic optimization.

The algorithm is also inherently robust, in the sense that to compute a solution of the SDP to accuracy ε , it suffices to be able to prepare approximations to accuracy $O(\varepsilon)$ of the Gibbs state of Hamiltonians given by linear combinations of the input matrices. Moreover we believe the constants and the dependence on R and δ might be substantially improved by a more careful analysis. If this turns out to be indeed the case, we expect the algorithm to be a promising candidate for a relevant application of small quantum computers (even without the need for error correction).

This work leaves several open questions for future work. For example:

- The algorithm has very poor scaling in terms of R and δ . It is a pressing open question to improve its running time in terms of these parameters. Also can we close the gap in terms of n and m between the lower bound ($\Omega(\sqrt{n} + \sqrt{m})$) and the algorithm ($O(\sqrt{nm})$)?
- Although quadratic speed-ups in terms of n and m are the best possible in the worst case, it is an interesting question whether more significant speed-ups are possible in specific instances. We show that in a certain sense this is the case for SDPs with low-rank matrices. Are there other interesting examples? How large are the speed-ups on average (for example choosing the input matrices at random from a given distribution)? Even more interesting is to explore whether there is a SDP of practical interest for which we might have larger quantum speed-ups.
- Is there a practical setting for which the second model of input is relevant? More generally, can the quantum speed-ups we consider be applied to obtain faster solutions for problems in which SDPs are used as an algorithmic tool?
- How robust is the algorithm to noise? Can we run it without the need of quantum error correction in analogy to what has been proposed for quantum annealing? Is there an improvement of the algorithm which would be suitable for a small-scale quantum computer (with hundreds of physical qubits)?
- The multiplicative method is an important algorithmic technique classically. In this paper we give an application of the matrix multiplicative weight method to quantum algorithms. Are there more applications?
- Can we enlarge the class of optimization problems having a quantum speed-up beyond SDPs? In particular, can we get quantum speed-ups for optimizing general convex functions over convex sets (assuming we have an efficient oracle for membership in the set)?
- In practice the preferred algorithms for solving SDPs are based on the interior point method. Can we also find a quantum algorithm for SDPs based on it?

3 Analysis of the Quantum Algorithm

3.1 The Arora-Kale Algorithm

The quantum algorithm builds on a classical algorithm of Arora and Kale for solving SDPs, which we now review. One element of their approach is an auxiliary algorithm termed $\text{ORACLE}(\rho)$, which given a density matrix ρ , searches for a vector y from the polytope

$$\mathcal{D}_\alpha := \{y \in \mathbb{R}^m : y \geq 0, b \cdot y \leq \alpha\} \quad (20)$$

such that

$$\sum_{j=1}^m y_j \text{tr}(A_j \rho) \geq \text{tr}(C \rho), \quad (21)$$

or outputs *fail* if no such vector exists.

The running time of their algorithm also depends on the so-called width ω of the SDP, defined as

$$\omega := \max_{y \in \mathcal{D}_\alpha} \left\| \sum_j y_j A_j - C \right\|. \quad (22)$$

We note the bound:⁸

$$\begin{aligned} \omega &= \max_{y \in \mathcal{D}_\alpha} \left\| \sum_j y_j A_j - C \right\| \\ &\leq \max_{y \in \mathcal{D}_\alpha} \sum_j y_j + 1 \\ &\leq \max_{y \in \mathcal{D}_\alpha} \sum_j b_j y_j + 1 \\ &\leq \alpha + 1 \\ &\leq R + 1. \end{aligned} \quad (23)$$

The Arora-Kale algorithm is the following:

Algorithm 7 (Arora-Kale Algorithm for SDPs).

Set $\rho^{(1)} = I/n$. Let $\varepsilon = \frac{\delta \alpha}{2R^2}$, and let $\varepsilon' = \ln(1 - \varepsilon)$. Let $T \geq \frac{16R^4 \ln(n)}{\alpha^2 \delta^2}$. For $t = 1, \dots, T$:

1. Run $\text{ORACLE}(\rho^{(t)})$. If it fails, stop and output $\rho^{(t)}$.
2. Else, let $y^{(t)}$ be the vector generated by $\text{ORACLE}(\rho^{(t)})$.
3. Let $M^{(t)} = (\sum_{j=1}^m A_j y_j^{(t)} - C + \omega I) / 2\omega$
4. Compute $W^{(t+1)} = \exp(-\varepsilon' (\sum_{\tau=1}^t M^{(\tau)}))$.
5. Set $\rho^{(t+1)} = \frac{W^{(t+1)}}{\text{tr}(W^{(t+1)})}$ and continue.

⁸Assuming $b_i \geq 1$.

The central idea behind the algorithm is a variant of the multiplicative weight method for positive semidefinite matrices. Let us denote by $\lambda_n(X)$ the minimum eigenvalue of the $n \times n$ Hermitian matrix X . Arora and Kale proved the following:

Lemma 8. [Matrix Multiplicative Weights method; Theorem 10 of [9]] Let $M^{(t)}$ be such that $0 \leq M^{(t)} \leq I$ for every t . Fix $\varepsilon < \frac{1}{2}$, and let $\varepsilon' = -\ln(1 - \varepsilon)$. Define $W^{(t)} = \exp\left(-\varepsilon' \left(\sum_{\tau=1}^{t-1} M^{(\tau)}\right)\right)$ and the density matrices $\rho^{(t)} = \frac{W^{(t)}}{\text{tr}(W^{(t)})}$. Then

$$\sum_{t=1}^T \text{tr}(M^{(t)} \rho^{(t)}) \leq (1 + \varepsilon) \lambda_n \left(\sum_{t=1}^T M^{(t)} \right) + \frac{\ln(n)}{\varepsilon}. \quad (24)$$

Let $e_1 := (1, 0, \dots, 0)$. Then the main result of [9] is the following (as a warm up to the proof of correctness of the quantum algorithm we reproduce the argument of Arora and Kale below):

Theorem 9 (Theorem 1 of [9]). Suppose ORACLE never fails for $T = \frac{16R^4 \ln(n)}{\alpha^2 \delta^2}$ iterations. Then $\bar{y} = \frac{\delta\alpha}{R} e_1 + \frac{1}{T} \sum_{t=1}^T y^{(t)}$ is dual feasible with objective value at most $\alpha(1 + \delta)$.

Proof. By the definition of ORACLE and the fact that $A_1 = I$, $b_1 = R$, we have

$$\bar{y}.b = \delta\alpha + \frac{1}{T} \sum_{t=1}^T y^{(t)}.b \leq \alpha(1 + \delta). \quad (25)$$

So it remains to show that \bar{y} is dual feasible. Again by the properties of ORACLE, $\bar{y} \geq 0$. We now use Lemma 8 to show $\sum_{j=1}^m \bar{y}_j A_j \geq C$. Indeed

$$\lambda_n \left(\sum_{j=1}^m \bar{y}_j A_j - C \right) \quad (26)$$

$$= \lambda_n \left(\frac{1}{T} \sum_{t=1}^T \sum_{j=1}^m y_j^t A_j - C \right) + \frac{\delta\alpha}{R} \quad (27)$$

$$= 2\omega \lambda_n \left(\frac{1}{T} \sum_{t=1}^T \left(\sum_{j=1}^m y_j^t A_j - C + \omega I \right) / 2\omega \right) - \omega + \frac{\delta\alpha}{R} \quad (28)$$

$$\stackrel{(i)}{\geq} \frac{2\omega}{(1 + \varepsilon)} \frac{1}{T} \sum_{t=1}^T \text{tr} \left(\rho^{(t)} \left(\sum_{j=1}^m y_j^t A_j - C + \omega I \right) / 2\omega \right) - \frac{2\omega \ln(n)}{T(1 + \varepsilon)\varepsilon} - \omega + \frac{\delta\alpha}{R} \quad (29)$$

$$\stackrel{(ii)}{\geq} \frac{\omega}{(1 + \varepsilon)} - \frac{2\omega \ln(n)}{T(1 + \varepsilon)\varepsilon} - \omega + \frac{\delta\alpha}{R} \quad (30)$$

$$= -\frac{2\omega \ln(n)}{T(1 + \varepsilon)\varepsilon} - \frac{\varepsilon\omega}{(1 + \varepsilon)} + \frac{\delta\alpha}{R} \quad (31)$$

$$\geq -\frac{2\omega \ln(n)}{T(1 + \varepsilon)\varepsilon} + \frac{\delta\alpha}{2R} \quad (32)$$

$$\stackrel{(iii)}{\geq} 0. \quad (33)$$

Inequality (iii) follows from Eq. (23) and the choices of T and ε . Inequality (ii) follows from Eq. (21) in the definition of ORACLE which gives

$$\frac{1}{T} \sum_{t=1}^T \text{tr} \left(\rho^{(t)} \left(\sum_{j=1}^m y_j^t A_j - C + \omega I \right) / 2\omega \right) \geq \frac{1}{2}. \quad (34)$$

Finally inequality (i) follows from the Matrix Multiplicative Weight method (Lemma 8), which can be applied since by the definition of ω ,

$$0 \leq \left(\sum_{j=1}^m y_j^t A_j - C + \omega I \right) / 2\omega \leq I. \quad (35)$$

□

3.2 Approximately Implementing ORACLE by Gibbs Sampling

As a step towards the quantum algorithm, we now give an explicit implementation of the ORACLE auxiliary algorithm. The idea is to use the fact that by Jaynes' principle, we can w.l.o.g. take the output y of ORACLE(ρ) to be (up to normalization) a Gibbs probability distribution over the two constraints, i.e., a distribution over $[m]$ of the form

$$q_{\rho, \lambda, \nu}(i) := \frac{\exp(\lambda \text{tr}(A_i \rho) + \mu \sum_i b_i)}{\sum_i \exp(\lambda \text{tr}(A_i \rho) + \mu \sum_i b_i)}, \quad (36)$$

for real numbers λ, μ .

The following is a special case of Lemma 4.6 of [19] (obtained by taking the reference state to be maximally mixed) and is an approximate version of Jaynes' principle with a quantitative control of the parameters of the Hamiltonian in the Gibbs state (in contrast, in the original Jaynes' principle there is no control over the size of the interaction strengths of the Hamiltonian)

Let $\mathcal{M}(\mathbb{C}^n)$ and $\mathcal{D}(\mathbb{C}^n)$ be the set of Hermitian and density matrices over \mathbb{C}^n . Let $\mathcal{T} \subseteq \mathcal{M}(\mathbb{C}^n)$ be a compact set of matrices. We set

$$\Delta(\mathcal{T}) := \sup_{A \in \mathcal{T}} \|A\|. \quad (37)$$

For $A \in \mathcal{M}(\mathbb{C}^n)$, we define the associated dual norm

$$[A]_{\mathcal{T}} := \sup_{B \in \mathcal{T}} \text{tr}(BA). \quad (38)$$

Lemma 10. (Lemma 4.6 of [19]) *For every $\kappa > 0$, the following holds. Let $\mathcal{T} \subseteq \mathcal{M}(\mathbb{C}^m)$ be a compact set of matrices and let $\pi \in \mathcal{D}(\mathbb{C}^m)$ be a density matrix. If one defines $\gamma = \lceil \frac{8}{\kappa^2} \log(m) \Delta(\mathcal{T})^2 \rceil$ then there exist $X_1, \dots, X_\gamma \in \mathcal{T}$ such that*

$$\tilde{\pi} := \frac{\exp\left(-\frac{\kappa}{4\Delta(\mathcal{T})^2} \sum_{i=1}^{\gamma} X_i\right)}{\text{tr}\left(\exp\left(-\frac{\kappa}{4\Delta(\mathcal{T})^2} \sum_{i=1}^{\gamma} X_i\right)\right)} \quad (39)$$

satisfies

$$[\pi - \tilde{\pi}]_{\mathcal{T}} \leq \kappa. \quad (40)$$

The finitary version of Jaynes' principle above allows us to implement ORACLE assuming we have access to samples from the distributions $q_{\rho,\lambda,\mu}$. In fact, for the quantum algorithm it will be useful to prove a generalization in which we only assume we can sample from distributions $\bar{q}_{\rho,\lambda,\mu}$ close in variational distance to $q_{\rho,\lambda,\mu}$.

For an integer k , let

$$q_{\rho,k} := q_{\rho, -\frac{\kappa}{4R^2}k, -\frac{\kappa}{4R^2}(\gamma-k)}, \quad (41)$$

with

$$\gamma := \left\lceil \frac{8}{\kappa^2} \log(m) R^2 \right\rceil. \quad (42)$$

Algorithm 11 (Instantiation of ORACLE(ρ) by Sampling).

Input: Samples from distributions $\bar{q}_{\rho,k}$ such that $\|\bar{q}_{\rho,k} - q_{\rho,k}\|_1 \leq \nu$ and real numbers $\{e_i\}_{i=1}^{m+1}$ such that $|e_i - \text{tr}(A_i\rho)| \leq \nu$. A parameter $\kappa > 0$.

Output: Samples from distribution $y/\|y\|_1$ and value of $\|y\|_1$ satisfying Eqs. (43) and (44).

Let $M = 80 \log^{1+\xi}(8R^2nm/\varepsilon)/\varepsilon^2$. For $k = 1, \dots, \gamma$ and $N = 1, \dots, \lceil \frac{\alpha}{\kappa} \rceil$:

- Sample $i_1, \dots, i_M \in [m]^M$ independently from the distribution $\bar{q}_{\rho,k}$.
- If $\frac{1}{M} \sum_{j=1}^M e_{i_j} \geq \frac{e_{m+1}}{\kappa N} - (\kappa + \nu)$ and $\frac{1}{M} \sum_{j=1}^M b_{i_j} \leq \frac{\alpha}{\kappa N} + R(\kappa + \nu)$, output samples from $\bar{q}_{\rho,k}$ and the number κN (as $y/\|y\|_1$ and $\|y\|_1$, respectively).

Lemma 12. Suppose ORACLE(ρ) does not fail. Then with probability larger than $1 - \exp(-\log^\xi(nm))$, Algorithm 11 outputs y such that

$$\sum_{i=1}^m b_i y_i \leq \alpha(1 + 2R(\kappa + \nu)), \quad (43)$$

and

$$\sum_{i=1}^m y_i \text{tr}(A_i\rho) \geq \text{tr}(C\rho) - 2\alpha(\kappa + \nu). \quad (44)$$

Proof. By the Chernoff bound and the union bound over all all $k \in [\gamma]$, with probability at least $1 - \exp(-\log^\xi(nm))$, sampling i_1, \dots, i_M independently from $\bar{q}_{\rho,k}$ guarantees that

$$\left| \sum_{i=1}^m \bar{q}_{\rho,k}(i) \text{tr}(A_i\rho) - \frac{1}{M} \sum_{j=1}^M e_{i_j} \right| \leq \kappa + \nu, \quad (45)$$

and

$$\left| \sum_{i=1}^m \bar{q}_{\rho,k}(i) b_i - \frac{1}{M} \sum_{j=1}^M b_{i_j} \right| \leq \kappa + \nu, \quad (46)$$

for all $k \leq \gamma$.

Let $k \leq \gamma, N \leq \lceil \frac{R}{\kappa} \rceil$ be the smallest integers (assuming they exist) such that

$$\sum_{i=1}^m \bar{q}_{\rho,k}(i) \operatorname{tr}(A_i \rho) \geq \frac{\operatorname{tr}(C\rho)}{\kappa(N+1)} - (\kappa + \nu), \quad (47)$$

and

$$\sum_{i=1}^m \bar{q}_{\rho,k}(i) b_i \leq \frac{\alpha}{\kappa N} + R(\kappa + \nu). \quad (48)$$

Then by Eqs. (45) and (46), with probability larger than $1 - \exp(-\log^{\xi}(nm))$ over the choice of i_1, \dots, i_M ,

$$\frac{1}{M} \sum_{j=1}^M e_{i_j} \geq \frac{\operatorname{tr}(C\rho)}{\kappa(N+1)} - 2(\kappa + \nu), \quad (49)$$

and

$$\frac{1}{M} \sum_{j=1}^M b_{i_j} \leq \frac{\alpha}{\kappa N} + 2R(\kappa + \nu), \quad (50)$$

where we used $\max_i |b_i| = R$. The algorithm will then output $y = \kappa N \bar{q}_{\rho,k}$, which satisfies Eqs. (43) and (44).

It remains to prove the existence of at least one pair $k \leq \gamma, N \leq \lceil \frac{\alpha}{\kappa} \rceil$ satisfying Eqs. (47) and (48). Let y^* be an output of $\text{ORACLE}(\rho)$. Let us apply Lemma 10 with $\pi := \sum_i y_i^* |i\rangle \langle i| / \|y^*\|_1$ and $\mathcal{T} = \{X := \sum_i \operatorname{tr}(A_i \rho) |i\rangle \langle i|, Y := \sum_i b_i |i\rangle \langle i|\}$. Note that $\Delta(\mathcal{T}) = R$. We find there is an integer $k \leq \gamma$ such that

$$\tilde{\pi} := \frac{\exp\left(-\frac{\kappa}{4R^2}(kX + (\gamma - k)Y)\right)}{\operatorname{tr}\left(\exp\left(-\frac{\kappa}{4R^2}(kX + (\gamma - k)Y)\right)\right)} \quad (51)$$

satisfies

$$[\pi - \tilde{\pi}]_{\mathcal{T}} \leq \kappa. \quad (52)$$

Let N be the integer which minimizes $|\kappa N' - \|y^*\|_1|$ over $N' \in \lceil \lceil \frac{\alpha}{\kappa} \rceil \rceil$. By the bound $\|y^*\| \leq \sum_i b_i y_i^* \leq \alpha$, we have

$$|\kappa N - \|y^*\|_1| \leq \kappa. \quad (53)$$

Then

$$\begin{aligned} \sum_{j=1}^m \bar{q}_{\rho,k_{\text{opt}}}(j) \operatorname{tr}(A_j \rho) &\geq \sum_{j=1}^m \frac{y_j^*}{\|y^*\|_1} \operatorname{tr}(A_j \rho) - (\kappa + \nu) \\ &\geq \frac{1}{\|y^*\|_1} \operatorname{tr}(C\rho) - (\kappa + \nu) \\ &\geq \frac{1}{\kappa(N+1)} \operatorname{tr}(C\rho) - (\kappa + \nu), \end{aligned} \quad (54)$$

where the first inequality follows from Eq. (52) and the fact that $\bar{q}_{\rho,k}$ is ν -close to $q_{\rho,k}$, the second from Eq. (21), and the last from Eq. (53).

Likewise,

$$\begin{aligned}
\sum_{j=1}^m \bar{q}_{\rho, k_{\text{opt}}} b_i &\leq \sum_{j=1}^m \frac{y_i^*}{\|y^*\|_1} b_i + R(\kappa + \nu) \\
&\leq \frac{1}{\|y^*\|_1} \alpha + R(\kappa + \nu) \\
&\leq \frac{1}{\kappa(N-1)} \alpha + R(\kappa + \nu).
\end{aligned} \tag{55}$$

□

4 Implementing the Oracle for $\bar{h}(\rho, \lambda, \mu)$

In this section we explain how to implement the oracle that outputs the entries of the Hamiltonian $\bar{h}(\rho, \lambda, \mu)$ defined in Eq. (7). We start with the following standard result in quantum algorithms:

Lemma 13. *Given a s -sparse $n \times n$ Hermitian matrix A with $\|A\| \leq 1$ and a density matrix ρ , with probability larger than $1 - p_e$, one can compute $\text{tr}(\rho A)$ with additive error ε in time $O(s\varepsilon^{-2} \log^4(ns/(p_e\varepsilon)))$ using $O(\log(1/p_e)\varepsilon^{-2})$ copies of ρ .*

Proof. Using the Hamiltonian simulation technique of Ref. [25], the phase estimation algorithm takes time $(s \log^4(ns/\varepsilon))$ to measure to accuracy $\varepsilon/2$ the energy of A_i in the state ρ . By the Chernoff bound, repeating the process $O(1/\varepsilon^2)$ times allows us to obtain an estimation for $\text{tr}(\rho A_i)$ to accuracy ε . □

Lemma 14. *Using $O(\log(m/p_e)h_{\text{precision}}^{-2})$ copies of ρ and time $O(sh_{\text{precision}}^{-2} \log^4(nms/(p_e h_{\text{precision}})))$, one can implement $O[\bar{h}]$ with error probability p_e .*

Proof. Since by assumption we have an oracle for the $\{b_i\}$, we can focus on showing how to compute $\text{tr}(A_i\rho)$ to accuracy ν given access to an oracle for the entries of the A_i and copies of the state ρ . Lemma 13 shows how to compute, with probability at least $1 - p'_e$, an estimate of $\text{tr}(A_i\rho)$ to accuracy $h_{\text{precision}}$ in time $O(sh_{\text{precision}}^{-2} \log^4(ns/(p'_e h_{\text{precision}})))$ using $O(\log(1/p'_e)h_{\text{precision}}^{-2})$ copies of ρ . Suppose the input of the oracle is $\sum_{i=1}^m c_i |i\rangle$. Then its output is $\sum_{i=1}^m c_i |i, e_i\rangle$ with e_i the estimation of $\text{tr}(A_i\rho)$. We know each e_i is within $h_{\text{precision}}$ from the true value with probability at least $1 - p'_e$. Therefore by the union bound all of the e_i are $h_{\text{precision}}$ -close with probability $1 - mp'_e$. □

5 The Quantum Algorithm: Correctness and Complexity

We are ready to prove:

Theorem 15 (restatement of Theorem 4). *Algorithm 3 runs in time*

$$\tilde{O}\left(\frac{R^{21}}{\delta^{11}} G_h G_M\right) + \tilde{O}\left(\frac{R^{13}}{\delta^5} T_{\text{Meas}}\right). \tag{56}$$

The algorithm fails with probability at most

$$O((R/\delta)^{18} (nm)^{10} \exp(-\log^{\tilde{c}}(nm))). \tag{57}$$

Assuming it does not fail, if it outputs Larger, the optimal objective value is larger than $(1 - \delta)\alpha$. Otherwise it outputs a sample of a probability distribution p and a real number L such that $y = Lp$ is dual feasible and $\sum_i y_i b_i \leq (1 + \delta)\alpha$.

Proof. Correctness of Algorithm: We let $p_e = \exp(-\log^{\xi}(nm))$, with p_e the error probability for the oracle \bar{h} . If at some call to \bar{h} , the output is wrong, we declare the algorithm failed. Let us first assume that the oracle \bar{h} outputs the correct value in all calls. Later we will show the oracle to \bar{h} is used at most $O((R/\delta)^{10}(nm)^{10} \exp(-\log^{\xi}(nm)))$ times, the algorithm fails due to a faulty \bar{h} with probability at most $O((R/\delta)^{18}(nm)^{10} \exp(-\log^{\xi}(nm)))$.

Let us first consider the case in which for every $t \leq T$, after step 2, $y^{(t)}$ is not equal to $(0, \dots, 0)$. Then the algorithm has to output a sample from $y/\|y\|_1$ and the value of $\|y\|_1$, with

$$\bar{y} = \frac{\delta\alpha}{2R}e_1 + \frac{1}{T} \sum_{t=1}^T y^{(t)}. \quad (58)$$

Note $\|y^{(t)}\|_1 = \varepsilon N_t$, so we know all of them. Therefore we can compute $\|\bar{y}\|_1$. Since we have samples from $y^{(t)}$ for all $t \leq T$, we can sample from \bar{y} , which is a convex combination of them (and where we know the mixing probability distribution).

Lemma 14 with $h_{\text{precision}} = \varepsilon/2$ and Lemma 12 with $\nu = \kappa = \varepsilon/2$ show Step 2 of the algorithm implements $\text{ORACLE}(\rho^{(t)})$ as in Algorithm 11. Then for each t , Step 2 outputs a vector $y^{(t)}$ such that, with probability at least $1 - \exp(-\log^{\xi}(nm))$,

$$\sum_{i=1}^m b_i y_i^{(t)} \leq \alpha(1 + 2R\varepsilon), \quad (59)$$

and

$$\sum_{i=1}^m y_i^{(t)} \text{tr}(A_i \rho^{(t)}) \geq \text{tr}(C \rho^{(t)}) - 2\alpha\varepsilon. \quad (60)$$

The remaining steps of the algorithm run the Matrix Multiplicative Weight method. One difference (which will be important to keep the quantum complexity of the algorithm low), is the sparsification of the pay-off matrix $M^{(t)}$. Let

$$\hat{M}^{(t)} := \left(\sum_{i=1}^m y_i^{(t)} A_i - C + 2\alpha I \right) / 4\alpha. \quad (61)$$

Since

$$\left\| \sum_{i=1}^m y_i^{(t)} A_i - C \right\| \leq \sum_{i=1}^m y_i^{(t)} + 1 \leq \sum_{i=1}^m b_i y_i^{(t)} + 1 \leq \alpha + 1, \quad (62)$$

we have $0 \leq \hat{M}^{(t)} \leq I$.

By the Matrix Hoeffding bound (Lemma 16), with probability larger than $1 - \exp(-\log^{\xi}(nm))$,

$$\left\| \hat{M}^{(t)} - M^{(t)} \right\| \leq \frac{1}{4T}. \quad (63)$$

Then by Lemma 17 and the fact the Gibbs sampler has error $\varepsilon/2$, we find

$$\left\| \rho^{(t)} - \hat{\rho}^{(t)} \right\|_1 \leq \varepsilon, \quad (64)$$

with

$$\hat{\rho}^{(t)} := \frac{\exp(-\varepsilon'((\sum_{\tau=1}^t \hat{M}^{(\tau)}))}{\text{tr}(\exp(-\varepsilon'((\sum_{\tau=1}^t \hat{M}^{(\tau)})))}. \quad (65)$$

We are ready to show the correctness of the algorithm. Since $\eta = \frac{\delta\alpha}{2R}$, we find

$$\begin{aligned} \bar{y}.b &= R\eta + \frac{1}{T} \sum_{t=1}^T y^{(t)}.b \\ &\leq R\eta + \alpha(1 + 2R\varepsilon) \\ &= \alpha(1 + \delta/2 + 2R\varepsilon) \\ &\leq (1 + \delta)\alpha. \end{aligned} \quad (66)$$

Let us now show that \bar{y} is dual feasible. It is clear that $\bar{y} \geq 0$. In analogy with Eq. (26), we have

$$\lambda_n \left(\sum_{j=1}^m \bar{y}_j A_j - C \right) \quad (67)$$

$$= \lambda_n \left(\frac{1}{T} \sum_{t=1}^T \sum_{j=1}^m y_j^t A_j - C \right) + \eta \quad (68)$$

$$= 4\alpha\lambda_n \left(\frac{1}{T} \sum_{t=1}^T \left(\sum_{j=1}^m y_j^t A_j - C + 2\alpha I \right) / 4\alpha \right) - 2\alpha + \eta \quad (69)$$

$$\stackrel{(i)}{\geq} \frac{4\alpha}{(1+\varepsilon)} \frac{1}{T} \sum_{t=1}^T \text{tr} \left(\hat{\rho}^{(t)} \left(\sum_{j=1}^m y_j^t A_j - C + 2\alpha I \right) / 4\alpha \right) - \frac{4\alpha \ln(n)}{T(1+\varepsilon)\varepsilon} - 2\alpha + \eta \quad (70)$$

$$\stackrel{(ii)}{\geq} \frac{4\alpha}{(1+\varepsilon)} \frac{1}{T} \sum_{t=1}^T \text{tr} \left(\rho^{(t)} \left(\sum_{j=1}^m y_j^t A_j - C + 2\alpha I \right) / 4\alpha \right) - \frac{4\alpha\varepsilon}{(1+\varepsilon)} - \frac{4\alpha \ln(n)}{T(1+\varepsilon)\varepsilon} - 2\alpha + \eta \quad (71)$$

$$\stackrel{(iii)}{\geq} \frac{4\alpha}{(1+\varepsilon)} \left(\frac{1}{2} - 2\alpha\varepsilon \right) - \frac{4\alpha\varepsilon}{(1+\varepsilon)} - \frac{4\alpha \ln(n)}{T(1+\varepsilon)\varepsilon} - 2\alpha + \eta \quad (72)$$

$$= -\frac{6\alpha\varepsilon}{(1+\varepsilon)} - \frac{8\alpha^2\varepsilon}{(1+\varepsilon)} + \frac{\alpha\delta}{2R} - \frac{4\alpha \ln(n)}{T(1+\varepsilon)\varepsilon}$$

$$\geq \frac{\delta\alpha}{4R} - \frac{4\alpha \ln(n)}{T(1+\varepsilon)\varepsilon} \quad (73)$$

$$\geq 0. \quad (74)$$

where we used that $\alpha \geq 1$,

$$T = \frac{16R \ln(n)}{\delta\varepsilon}, \quad (75)$$

and

$$\varepsilon = \frac{\delta}{28R^2}. \quad (76)$$

Inequality (i) follows from the Matrix Multiplicative Weight method (Lemma 8), which can be applied since by Eq. (62),

$$0 \leq \left(\sum_{j=1}^m y_j^t A_j - C + 2\alpha I \right) / 4\alpha \leq I. \quad (77)$$

Inequality (ii) follows from Eq. (64), while Inequality (iii) follows from Eq. (60).

Let us now turn to the case in which there is a $t \in [T]$ such that the loop in Step 2 outputs $y^{(t)} = (0, \dots, 0)$. Then by the Chernoff bound and Lemma 10 we find that for every $y \geq 0$ s.t.

$$\sum_{i=1}^m y_i \operatorname{tr}(A_i \rho) \geq \operatorname{tr}(C \rho) - 3\alpha\epsilon, \quad (78)$$

we must have

$$y \cdot b \leq \alpha(1 - 3R\epsilon). \quad (79)$$

By duality of linear programming it follows the maximum over $\lambda \geq 0$ of

$$\lambda(\operatorname{tr}(C \rho^{(t)}) - 3\alpha\epsilon) \quad (80)$$

subject to the constraints

$$\lambda \operatorname{tr}(A_i \rho^{(t)}) \leq b_i, \quad i \in [m] \quad (81)$$

must be larger than $\alpha(1 - 3R\epsilon)$. Note that $\lambda = \lambda \operatorname{tr}(\rho^{(t)}) \leq R$. Then defining $X = \lambda_{\text{optimal}} \rho^{(t)}$ (with λ_{optimal} the optimal value of λ for the LP above), we find that $\operatorname{tr}(A_i X) \leq b_i$ for all $i \in [m]$ and

$$\operatorname{tr}(CX) \geq \alpha(1 - 3R\epsilon) \geq (1 - \delta)\alpha. \quad (82)$$

Finally, let us bound the probability that the algorithm fails. This can be due to three causes. The first is a faulty oracle call for \bar{h} . This is upper bounded by

$$O((R/\delta)^{18} (nm)^{10} \exp(-\log^{\bar{\epsilon}}(nm))). \quad (83)$$

The second is due to an error in estimating the values of $\{\operatorname{tr}(A_i \rho^{(t)})\}$ in Step 2 of the algorithm. The third is the random sampling used to construct $M^{(t)}$. By the union bound the error probability of both are also bounded by Eq. (83).

Run-time Analysis: We remind the reader we are assuming that $\alpha \geq 1$.

The cost of running step 2 is $\tilde{O}(\alpha R^2 / \epsilon^3) \leq \tilde{O}(R^9 / \delta^3)$ times the sum of cost of the following: (i) performing M times the Gibbs sampling of the Hamiltonian \bar{h} with cost $M G_{\bar{h}}$, (ii) sampling M times from the resulting distribution with cost $O(M)$, and (iii) computing estimates to the expectation values of the input matrices on the state $\rho^{(t)}$, with cost $(M + 1) T_{\text{Meas}}$. Therefore the total cost of step 2 (for a particular $t \in [T]$) is

$$\tilde{O}\left(\frac{R^9}{\delta^3} M (G_{\bar{h}} + T_{\text{Meas}})\right) = \tilde{O}\left(\frac{R^{13}}{\delta^5} (G_{\bar{h}} + T_{\text{Meas}})\right) \quad (84)$$

The cost of step 4 is $(Q + 1) G_{\bar{h}} = \tilde{O}((R^6 / \delta^4) G_{\bar{h}})$. The cost of step 5 is $Q = \tilde{O}(R^6 / \delta^4)$. The cost of step 6-7 is

$$C_t G_M = \tilde{O}\left(\frac{\alpha R^2}{\epsilon^3 \epsilon^2} \left(\frac{1}{\epsilon^2} + \frac{R^6}{\delta^4}\right) G_{\bar{h}}\right) G_M + \tilde{O}\left(\frac{2\alpha R^2}{\epsilon^3 \epsilon^4}\right) G_M = \left(\frac{R^{19}}{\delta^9}\right) G_{\bar{h}} G_M. \quad (85)$$

As each step is repeated $T = \tilde{O}(R^3 / \delta^2)$ times, the total cost is

$$\tilde{O}\left(\frac{R^{21}}{\delta^{11}} G_{\bar{h}} G_M\right) + \tilde{O}\left(\frac{R^{13}}{\delta^5} T_{\text{Meas}}\right). \quad (86)$$

□

Lemma 16. (Matrix Hoeffding Bound, Theorem 2.8 of [26]) Suppose Z_1, \dots, Z_k are independent random $d \times d$ Hermitian matrices satisfying $\mathbb{E}[Z_i] = 0$ and $\|Z_i\| \leq \lambda$. Then

$$\Pr \left[\left\| \frac{1}{k} \sum_{i=1}^k Z_i \right\| \geq \delta \right] \leq d.e^{-\frac{k\delta^2}{8\lambda^2}}. \quad (87)$$

Lemma 17. Let H, H' be Hermitian matrices. Then

$$\left\| \frac{e^H}{\text{tr}(e^H)} - \frac{e^{H'}}{\text{tr}(e^{H'})} \right\|_1 \leq 2 \left(e^{\|H-H'\|} - 1 \right). \quad (88)$$

Proof. We can assume w.l.o.g. that $\text{tr}(e^H) \geq \text{tr}(e^{H'})$.

Let M be an arbitrary operator with $\|M\| \leq 1$. We write

$$\text{tr} \left(M \frac{e^H}{\text{tr}(e^H)} \right) = \frac{\text{tr}(e^{H'})}{\text{tr}(e^H)} \text{tr} \left(M \frac{e^{H'}}{\text{tr}(e^{H'})} \left(\mathcal{T} \exp \left(\varepsilon \int_0^1 dt e^{-tH'} (H - H') e^{tH'} \right) \right) \right) \quad (89)$$

with \mathcal{T} the time-ordered operator, i.e.

$$\begin{aligned} & \mathcal{T} \exp \left(\int_0^1 dt e^{-tH'} (H - H') e^{tH'} \right) \\ & := I + \int_0^1 dt e^{-tH'} (H - H') e^{tH'} + \int_0^1 dt_1 \int_0^{t_1} dt_2 e^{-t_1H'} (H - H') e^{(t_1-t_2)H'} (H - H') e^{t_2H'} + \dots, \end{aligned} \quad (90)$$

where the times are such that $1 \geq t_1 \geq \dots \geq t_k$.

We now follow closely the argument in the Appendix of [21]. Write

$$\frac{\text{tr}(e^H)}{\text{tr}(e^{H'})} \text{tr} \left(M \frac{e^H}{\text{tr}(e^H)} \right) = \sum_{k=0}^{\infty} T_k, \quad (91)$$

with

$$\begin{aligned} & T_k \\ & := \text{tr} \left(M \frac{e^{H'}}{\text{tr}(e^{H'})} \left(\int_0^1 dt_1 \int_0^{t_1} dt_2 \dots \int_0^{t_k} dt_k e^{-t_1H'} (H - H') e^{(t_1-t_2)H'} (H - H') \dots (H - H') e^{t_kH'} \right) \right) \end{aligned} \quad (92)$$

Note

$$T_0 = \text{tr} \left(M \frac{e^{H'}}{\text{tr}(e^{H'})} \right). \quad (93)$$

Consider the k -th term in the series (for $k \geq 1$). We can bound it by

$$\begin{aligned} T_k & \leq \frac{1}{\text{tr}(e^{H'})} \int_0^1 dt_1 \int_0^{t_1} dt_2 \dots \int_0^{t_k} dt_k \\ & \quad \left\| e^{(1-t_1)H'} (H - H') e^{(t_1-t_2)H'} (H - H') \dots (H - H') e^{(t_{k-1}-t_k)H'} (H - H') e^{t_kH'} \right\|_1 \end{aligned} \quad (94)$$

We note there are k terms equal to $H - H'$ and k terms given by exponentials $\exp(\delta_i H')$, for positive δ_i . Hölder's inequality give $\|X_1 \dots X_l\|_1 \leq \prod_i \|X_i\|_{p_i}$, for the p_i norms of the matrices, with $\sum_i p_i^{-1} = 1$. We apply it to the expression above, taking the $p = \infty$ norm for the $(H - H')$ terms and the $1/\delta_i$ norm for the $e^{\delta_i H'}$ terms. Since $\sum_i \delta_i = 1$, we can apply the inequality. We find

$$T_k \leq \|H - H'\|^k \int_0^1 dt_1 \int_0^{t_1} dt_2 \dots \int_0^{t_k} dt_k \leq \frac{\|H - H'\|^k}{k!}. \quad (95)$$

Therefore

$$\left| \frac{\text{tr}(e^H)}{\text{tr}(e^{H'})} \text{tr} \left(M \frac{e^H}{\text{tr}(e^H)} \right) - \text{tr} \left(M \frac{e^{H'}}{\text{tr}(e^{H'})} \right) \right| \leq \sum_{k=1}^{\infty} |T_k| \leq e^{\|H - H'\|} - 1 \quad (96)$$

The result follows from the Golden-Thompson inequality, which implies

$$\text{tr}(e^H) \leq \text{tr}(e^{H'}) e^{\|H - H'\|}. \quad (97)$$

□

Finally let us prove

Corollary 18 (Restatement Corollary 5). *Using the Gibbs Sampler from Ref. [13], Algorithm 3 runs in time $\tilde{O}(n^{\frac{1}{2}} m^{\frac{1}{2}} s^2 R^{32} / \delta^{18})$.*

Proof. The result is a consequence of Theorem 4 and the main result of [13], which gives a method to prepare an ε -approximation to $e^{\beta H} / \text{tr}(e^{\beta H})$ for a s' -sparse H with $\|H\| \leq 1$ using

$$\tilde{O}(\sqrt{\dim(H)} \beta s' / \varepsilon). \quad (98)$$

calls to the oracle and two-qubit gates.⁹

We apply this Gibbs sampler to steps 2, 4 and 7 of the quantum algorithm. In order to estimate the running time of each of these applications, we must estimate the associated β and sparsity. In steps 2 and 4, the associated β is upper bounded by $O(1/\varepsilon) = O(R^2/\delta)$ and the sparsity is one. In step 7 the associated β is upper bounded by $\varepsilon' T = \tilde{O}(R/\delta)$, while the sparsity is upper bounded by TQ times the sparsity of each of the input matrices, which gives $\tilde{O}(sR^9/\delta^6)$. Finally, since to query an element the Hamiltonian we need to query each of the A'_i 's matrices appearing in the decomposition, we have a cost of Ts for querying an element of the Hamiltonian. Therefore we have the bounds

$$G_{\bar{h}} \leq \tilde{O}(\sqrt{m} R^2 / \delta) \quad (99)$$

and

$$G_M \leq \tilde{O}(\sqrt{ns^2} R^9 / \delta^6). \quad (100)$$

□

⁹The result of Ref. [13] is presented in the special case of a local Hamiltonian. However one can check that the only property of the Hamiltonian needed is that $U(t) = e^{-itH}$ can be implemented to error ε by a circuit of size $t \text{poly}(n, 1/\varepsilon)$. Since by [25] s -sparse Hamiltonians can be implemented by a circuit of size $st \text{poly}(n, \log(1/\varepsilon))$, the result can be applied to them.

5.1 Exponential Speed-ups for Low-Rank SDPs

We start with the following result:

Lemma 19. *In Input Model 2, there is a Gibbs Sampler for $\varepsilon'(\sum_{\tau=1}^t M^{(\tau)})$ with error δ running in time $\text{poly}(\log(n), \log(m), r, 1/\delta)$, with $r = \max(\max_i \text{rank}(A_i), \text{rank}(C))$.*

Proof. We are interested in constructing the density matrix $\exp(\varepsilon'(\sum_{\tau=1}^t M^{(\tau)})) / \text{tr}(\exp(\varepsilon'(\sum_{\tau=1}^t M^{(\tau)})))$. Note that the state is unaltered if we add a multiple of the identity to the matrix in the exponential. Thus it suffices to show how to construct the Gibbs state of the matrix $K := \varepsilon' \sum_{\tau=1}^t N^{(\tau)}$, with

$$N^{(\tau)} = \left(\varepsilon N Q^{-1} \sum_{j=1}^Q A_{ij} - C \right) / 4\alpha. \quad (101)$$

Let

$$K = \sum_{i=1}^{r_K} \lambda_i |\psi_i\rangle \langle \psi_i| \quad (102)$$

be its spectral decomposition. We have the bound $r_K \leq r \text{poly}(\log(n), \log(m))$. We will show how to prepare on a quantum computer in time $\text{poly}(r, \log(n), \log(m), 1/\delta)$ an $\delta/2$ -approximation of each of the eigenstates $|\psi_i\rangle$ (in trace norm) and how to compute each of the eigenvalues λ_i , with accuracy $1 - \delta/r_K$. Then we can construct the Gibbs state

$$\frac{\exp(K)}{\text{tr}(\exp(K))} = \frac{1}{\sum_i e^{\lambda_i}} \sum_{i=1}^{r_K} e^{\lambda_i} |\psi_i\rangle \langle \psi_i|, \quad (103)$$

by mixing the states $\{|\psi_i\rangle\}$ with probabilities $\{e^{\lambda_i} / \sum_i e^{\lambda_i}\}$.

The Hermitian matrix K is given by a linear combination of some of the A_i 's and C , i.e.,

$$K = \sum_{i=0}^{N_K} \mu_i A_{k_i}, \quad (104)$$

with $N_K \leq \text{polylog}(n, m)$, for a subset $\{A_{k_1}, \dots, A_{k_{N_K}}\}$ of $\{A_1, \dots, A_m\}$ and with $A_{k_0} = C$.

We have access to copies of the eigenstates $|\eta_l^{k_i}\rangle$ of A_{k_i} and the associated eigenvalues $\kappa_l^{k_i}$ with accuracy $\nu/2$. Then using the swap test [28] we can compute all overlaps $\langle \eta_l^{k_i} | \eta_{l'}^{k_{i'}} \rangle$ with error ν in time $\text{poly}(r, \log(n), \log(m), 1/\nu)$.

By Eq. (104), we can write

$$|\psi_j\rangle = \sum_{i,l} c_{i,l}^j |\eta_l^{k_i}\rangle \quad (105)$$

for complex numbers $c_{i,l}^j$. So all the eigenstates of K are in the span of the $|\eta_l^{k_i}\rangle$. But since we know all the overlaps $\langle \eta_l^{k_i} | \eta_{l'}^{k_{i'}} \rangle$, it means we can compute the coefficients $c_{i,l}^j$ and the eigenvalues λ_i (to precision δ) in time $\text{poly}(r, \log(n), \log(m), 1/\delta)$, by diagonalizing the matrix on a classical computer once we write it in the basis given by the $\{|\eta_l^{k_i}\rangle\}$.

Finally we can prepare the quantum states $|\psi_j\rangle$ as follows. First we prepare the state (up to normalization) $\sum_{i,l} c_{i,l}^j |i, l\rangle$. Then using the ability to prepare copies of $|\eta_l^{k_i}\rangle$, we use it to prepare

$\sum_{i,l} c_{i,l}^j |i, l, \eta_l^{k_i}\rangle$. Then we measure the first two registers (each of dimension $\text{poly}(r, \log(n), \log(m))$) in the Fourier basis and accept if the outcome obtained projects the state onto the uniform superposition over basis states. Conditioned on accepting (which happens with probability larger than $1 - 1/\text{poly}(r, \log(n), \log(m))$), we obtain $|\psi_j\rangle$. We then repeat $\text{poly}(r, \log(n), \log(m))$ times until succeeding. \square

We are ready to prove:

Theorem 20. (Restatement Theorem 6) *In Input Model 2, for SDPs with rank r input matrices, Algorithm 7 runs in time $\text{poly}(\log(n), r, R, \alpha, 1/\delta)m^{\frac{1}{2}}$.*

Proof. We follow Algorithm 3, using Lemma 19 to implement the Gibbs Sampler for $\varepsilon'(\sum_{\tau=1}^t M^{(\tau)})$. All the steps remain unchanged, apart from the estimation of the expectation values $\{\text{tr}(A_i \rho^{(t)})\}$ in Step 2. Since we do not have oracles for the entries of A_i , we cannot use Lemma 13. However, in [22] (see also [27]) it was shown how given copies of a n -dimensional state π one can implement a κ -approximation to $e^{-it\pi}$, in operator norm, in time $\text{poly}(t, 1/\kappa, \log(n))$. Therefore given that we have copies of $A_i^{+/-} / \text{tr}(A_i^{+/-})$, know the values of $\text{tr}(A_i^{+/-})$ (which are all smaller than r since $\|A_i^{+/-}\| \leq 1$ and their rank is less than r), we can implement e^{-itA_i} with error κ in time $\text{poly}(t, 1/\kappa, \log(n))$, which by Trotter decomposition allows us to implement e^{-itA_i} with error κ in time $\text{poly}(t, 1/\kappa, \log(n))$. Then we can proceed as in Lemma 13, using phase estimation repeatedly to estimate $\text{tr}(A_i \rho)$. \square

Acknowledgments

We thank Joran van Apeldoorn, Ronald de Wolf, Andras Gilyen, Aram Harrow, Sander Gribling, Matt Hastings, Cedric Yen-Yu Lin, Ojas Parekh, and David Poulin for interesting discussions and useful comments on the paper.

A Reduction to $b_i \geq 1$

Here we prove:

Lemma 21 (Restatement Lemma 2). *One can sample from a δ -optimal solution of the SDP given by Eq. (2) (with dimension n , m variables, size parameter R and upper bound on optimal solution vector r) given the ability to sample from a δ/r -optimal solution of the SDP given by Eq. (2) (with dimension $n + 1$, $m + 1$ variables and size parameter $2R + 1$) in which $b_i \geq 1$ for all $i \in [m]$.*

Proof. Given the SDP of Eq. (2), we define the following related SDP with $m + 1$ variables in $n + 1$ dimensions:

$$\begin{aligned} & \min \sum_{i=1}^m b_i y_i + (R + 1) \sum_{i=1}^{m+1} y_i \\ & \sum_{i=1}^m y_i \begin{bmatrix} A_i & 0 \\ 0 & 1 \end{bmatrix} + y_{m+1} \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \geq \begin{bmatrix} C & 0 \\ 0 & r \end{bmatrix} \\ & y \geq 0, \end{aligned} \tag{106}$$

for r an upper bound on $\sum_i z_i$, for an optimal solution (z_1, \dots, z_m) of the SDP given by Eq. (2).

Note that since $\max |b_i| = R$, the vector defining the objective function of the SDP of Eq. (106) has all elements larger or equal than one.

Let (y_1, \dots, y_{m+1}) be a δ -optimal solution of the SDP given by Eq. (106). We claim (y_1, \dots, y_m) is an δ -optimal solution of the original SDP given by Eq. (2). Indeed we have that

$$\sum_{i=1}^m y_i A_i \geq C \quad (107)$$

so (y_1, \dots, y_m) is feasible. It remains to show that

$$\sum_{i=1}^m b_i y_i \leq \text{opt} + \delta, \quad (108)$$

with opt the optimal value of the SDP of Eq. (2). Suppose it was not the case and that

$$\sum_{i=1}^m b_i y_i > \text{opt} + \delta. \quad (109)$$

Let us find a contradiction.

Let (z_1, \dots, z_m) be the optimal solution to the SDP of Eq. (2) defined above and consider the following solution to the SDP given by Eq. (106):

$$(y'_1, \dots, y'_{m+1}) := \left(z_1, \dots, z_m, \sum_{i=1}^{m+1} y_i - \sum_{i=1}^m z_i \right). \quad (110)$$

Note that since

$$\sum_{i=1}^{m+1} y_i \geq r \geq \sum_{i=1}^m z_i, \quad (111)$$

it follows

$$\sum_{i=1}^{m+1} y_i - \sum_{i=1}^m z_i \geq 0, \quad (112)$$

so all the elements are non-negative. Moreover, the matrix inequality constraint is satisfied since

$$\sum_{i=1}^m y'_i A_i = \sum_{i=1}^m z_i A_i \geq C \quad (113)$$

and

$$\sum_{i=1}^{m+1} y'_i = \sum_{i=1}^{m+1} y_i \geq r. \quad (114)$$

Finally since

$$\sum_{i=1}^m b_i y'_i = \sum_{i=1}^m b_i z_i, \quad (115)$$

we find

$$\left(\sum_{i=1}^m b_i y'_i + (R+1) \sum_{i=1}^{m+1} y'_i \right) - \left(\sum_{i=1}^m b_i y_i + (R+1) \sum_{i=1}^{m+1} y_i \right) = \text{opt} - \sum_{i=1}^m b_i y_i < -\delta, \quad (116)$$

which contradicts the assumption that (y_1, \dots, y_{m+1}) is δ -optimal.

Finally note that C has norm $\|C\| = \max(1, r)$ which might be larger than one. Therefore we solve an associated SDP with a rescaled C by $1/r$. To solve the original SDP with additive error δ , we must solve the rescaled SDP with additive error δ/r .

□

References

- [1] P.W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review* 41.2, 303 (1999).
- [2] L.K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.* 79, 325 (1997).
- [3] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review* 38, 49 (1996).
- [4] M.X. Goemans. Semidefinite programming in combinatorial optimization. *Mathematical Programming* 79, 143 (1997).
- [5] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [6] M.X. Goemans and D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM* 42, 1115 (1995).
- [7] Y.T. Lee, A. Sidford, and S.C. Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. *IEEE 56th Annual Symposium on the Foundations of Computer Science (FOCS)*, 2015.
- [8] S. Arora, E. Hazan and S. Kale. Fast algorithms for approximate semidefinite programming using the multiplicative weights update method. *46th Annual IEEE Symposium on Foundations of Computer Science*, 2005. *FOCS* 2005.
- [9] S. Arora and S. Kale. A combinatorial, primal-dual approach to semidefinite programs. *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. ACM, 2007.
- [10] S. Arora, E. Hazan and S. Kale. The Multiplicative Weights Update Method: a Meta-Algorithm and Applications. *Theory of Computing* 8, 121 (2012).
- [11] K. Temme et al. Quantum metropolis sampling. *Nature* 471, 87 (2011).
- [12] M.H. Yung and A. Aspuru-Guzik. A quantum–quantum Metropolis algorithm. *Proceedings of the National Academy of Sciences* 109, 754 (2012).
- [13] D. Poulin and P. Wocjan. Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer. *Phys. Rev. Lett.* 103, 220502 (2009).
- [14] A.N. Chowdhury and R.D. Somma. Quantum algorithms for Gibbs sampling and hitting-time estimation. *arXiv preprint arXiv:1603.02940* (2016).

- [15] M. Kastoryano and F.G.S.L. Brandao. Quantum Gibbs Samplers: the commuting case. *Comm. Math. Phys.* 344, 915 (2016).
- [16] F.G.S.L. Brandao and M. Kastoryano. Finite correlation length implies efficient preparation of quantum thermal states. In preparation.
- [17] G. Brassard et al. Quantum amplitude amplification and estimation. *Contemporary Mathematics* 305, 53 (2002).
- [18] E.T. Jaynes. *Information Theory and Statistical Mechanics II*. *Phys. Rev.* 108, 171 (1957).
- [19] J.R. Lee, P. Raghavendra, and D. Steurer. Lower bounds on the size of semidefinite programming relaxations. *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*. ACM, 2015.
- [20] R. Ahlswede, A. Winter. Strong Converse for Identification via Quantum Channels". *IEEE Trans. Information Theory* 48, 569 (2003).
- [21] N.E. Sherman, T. Devakul, M.B. Hastings, R.R.P. Singh. *Phys. Rev. E* 93, 022128 (2016).
- [22] S. Lloyd, M. Mohseni, P. Rebentrost. Quantum Principal Component Analysis. *Nature Physics* 10, 631 (2014).
- [23] A. Childs and R. Kothari. Limitations on the simulation of non-sparse Hamiltonians. arXiv preprint arXiv:0908.4398 (2009).
- [24] W.R. Gilks. *Markov chain monte carlo*. John Wiley and Sons, Ltd. Chicago (2005).
- [25] D.W. Berry, A.M. Childs, R. Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science (FOCS 2015)*, 792 (2015).
- [26] J. A. Tropp. User-friendly tail bounds for sums of random matrices, 2010, arXiv:1004.4389.
- [27] S. Kimmel, C. Yen-Yu Lin, G. Hao Low, M. Ozols, T.J. Yoder. Hamiltonian Simulation with Optimal Sample Complexity. arXiv:1608.00281.
- [28] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.* 87(16):167902, 2001. quant-ph/0102001.