

A QUANTUM ALGORITHM FOR SIMULATING NON-SPARSE HAMILTONIANS

CHUNHAO WANG^{1,2} ^a LEONARD WOSSNIG³ ^b

¹ *Institute for Quantum Computing and School of Computer Science, University of Waterloo*

² *Department of Computer Science, University of Texas at Austin*

³ *Department of Computer Science, University College London*

Received November 25, 2019

Revised June 2, 2020

We present a quantum algorithm for simulating the dynamics of Hamiltonians that are not necessarily sparse. Our algorithm is based on the input model where the entries of the Hamiltonian are stored in a data structure in a quantum random access memory (qRAM) which allows for the efficient preparation of states that encode the rows of the Hamiltonian. We use a linear combination of quantum walks to achieve poly-logarithmic dependence on precision. The time complexity of our algorithm, measured in terms of the circuit depth, is $O(t\sqrt{N}\|H\| \text{polylog}(N, t\|H\|, 1/\epsilon))$, where t is the evolution time, N is the dimension of the system, and ϵ is the error in the final state, which we call precision. Our algorithm can be directly applied as a subroutine for unitary implementation and quantum linear systems solvers, achieving $\tilde{O}(\sqrt{N})$ dependence for both applications.

Keywords: Hamiltonian simulation, quantum algorithm, linear combination of unitaries, quantum walk, quantum machine learning, qRAM, linear systems solver

Communicated by: R Cleve & J Eisert

1 Introduction

1.1 Background and main results

Hamiltonian simulation is the problem of simulating the dynamics of quantum systems, which is the original motivation for quantum computers [1, 2]. It has been shown to be BQP-hard and is hence conjectured not to be classically solvable in polynomial time, since such an algorithm would imply an efficient classical algorithm for any problem with an efficient quantum algorithm, including integer factorization [3].

Different input models have been considered in previous quantum algorithms for simulating Hamiltonian evolution. The *local Hamiltonian* model is specified by the local terms of a Hamiltonian. The *sparse-access* model for a sparse Hamiltonian H is specified by the following two oracles:

$$O_S|i, j\rangle|z\rangle \mapsto |i, j\rangle|z \oplus S_{i,j}\rangle, \text{ and} \tag{1}$$

$$O_H|i, j\rangle|z\rangle \mapsto |i, j\rangle|z \oplus H_{i,j}\rangle, \tag{2}$$

^achunhao.wang@uwaterloo.ca

^bleonard.wossnig@cs.ucl.ac.uk

where $S_{i,j}$ is the j -th nonzero entry of the i -th row and \oplus denotes the bit-wise XOR. The *linear combination of unitaries (LCU)* model is specified by a decomposition of a Hamiltonian as a linear combination of unitaries and we are given the coefficients and access to each unitary. Following the first proposal by Lloyd [4] for local Hamiltonians, Aharonov and Ta-Shma gave an efficient algorithm for sparse Hamiltonians [5]. Subsequently, many algorithms have been proposed which improved the runtime [6, 7, 8, 9, 10, 11, 12, 13, 14, 15], mostly in the sparse-access model, and have recently culminated in a few works with optimal dependence on all (or nearly all) parameters for sparse Hamiltonians [16, 17, 18].

While the above-mentioned input models arise naturally in many physics applications and matrix arithmetic applications (i.e., we have access to the local terms of a Hamiltonian or each entry of a Hamiltonian can be efficiently computed), in many machine learning applications, it is more convenient to work with a different input model, namely, the *quantum random access memory (qRAM)* model, where we assume that the entries of a Hamiltonian are stored in a data structure as in [19] and we have quantum access to the memory. As we receive and process the input data, with little overhead, they can be stored in the data structure, and then the qRAM will facilitate preparing quantum states corresponding to the input data. The use of the qRAM model has been successfully demonstrated in many applications such as quantum principal component analysis [20], quantum support vector machines [21], and quantum recommendation systems [19].

In this work, we consider the qRAM model for obtaining information about the Hamiltonian that is not necessarily sparse. Quantum access to this data structure allows us to efficiently prepare states that encode the rows of the Hamiltonian. Using the ability to prepare these states in combination with a quantum walk [16], we give the *first* (to the best of our knowledge) Hamiltonian simulation algorithm in the qRAM model whose time complexity has $\tilde{O}(\sqrt{N})$ dependence^c for non-sparse Hamiltonians of dimensionality N . As a subroutine of quantum linear systems solver given by [22], our result directly implies a quantum linear systems solver in the qRAM model with square-root dependence on dimension and poly-logarithmic dependence on precision, which exponentially improves the precision dependence of the quantum linear systems solver by [23] with the same input model. Since solving linear systems is a fundamental procedure of many machine learning tasks, our algorithm has extensive potential applications in *quantum machine learning*.

In [19], a quantum algorithm for recommendation systems was introduced based on an explicit description of a data structure, which resulted in a fast quantum algorithm for estimating singular values for any real matrix. This fast singular value estimation algorithm was used in a quantum algorithm for solving dense linear systems [23]. The data structure in [19, 23] allows us to prepare states that correspond to row vectors. In our algorithm for the Hamiltonian simulation problem, the main hurdle is to efficiently prepare the states which allow for a quantum walk corresponding to $e^{-iH/\|H\|_1}$. These states are quite different from those in [19, 7, 16]: the states required by [16] allow for a quantum walk corresponding to $e^{-iH/(d\|H\|_{\max})}$, where d is the row-sparsity of H , and their states can be prepared with $O(1)$ queries to the sparse-access oracle. However, the states required by our algorithm (See Eq. (7)) is less structural and it is not known how to prepare them with poly-logarithmic cost in the sparse-access model. In the qRAM model, we assume the entries of H are stored

^cIn this paper, we use $\tilde{O}(\cdot)$ to hide poly-logarithmic factors.

in a data structures, which permits state preparation with time complexity (circuit depth) $O(\text{polylog}(N))$. The precise definition of the data structure is presented in Definition 1.

Using the efficient state preparation procedure, we implement the linear combination of quantum walks in order to simulate the time evolution for non-sparse Hamiltonians with only poly-logarithmic dependence on precision. The main result of this work is summarized in the following theorem, which we prove in Sec. 3.

Theorem 1 (Non-sparse Hamiltonian Simulation). *Let $H \in \mathbb{C}^{N \times N}$ (with $N = 2^n$) be a Hermitian matrix stored in the data structure as specified in Definition 1. There exists a quantum algorithm for simulating the evolution of H for time t and error ϵ with time complexity (circuit depth)*

$$O\left(t \|H\|_1 n^2 \log^{5/2}(t \|H\|_1 / \epsilon) \frac{\log(t \|H\| / \epsilon)}{\log \log(t \|H\| / \epsilon)}\right). \tag{3}$$

In this work, we use the notation $\|\cdot\|_1$ to denote the induced 1-norm (i.e., maximum absolute row-sum norm), defined as $\|H\|_1 = \max_j \sum_{k=0}^{N-1} |H_{jk}|$; we use $\|\cdot\|$ to denote the spectral norm, and use $\|\cdot\|_{\max}$ to denote the max norm, defined as $\|H\|_{\max} = \max_{i,j} |H_{jk}|$.

By the fact that $\|H\|_1 \leq \sqrt{N} \|H\|$ (see [24], alternatively, a more generalized version of this relation is shown in Appendix A), we immediately have the following corollary.

Corollary 1. *Let $H \in \mathbb{C}^{N \times N}$ (where $N = 2^n$) be a Hermitian matrix stored in the data structure as specified in Definition. 1. There exists a quantum algorithm for simulating the evolution of H for time t and error ϵ with time complexity (circuit depth)*

$$O\left(t \sqrt{N} \|H\| n^2 \log^{5/2}(t \sqrt{N} \|H\| / \epsilon) \frac{\log(t \|H\| / \epsilon)}{\log \log(t \|H\| / \epsilon)}\right). \tag{4}$$

Remarks:

1. In Theorem 1, the circuit depth scales as $\tilde{O}(\sqrt{N})$. However, the gate complexity could be as large as $O(N^{2.5} \log^2(N))$ in general because of the addressing scheme which allows for *quantum access to classical data* stored in the data structure as specified in Definition 1. If some structure of H is promised (e.g., the entries of H repeat in some pattern), the addressing scheme could be implemented efficiently.
2. If H is d -sparse (i.e., H has at most d non-zero entries in each row), the time complexity (circuit depth) of our algorithm is

$$O\left(t \sqrt{d} \|H\| n^2 \log^{5/2}(t \sqrt{d} \|H\| / \epsilon) \frac{\log(t \|H\| / \epsilon)}{\log \log(t \|H\| / \epsilon)}\right). \tag{5}$$

This follows from Theorem 1 and the fact that $\|H\|_1 \leq \sqrt{d} \|H\|$ (as shown in Appendix A).

3. The techniques of our algorithm also work for the sparse-access model where we use standard state preparation techniques to prepare the state as in Eq. (7). Compared to the qRAM model, the time complexity of the state preparation for Eq. (7) in the sparse-access model incurs an additional $O(d)$ factor (for computing σ_j). Now, to simulate H for time t in the sparse-access model, the dependence of the time complexity on t , d , and $\|H\|$ becomes $O(td^{1.5} \|H\|)$ as $\|H\|_1 \leq \sqrt{d} \|H\|$. Hence our techniques have no advantage over previous results in the sparse-access model.

1.2 Related work

Hamiltonian simulation with $\|H\|$ dependence. For non-sparse Hamiltonians, a suitable model is the black-box model (a special case of the sparse-access model): querying the oracle with an index-pair $|i, j\rangle$ returns the corresponding entry of the Hamiltonian H , i.e., O_H defined in Eq. (2). With access to a black-box Hamiltonian, simulation with error ϵ can then be provably performed with query complexity $O((\|H\|t)^{3/2}N^{3/4}/\sqrt{\epsilon})$ for dense Hamiltonians [7], and it was also empirically observed in [7] that for several classes of Hamiltonians, $O(\sqrt{N}\log(N))$ queries suffice. Whether this $\tilde{O}(\sqrt{N})$ dependence holds for all Hamiltonians was left as an open problem. After the first version of this work was made public, this open problem was almost resolved by Low [18], where he proposed a quantum algorithm for simulating black-box Hamiltonians with time complexity $O((t\sqrt{N}\|H\|)^{1+o(1)}/\epsilon^{o(1)})$. Although our input model is stronger than the black-box model, our work distinguishes itself since the two models are still comparable in many quantum machine learning applications and our work gives a better complexity.

Hamiltonian simulation with $\|H\|_{\max}$ dependence. Since the qRAM model is stronger than the black-box model and the sparse-access model, previous quantum algorithms such as [7, 16] can be directly used to simulate Hamiltonians in the qRAM model. For a d -sparse Hamiltonian, the circuit depth of the black-box Hamiltonian simulation is given by $\tilde{O}(td\|H\|_{\max})$ as shown in [7, 16]. When H is non-sparse, their results imply the scaling $\tilde{O}(tN\|H\|_{\max})$. In applications where $\|H\|_{\max}$ is a measure of cost, our result has no advantage against theirs, as the inequality $\|H\| \leq \sqrt{N}\|H\|_{\max}$ implies $\tilde{O}(t\sqrt{H}\|H\|) = \tilde{O}(tN\|H\|_{\max})$. However, in the case where $\|H\|$ is a measure of cost, such as solving linear systems [25, 22], and black-box unitary implementation [7], our result has a quadratic improvement in the dimensionality dependence, as the inequality $\|H\|_{\max} \leq \|H\|$ implies $\tilde{O}(tN\|H\|_{\max}) = \tilde{O}(tN\|H\|)$.

Hamiltonian simulation in the qRAM model. Shortly after the first version of this paper was made public, Chakraborty, Gilyén, and Jeffery [26] independently proposed a quantum algorithm for simulating non-sparse Hamiltonians based on the similar qRAM input model which achieved the same time complexity as our method. Their work is based on a very general input model, namely, the block-encoding model, which was originally proposed in [13]. The block-encoding model assumes we are given a unitary (H/α) that contains H/α in its upper-left block. Then the evolution e^{-iHt} can be simulated in $\tilde{O}(\alpha\|H\|t)$ time. It was shown in [13] that working with the sparse-access model of a d -sparse Hamiltonian H , a block-encoding of H with $\alpha = d$ can be efficiently implemented, and it hence implies a simulation algorithm with time complexity $\tilde{O}(d\|H\|t)$. One of the main results in [26] is that working with the qRAM model of a d -sparse Hamiltonian H , a block-encoding with $\alpha = \sqrt{d}$ can be efficiently implemented, which yields a simulation algorithm with time complexity $\tilde{O}(\sqrt{d}\|H\|t)$. This result has the same complexity as ours, but their techniques are more general. In fact, the techniques of [26] has been generalized in [27] to a quantum framework for implementing singular value transformation of matrices. Another merit of [26] is that they gave detailed analysis for application to a quantum linear systems solvers in the qRAM model.

Quantum-inspired classical algorithms for Hamiltonian simulation. Recently (after the first version of this paper was made public), several classical algorithms [28, 29] have been proposed for simulating Hamiltonians in a classical input model, namely, the *sampling and query access*, which is comparable to the qRAM model we use. The algorithm in [28] is efficient in the low-rank and sparse regime, and the algorithm in [29] is efficient when H is low-rank: their time complexity scales as $\text{poly}(t, \|H\|_F, 1/\epsilon)$, where $\|H\|_F$ is the Frobenius norm of H . While these results have ruled out the possibility of exponential speedups of our algorithm in the low-rank regime, noting that the degrees in the polynomials of the time complexity of these classical algorithms are large, our work still has a polynomial speedup over classical algorithm (for low-rank Hamiltonians and dense Hamiltonians).

To summarize this subsection, we provide Table 1 for state-of-the-art algorithms for Hamiltonian simulation in different models and how our techniques compare with theirs in the qRAM model.

Table 1. Comparing our result $O(t\sqrt{d}\|H\| \text{polylog}(t, d, \|H\|, 1/\epsilon))$ with other quantum and classical algorithms for different models. Since the qRAM model is stronger than the sparse-access model and the classical sampling and query access model, we consider the advantage of our algorithm against others when they are directly applied to the qRAM model.

Model	State-of-the-art	Advantage of our results
Sparse-access with $\ H\ _{\max}$ dependence	$\tilde{O}(td\ H\ _{\max})$ [16]	No advantage
Sparse-access with $\ H\ $ dependence	$O((t\sqrt{d}\ H\)^{1+o(1)}/\epsilon^{o(1)})$ [18]	Subpolynomial improvement in t, d ; exponential improvement in ϵ
qRAM	$\tilde{O}(t\sqrt{d}\ H\)$ [26]	Same result
Classical sampling and query access	$\text{poly}(t, \ H\ _F, 1/\epsilon)$ [29]	Polynomial speedup

1.3 Applications

Unitary implementation. One immediate application of simulating non-sparse Hamiltonians is the unitary implementation problem: given access to the entries of a unitary U , the objective is to construct a quantum circuit to approximate U with precision ϵ . As proposed in [7, 30], unitary implementation can be reduced to Hamiltonian simulation by considering the Hamiltonian

$$H = \begin{pmatrix} 0 & U \\ U^\dagger & 0 \end{pmatrix}, \quad (6)$$

and the fact that $e^{-iH\pi/2}|1\rangle|\psi\rangle = -i|0\rangle U|\psi\rangle$. If the entries of a unitary matrix can be accessed by a black-box query oracle, it is shown in [7] that this unitary operator can be implemented with $O(N^{2/3}\text{polylog}(N)\text{poly}(1/\epsilon))$ queries to the black box. Assume now that the entries of U are stored in a data structure such as in Definition 1. Using the same reduction to Hamiltonian simulation, our Hamiltonian simulating algorithm implies an implementation of U with time complexity (circuit depth) $O(\sqrt{N}\text{polylog}(N, 1/\epsilon))$.

Quantum linear systems solver. Simulating non-sparse Hamiltonians can also be used as a subroutine for solving linear systems of equations for non-sparse matrices. The essence for solving a linear system $A|x\rangle = |b\rangle$ is to apply A^{-1} on $|b\rangle$, assuming here for simplicity that $|b\rangle$ is entirely in the column-space of A . When A is d -sparse, it is shown in [22] that A^{-1} can be approximated as a linear combination of unitaries of the form e^{-iAt} . An efficient quantum algorithm for Hamiltonian simulation such as [16] can then be used as a subroutine, so that this linear system can be solved with gate complexity $O(d\kappa^2 \text{polylog}(N, \kappa/\epsilon))$. If A is non-sparse, the time complexity of their algorithm scales as $\tilde{O}(N)$. Based on the data structure in [19], a quantum algorithm for solving linear systems for non-sparse matrices was described in [23], with time complexity (circuit depth) $O(\kappa^2 \sqrt{N} \text{polylog}(N)/\epsilon)$. If we assume a similar input model as in [19, 23], using our Hamiltonian simulation algorithm, together with the linear combinations of unitaries (LCU) decompositions in [22], we have a quantum algorithm for solving linear systems for non-sparse matrices with time complexity (circuit depth) $O(\kappa^2 \sqrt{N} \text{polylog}(\kappa/\epsilon))$, which is an exponential improvement in error dependence compared to [23].

In the remainder of this paper, we first define our data structure, then describe the algorithm in detail and finally prove the main results. We then finish with a summary of this work and a discussion of the benefits and possible drawbacks of our algorithm.

2 Data structure and quantum walk

We first give the precise definition of the data structure that stores the entries of the Hamiltonian, and then show how this data structure can be used to prepare states that will allow us to perform fast Hamiltonian simulation even for dense matrices. The data structure was introduced in in [19].

Definition 1 (Data Structure). *Let $H \in \mathbb{C}^{N \times N}$ be a Hermitian matrix (where $N = 2^n$), $\|H\|_1$ being the maximum absolute row-sum norm, and $\sigma_j := \sum_k |H_{jk}|$. Each entry H_{jk} is represented with b bits of precision. Define D as an array of N binary trees D_j for $j \in \{0, \dots, N-1\}$. Each D_j corresponds to the row H_j , and its organization is specified by the following rules.*

1. *The leaf node k of the tree D_j stores the value^d H_{jk}^* corresponding to the index-entry pair (j, k, H_{jk}) .*
2. *For the level immediately above the bottom level, i.e., the leaves, and any node level above the leaves, the data stored is determined as follows: suppose the node has two children storing data a , and b respectively (note that a and b are complex numbers). Then the entry that is stored in this node is given by $(|a| + |b|)$.*

An example of the above data structure is shown Fig. 1. Note that for each binary tree D_j in the data structure, the value stored in an internal (non-leaf) node is a real number, while for a leaf node, the value stored is a complex number. The root node of D_j stores the value $\sum_{k=0}^{N-1} |H_{jk}^*|$ and we can calculate the value $\|H\|_1 - \sigma_j$ in constant time. In addition, $\|H\|_1$ can be computed by taking the maximum value of the roots of all the binary trees, which can be done during the construction of the data structure.

^dNote that the conjugation here is necessary. See Eq. (16).




Fig. 1. A example of the data structure that allows for efficient state preparation using a logarithmic number of conditional rotations.

Using this data structure, we can efficiently perform the mapping described in the following technical lemma for efficient state preparation.

Lemma 1 (State Preparation). *Let $H \in \mathbb{C}^{N \times N}$ be a Hermitian matrix (where $N = 2^n$) stored in the data structure as specified in Definition 1. Each entry H_{jk} is represented with b bits of precision. Then the following holds*

1. Let $\|H\|_1$ be defined as above. A quantum computer that has access to the data structure can perform the following mapping for $j \in \{0, \dots, N - 1\}$,

$$|j\rangle|0^{\log N}\rangle|0\rangle \mapsto \frac{1}{\sqrt{\|H\|_1}}|j\rangle \sum_{k=0}^{N-1} |k\rangle \left(\sqrt{H_{jk}^*}|0\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{N}}|1\rangle \right), \tag{7}$$

with time complexity (circuit depth) $O(n^2 b^{5/2} \log b)$, where $\sigma_j = \sum_k |H_{jk}|$, and the square-root satisfies $\sqrt{H_{jk}}(\sqrt{H_{jk}^*})^* = H_{jk}$.

2. The size of the data structure containing all N^2 complex entries is $O(N^2 \log^2(N))$.

In order to perform the mapping we will need the following Lemma, which enables us to efficiently implement the conditional rotations with complex numbers.

Lemma 2. *Let $\theta, \phi_0, \phi_1 \in \mathbb{R}$ and let $\tilde{\theta}, \tilde{\phi}_0, \tilde{\phi}_1$ be the b -bit finite precision representation of θ, ϕ_0 , and ϕ_1 , respectively. Then there exists a unitary U that performs the following mapping:*

$$U : |\tilde{\phi}_0\rangle|\tilde{\phi}_1\rangle|\tilde{\theta}\rangle|0\rangle \mapsto |\tilde{\phi}_0\rangle|\tilde{\phi}_1\rangle|\tilde{\theta}\rangle \left(e^{i\tilde{\phi}_0} \cos(\tilde{\theta})|0\rangle + e^{i\tilde{\phi}_1} \sin(\tilde{\theta})|1\rangle \right). \tag{8}$$

Moreover, U can be implemented with $O(b)$ 1- and 2-qubit gates.

Proof. Define U as

$$U = \left(\sum_{\tilde{\phi}_0 \in \{0,1\}^b} |\tilde{\phi}_0\rangle\langle\tilde{\phi}_0| \otimes e^{i|0\rangle\langle 0|\tilde{\phi}_0} \right) \left(\sum_{\tilde{\phi}_1 \in \{0,1\}^b} |\tilde{\phi}_1\rangle\langle\tilde{\phi}_1| \otimes e^{i|1\rangle\langle 1|\tilde{\phi}_1} \right) \left(\sum_{\tilde{\theta} \in \{0,1\}^b} |\tilde{\theta}\rangle\langle\tilde{\theta}| \otimes e^{-iY\tilde{\theta}} \right), \tag{9}$$

where $Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$ is the Pauli Y matrix.

To implement the operator $\sum_{\tilde{\theta} \in \{0,1\}^b} |\tilde{\theta}\rangle\langle\tilde{\theta}| \otimes e^{-iY\tilde{\theta}}$, we use one rotation controlled on each qubit of the first register, with the rotation angles halved for each successive bit. The other two factors of U can be implemented in a similar way. Therefore, U can be implemented with $O(b)$ 1- and 2-qubit gates. \square

Before proving Lemma 1, we first describe the construction and the size of the data structure. Readers may refer to [19] for more details of this data structure.

- The data structure is built from N binary trees $D_i, i \in \{0, \dots, N-1\}$ and we start with an empty tree.
- When a new entry (i, j, H_{ij}) arrives, we create or update the leaf node j in the tree D_i , where the adding of the entry takes $O(\log(N))$ time, since the depth of the tree for $H \in \mathbb{C}^{N \times N}$ is at most $\log(N)$. Since the path from the root to the leaf is of length at most $\log(N)$ (under the assumption that $N = 2^n$), we have furthermore to update at most $\log(N)$ nodes, which can be done in $O(\log(N))$ time if we store an ordered list of the levels in the tree.
- The total time for updating the tree with a new entry is given by $\log(N) \times \log(N) = \log^2(N)$.
- The memory requirements for k entries are given by $O(k \log^2(N))$ as for every entry (j, k, H_{jk}) at least $\log(N)$ nodes are added and each node requires at most $O(\log(N))$ bits.

Now we are ready to prove Lemma 1.

Proof of Lemma 1. With this data structure, we can perform the mapping specified in Eq. (7), with the following steps. For each j , we start from the root of D_j . Starting with the initial state $|j\rangle|0^{\log N}\rangle|0\rangle$, first apply the rotation (according to the value stored in the root node and calculating the normalisation in one query) on the last register to obtain the state

$$\frac{1}{\sqrt{\|H\|_1}} |0\rangle|0^{\log N}\rangle \left(\sqrt{\sum_{k=0}^{N-1} |H_{jk}^*|} |0\rangle + \sqrt{\|H\|_1 - \sigma_j} |1\rangle \right). \quad (10)$$

Then a sequence of conditional rotations is applied on each qubit of the second register to obtain the state as in Eq. (7). At level ℓ of the binary tree D_j , a query to the data structure is made to load the data c (stored in the node) into a register in superposition, the rotation to perform is proportional to $(\sqrt{c}, \sqrt{(\|H\|_1 - \sigma_j)/2^\ell})$ (assuming at the root, $\ell = 0$, and for the leaves, $\ell = \log N$). Then the rotation angles will be determined by calculating the square root and trigonometric functions on the output of the query: this can be implemented with $O(b^{5/2})$ 1- and 2-qubit gates using simple techniques based on Taylor series and long multiplication as in [16], where the error is smaller than that caused by truncating to b bits. Then the conditional rotation is applied by the circuit described in Lemma 2, and the cost for the conditional rotation is $O(b)$. There are $n = \log(N)$ levels, so the cost excluding the implementation of the oracle is $O(nb^{5/2})$. To obtain quantum access to the classical data

structure, a quantum addressing scheme is required. One addressing scheme described in [31] can be used. Although the circuit size of this addressing scheme is $\tilde{O}(N)$ for each D_j , its circuit depth is $O(n)$. Therefore, the time complexity (circuit depth) for preparing the state in Eq. (7) is $O(n^2 b^{5/2} \log n)$.

We use the following rules to determine the sign of the square-root of a complex number: if H_{jk} is not a negative real number, we write $H_{jk} = r e^{i\varphi}$ (for $r \geq 0$ and $-\pi \leq \varphi \leq \pi$) and take $\sqrt{H_{jk}^*} = \sqrt{r} e^{-i\varphi/2}$; when H_{jk} is a negative real number, we take $\sqrt{H_{jk}^*} = \text{sign}(j-k) i \sqrt{|H_{jk}|}$ to avoid the sign ambiguity. With this recipe, we have $\sqrt{H_{jk}} (\sqrt{H_{jk}^*})^* = H_{jk}$. \square

In the following, we demonstrate a state preparation procedure based on the data structure in Fig. 1. In this example, the initial state (omitting the first register) is $|00\rangle|0\rangle$. Let $\sigma_j = |c_0| + |c_1| + |c_2| + |c_3|$. Apply the first rotation, we obtain the state

$$\begin{aligned} & \frac{1}{\sqrt{\|H\|_1}} |00\rangle \left(\sqrt{|c_0| + |c_1| + |c_2| + |c_3|} |0\rangle + \sqrt{\|H\|_1 - \sigma_j} |1\rangle \right) = \\ & \frac{1}{\sqrt{\|H\|_1}} \left(\sqrt{|c_0| + |c_1| + |c_2| + |c_3|} |00\rangle|0\rangle + \sqrt{\|H\|_1 - \sigma_j} |00\rangle|1\rangle \right). \end{aligned} \tag{11}$$

Then, apply a rotation on the first qubit of the first register conditioned on the last register, we obtain the state

$$\frac{1}{\sqrt{\|H\|_1}} \left(\left(\sqrt{|c_0| + |c_1|} |00\rangle + \sqrt{|c_2| + |c_3|} |10\rangle \right) |0\rangle + \left(\sqrt{\frac{\|H\|_1 - \sigma_j}{2}} (|00\rangle + |10\rangle) \right) |1\rangle \right). \tag{12}$$

Next, apply a rotation on the second qubit of the first register conditioned on the first qubit of the first register and last register, we obtain the desired state:

$$\begin{aligned} & \frac{1}{\sqrt{\|H\|_1}} \left(\sqrt{c_0} |00\rangle|0\rangle + \sqrt{c_1} |01\rangle|0\rangle + \sqrt{c_2} |10\rangle|0\rangle + \sqrt{c_3} |11\rangle|0\rangle + \right. \\ & \left. \sqrt{\frac{\|H\|_1 - \sigma_j}{4}} |00\rangle|1\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{4}} |01\rangle|1\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{4}} |10\rangle|1\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{4}} |11\rangle|1\rangle \right) = \\ & \frac{1}{\sqrt{\|H\|_1}} \left(|00\rangle \left(\sqrt{c_0} |0\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{4}} |1\rangle \right) + |01\rangle \left(\sqrt{c_1} |0\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{4}} |1\rangle \right) + \right. \\ & \left. |10\rangle \left(\sqrt{c_2} |0\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{4}} |1\rangle \right) + |11\rangle \left(\sqrt{c_3} |0\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{4}} |1\rangle \right) \right). \end{aligned} \tag{13}$$

Based on the data structure specified in Definition 1 and the efficient state preparation in Lemma 1, we construct a quantum walk operator for H as follows. First define the isometry T as

$$T = \sum_{j=0}^{N-1} \sum_{b \in \{0,1\}} (|j\rangle\langle j| \otimes |b\rangle\langle b|) \otimes |\varphi_{jb}\rangle, \tag{14}$$

with $|\varphi_{j_1}\rangle = |0\rangle|1\rangle$ and

$$|\varphi_{j_0}\rangle = \frac{1}{\sqrt{\|H\|_1}} \sum_{k=0}^{N-1} |k\rangle \left(\sqrt{H_{jk}^*} |0\rangle + \sqrt{\frac{\|H\|_1 - \sigma_j}{N}} |1\rangle \right), \tag{15}$$

where $\sigma_j = \sum_{k=0}^{N-1} |H_{jk}|$. Let S be the swap operator that maps $|j_0\rangle|b_0\rangle|j_1\rangle|b_1\rangle$ to $|j_1\rangle|b_1\rangle|j_0\rangle|b_0\rangle$, for all $j_0, j_1 \in \{0, \dots, N-1\}$ and $b_0, b_1 \in \{0, 1\}$. We observe that

$$\langle j|\langle 0|T^\dagger ST|k\rangle|0\rangle = \frac{\sqrt{H_{jk}} \left(\sqrt{H_{jk}^*} \right)^*}{\|H\|_1} = \frac{H_{jk}}{\|H\|_1}, \tag{16}$$

where the second equality is ensured by the choice of the square-root as in the proof of Lemma 1. This implies that

$$(I \otimes \langle 0|)T^\dagger ST(I \otimes |0\rangle) = \frac{H}{\|H\|_1}. \tag{17}$$

The quantum walk operator U is defined as

$$U = iS(2TT^\dagger - I). \tag{18}$$

A more general characterization of the eigenvalues of quantum walks is presented in [32]. Here we give a specific proof on the relationship between the eigenvalues of U and H as follows.

Lemma 3. *Let the unitary operator U be defined as in Eq. (18), and let λ be an eigenvalue of H with eigenstate $|\lambda\rangle$. It holds that*

$$U|\mu_\pm\rangle = \mu_\pm|\mu_\pm\rangle, \tag{19}$$

where

$$|\mu_\pm\rangle = (T + i\mu_\pm ST)|\lambda\rangle|0\rangle, \text{ and} \tag{20}$$

$$\mu_\pm = \pm e^{\pm i \arcsin(\lambda/\|H\|_1)}. \tag{21}$$

Proof. By the fact that $T^\dagger T = I$ and $(I \otimes \langle 0|)T^\dagger ST(I \otimes |0\rangle) = H/\|H\|_1$, and $(I \otimes \langle 1|)T^\dagger ST(I \otimes |0\rangle) = 0$, we have

$$U|\mu_\pm\rangle = \mu_\pm T|\lambda\rangle|0\rangle + i \left(1 + \frac{2\lambda i}{\|H\|_1} \mu_\pm \right) ST|\lambda\rangle|0\rangle. \tag{22}$$

In order for this state being an eigenstate, it must hold that

$$1 + \frac{2\lambda i}{\|H\|_1} \mu_\pm = \mu_\pm^2, \tag{23}$$

and the solution is

$$\mu_\pm = \frac{\lambda i}{\|H\|_1} \pm \sqrt{1 - \frac{\lambda^2}{\|H\|_1^2}} = \pm e^{\pm i \arcsin(\lambda/\|H\|_1)}. \tag{24}$$

□

3 Linear combination of unitaries and Hamiltonian simulation

To see how to convert the quantum walk operator U to Hamiltonian simulation, we first consider the generating function for the Bessel function, denoted by $J_m(\cdot)$. According to [33, (9.1.41)], we have

$$\sum_{m=-\infty}^{\infty} J_m(z)\mu_{\pm}^m = \exp\left(\frac{z}{2}\left(\mu_{\pm} - \frac{1}{\mu_{\pm}}\right)\right) = e^{iz\lambda/\|H\|_1}, \tag{25}$$

where the second equality follows from Eq. (21) and the fact that $\sin(x) = (e^{ix} - e^{-ix})/2i$. This leads to the following linear combination of unitaries:

$$V_{\infty} = \sum_{m=-\infty}^{\infty} \frac{J_m(z)}{\sum_{j=-\infty}^{\infty} J_j(z)} U^m = \sum_{m=-\infty}^{\infty} J_m(z) U^m = e^{izH/\|H\|_1}, \tag{26}$$

where the second equality follows from the fact that $\sum_{j=-\infty}^{\infty} J_j(z) = 1$.

Now we consider an approximation to $e^{-izH/\|H\|_1}$ in terms of Eq. (26):

$$V_k = \sum_{m=-k}^k \frac{J_m(z)}{\sum_{j=-k}^k J_j(z)} U^m. \tag{27}$$

Here the coefficients are normalized by $\sum_{j=-k}^k J_j(z)$ so that they sum to 1. This will minimize the approximation error (see the proof of Lemma 4 in Appendix B, and the normalization trick was originated in [16]). The eigenvalues of V_k are

$$\sum_{m=-k}^k \frac{J_m(z)}{\sum_{j=-k}^k J_j(z)} \mu_{\pm}^m. \tag{28}$$

Note that each eigenvalue of V_k does not depend on \pm as $J_{-m}(z) = (-1)^m J_m(z)$.

To bound the error in this approximation, we have the technical lemma, and the proof is shown in Appendix B.

Lemma 4. *Let V_k and V_{∞} be defined as above. There exists a positive integer k satisfying $k \geq |z|$ and*

$$k = O\left(\frac{\log(\|H\|/(\|H\|_1 \epsilon))}{\log \log(\|H\|/(\|H\|_1 \epsilon))}\right), \tag{29}$$

such that

$$\|V_k - V_{\infty}\| \leq \epsilon. \tag{30}$$

In the following, we provide technical lemmas for implementing linear combination of unitaries. Suppose we are given the implementations of unitaries U_0, U_1, \dots, U_{m-1} , and coefficients $\alpha_0, \alpha_1, \dots, \alpha_{m-1}$. Then the unitary

$$V = \sum_{j=0}^{m-1} \alpha_j U_j \tag{31}$$

can be implemented probabilistically by the technique called linear combination of unitaries (LCU) [34]. Provided $\sum_{j=0}^{m-1} |\alpha_j| \leq 2$, V can be implemented with success probability $1/4$. To achieve this, we define the multiplexed- U operation, which is denoted by multi- U , as

$$\text{multi-}U|j\rangle|\psi\rangle = |j\rangle U_j|\psi\rangle. \quad (32)$$

We summarize the probabilistic implementation of V in the following lemma, whose proof is shown in Appendix B.

Lemma 5. *Let multi- U be defined as above. If $\sum_{j=0}^{m-1} |\alpha_j| \leq 2$, then there exists a quantum circuit that maps $|0\rangle|0\rangle|\psi\rangle$ to the state*

$$\frac{1}{2}|0\rangle|0\rangle \left(\sum_{j=0}^{m-1} \alpha_j U_j |\psi\rangle \right) + \frac{\sqrt{3}}{2} |\Phi^\perp\rangle, \quad (33)$$

where $(|0\rangle\langle 0| \otimes |0\rangle\langle 0| \otimes I)|\Phi^\perp\rangle = 0$. Moreover, this quantum circuit uses $O(1)$ applications of multi- U and $O(m)$ 1- and 2-qubit gates.

Let W be the quantum circuit in Lemma 5, and let P be the projector defined as $P = |0\rangle\langle 0| \otimes |0\rangle\langle 0| \otimes I$. We have

$$PW|0\rangle|0\rangle|\psi\rangle = \frac{1}{2}|0\rangle|0\rangle \sum_{j=0}^{m-1} \alpha_j U_j |\psi\rangle. \quad (34)$$

If $\sum_{j=0}^{m-1} \alpha_j U_j$ is a unitary operator, one application of the oblivious amplitude amplification operator $-W(1-2P)W^\dagger(1-2P)W$ implements $\sum_{j=0}^{m-1} U_j$ with certainty [8]. However, in our application, the unitary operator \widetilde{W} implements an approximation of V_∞ in the sense that

$$P\widetilde{W}|0\rangle|0\rangle|\psi\rangle = \frac{1}{2}|0\rangle|0\rangle V_k |\psi\rangle, \quad (35)$$

with $\|V_k - V_\infty\| \leq \epsilon$. The following lemma shows that the error caused by the oblivious amplitude amplification is bounded by $O(\epsilon)$, and the proof is given in Appendix B.

Lemma 6. *Let the projector P be defined as above. If a unitary operator \widetilde{W} satisfies $P\widetilde{W}|0\rangle|0\rangle|\psi\rangle = \frac{1}{2}|0\rangle|0\rangle\widetilde{V}|\psi\rangle$ where $\|\widetilde{V} - V\| \leq \epsilon$. Then $\| -\widetilde{W}(I-2P)\widetilde{W}^\dagger(I-2P)\widetilde{W}|0\rangle|0\rangle|\psi\rangle - |0\rangle|0\rangle V|\psi\rangle \| = O(\epsilon)$.*

Now we are ready to prove Theorem 1.

Proof of Theorem 1. The proof we outline here follows closely the proof given in [16]. The intuition of this algorithm is to divide the simulation into $O(t\|H\|_1)$ segments, with each segment simulating $e^{-iH/2}$. To implement each segment, we use the LCU technique to implement V_k defined in Eq. (27), with coefficients $\alpha_m = J_m(z)/\sum_{j=-k}^k J_j(z)$. When $z = -1/2$, we have $\sum_{j=-k}^k |\alpha_j| < 2$. Actually, this holds for all $|z| \leq 1/2$ because

$$\begin{aligned} \sum_{j=-k}^k |\alpha_j| &\leq \sum_{j=-k}^k \frac{|J_j(z)|}{1 - 4\frac{|z/2|^{k+1}}{(k+1)!}} \leq \sum_{j=-k}^k \frac{|z/2|^{|j|}}{|j|!} \left(1 - \frac{4|z/2|^{k+1}}{(k+1)!}\right)^{-1} \\ &< \frac{8}{7} + \frac{16}{7} \sum_{j=1}^{\infty} \frac{1}{4^j j!} = \frac{8}{7} + \frac{16}{7} (\sqrt[4]{e} - 1) < 2, \end{aligned} \quad (36)$$

where the first inequality follows from the fact that $\sum_{j=-k}^k J_j(z) \geq 1 - 4|z/2|^{k+1}/(k+1)!$ (see [16]), the second inequality follows from the fact that $|J_m(z)| \leq |z/2|^{|m|}/|m|!$ (see [33, (9.1.5)]), and the third inequality uses the assumption that $|z| \leq 1/2$. Now, Lemmas 5 and 6 can be applied. By Eq. (25) and using Lemma 4, set

$$k = O\left(\frac{\log(\|H\|/(\|H\|_1 \epsilon'))}{\log \log(\|H\|/(\|H\|_1 \epsilon'))}\right), \tag{37}$$

and we obtain a segment that simulates $e^{-iH/(2\|H\|_1)}$ with error bounded by $O(\epsilon')$. Repeat the segment $O(t\|H\|_1)$ times with error $\epsilon' = \epsilon/(t\|H\|_1)$, and we obtain a simulation of e^{-iHt} with error bounded by ϵ . It suffices to take

$$k = O\left(\frac{\log(t\|H\|/\epsilon)}{\log \log(t\|H\|/\epsilon)}\right). \tag{38}$$

By Lemma 5, each segment can be implemented by $O(1)$ application of multi- U and $O(k)$ 1- and 2-qubit gates, as well as the cost for computing the coefficients α_m for $m \in \{-k, \dots, k\}$. The cost for each multi- U is k times the cost for implementing the quantum walk U . By Lemma 1, the state in Eq. (15) can be prepared with time complexity (circuit depth) $O(n^2 b^{5/2})$, where b is the number of bit of precision. To achieve the overall error bound ϵ , we choose $b = O(\log(t\|H\|_1/\epsilon))$. Hence the time complexity for the state preparation is $O(n^2 \log^{5/2}(t\|H\|_1/\epsilon))$, which is also the time complexity for applying the quantum walk U . Therefore, the time complexity for one segment is

$$O\left(n^2 \log^{5/2}(t\|H\|_1/\epsilon) \frac{\log(t\|H\|/\epsilon)}{\log \log(t\|H\|/\epsilon)}\right). \tag{39}$$

Considering $O(t\|H\|_1)$ segments, the time complexity is as claimed. □

Note that the coefficients $\alpha_{-k}, \dots, \alpha_k$ (for k defined in Eq. (37)) in Lemma 5 can be classically computed using the methods in [35, 36], and the cost is $O(k)$ times the number of bits of precision, which is $O(\log(t\|H\|/\epsilon))$. This is no larger than the quantum time complexity.

4 Discussion

We presented a quantum Hamiltonian simulation algorithm whose time complexity has $\tilde{O}(\sqrt{N})$ dependence for non-sparse Hamiltonians with access to a seemingly more powerful input model. Our technique for Hamiltonian simulation combines ideas from linear combination of quantum walks and an efficient memory model which prepares a special states to provide the improved performance. Our algorithm can be directly applied as a subroutine for the unitary implementation problem and for a quantum linear systems solver. Especially for the latter application, many quantum machine learning algorithms are based on solving linear systems in the same input model, and our algorithm implies exponentially improved error dependence over [23] for this application.

We note that the data structure in qRAM may require a large overhead in practice, if the data were not already stored in it. Furthermore it might be hard to implement such a data structure physically due to the exponential amount of quantum resources [37, 38, 39]. However, if the Hamiltonian is highly structured (i.e., the entries repeat in some pattern),

the memory model could be efficiently implemented. Another potential point of criticism is the required error rate of such a device, as some computations will require an error rate per gate of $O(1/\text{poly}(N))$ to retain a feasible error rate for applications [40]. Whereas, not all computations might need such low error rates [40] and hence the feasibility of our algorithm as a subroutine in an explicit algorithm must be further validated experimentally. To summarize, our algorithm inherits many problems of Grover's search algorithm for unordered classical data, as well as many qRAM-based quantum machine learning algorithms.

However, if the above mentioned caveats can be overcome, our algorithm still supplies a polynomial speedup over known quantum algorithms for Hamiltonian simulation and furthermore allows for a polynomial speedup in comparison with the best known classical algorithms for several practical problems, such as linear regression or linear systems.

Acknowledgements

We thank Richard Cleve and Simone Severini for the discussion and comments on this project. We also thank anonymous reviewers for their valuable suggestions and comments on this paper. CW acknowledges financial support by a David R. Cheriton Graduate Scholarship. LW acknowledges financial support by the Royal Society through a Research Fellow Enhancement Award.

References

1. Richard P Feynman. Simulating physics with computers. *International journal of theoretical physics*, 21(6-7):467–488, 1982.
2. Richard P Feynman. Quantum mechanical computers. *Foundations of physics*, 16(6):507–531, 1986.
3. Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Review*, 41(2):303–332, 1999.
4. Seth Lloyd. Universal quantum simulators. *Science*, pages 1073–1078, 1996.
5. Dorit Aharonov and Amnon Ta-Shma. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the 35th annual ACM symposium on Theory of Computing (STOC 2003)*, pages 20–29. ACM, 2003.
6. Dominic W Berry, Graeme Ahokas, Richard Cleve, and Barry C Sanders. Efficient quantum algorithms for simulating sparse Hamiltonians. *Communications in Mathematical Physics*, 270(2):359–371, 2007.
7. Dominic W Berry and Andrew M Childs. Black-box Hamiltonian simulation and unitary implementation. *Quantum Information & Computation*, 12(1-2):29–62, 2012.
8. Dominic W Berry, Andrew M Childs, Richard Cleve, Robin Kothari, and Rolando D Somma. Exponential improvement in precision for simulating sparse Hamiltonians. *Forum of Mathematics, Sigma*, 5, 2017.
9. Andrew M Childs. On the relationship between continuous-and discrete-time quantum walk. *Communications in Mathematical Physics*, 294(2):581–603, 2010.
10. Andrew M Childs and Nathan Wiebe. Hamiltonian simulation using linear combinations of unitary operations. *Quantum Information & Computation*, 12(11-12):901–924, 2012.
11. David Poulin, Angie Qarry, Rolando Somma, and Frank Verstraete. Quantum simulation of time-dependent Hamiltonians and the convenient illusion of hilbert space. *Physical Review Letters*, 106(17):170501, 2011.
12. Nathan Wiebe, Dominic W Berry, Peter Høyer, and Barry C Sanders. Simulating quantum dynamics on a quantum computer. *Journal of Physics A: Mathematical and Theoretical*, 44(44):445308, 2011.

13. Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019.
14. Dominic W Berry and Leonardo Novo. Corrected quantum walk for optimal Hamiltonian simulation. *Quantum Information & Computation*, 16(15-16):1295–1317, 2016.
15. Guang Hao Low and Isaac L Chuang. Hamiltonian simulation by uniform spectral amplification. *arXiv preprint arXiv:1707.05391*, 2017.
16. Dominic W Berry, Andrew M Childs, and Robin Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science (FOCS 2015)*, pages 792–809. IEEE, 2015.
17. Guang Hao Low and Isaac L Chuang. Optimal Hamiltonian simulation by quantum signal processing. *Physical Review Letters*, 118(1):010501, 2017.
18. Guang Hao Low. Hamiltonian simulation with nearly optimal dependence on spectral norm. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, pages 491–502, 2019.
19. Iordanis Kerenidis and Anupam Prakash. Quantum recommendation systems. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
20. Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature Physics*, 10(9):631, 2014.
21. Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13):130503, 2014.
22. Andrew M Childs, Robin Kothari, and Rolando D Somma. Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. *SIAM Journal on Computing*, 46(6):1920–1950, 2017.
23. Leonard Wossnig, Zhikuan Zhao, and Anupam Prakash. Quantum linear system algorithm for dense matrices. *Physical Review Letters*, 120(5):050502, 2018.
24. Andrew M. Childs and Robin Kothari. Limitations on the simulation of non-sparse Hamiltonians. *Quantum Information & Computation*, 10(7&8):669–684, 2010.
25. Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15):150502, 2009.
26. Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. The power of block-encoded matrix powers: Improved regression techniques via faster Hamiltonian simulation. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132, pages 33:1–33:14, 2019.
27. András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204. ACM, 2019.
28. Alessandro Rudi, Leonard Wossnig, Carlo Ciliberto, Andrea Rocchetto, Massimiliano Pontil, and Simone Severini. Approximating Hamiltonian dynamics with the Nyström method. *Quantum*, 4:234, 2020.
29. Nai-Hui Chia, András Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. *arXiv preprint arXiv:1910.06151*, 2019.
30. Stephen P Jordan and Pawel Wocjan. Efficient quantum circuits for arbitrary sparse unitaries. *Physical Review A*, 80(6):062301, 2009.
31. Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum random access memory. *Physical Review Letters*, 100(16):160501, 2008.
32. Mario Szegedy. Quantum speed-up of markov chain based algorithms. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2004)*, pages 32–41. IEEE, 2004.
33. Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, volume 55. Courier Corporation, 1964.

34. Robin Kothari. *Efficient algorithms in quantum query complexity*. PhD thesis, University of Waterloo, 2014.
35. W. G. Bickley, L. J. Comrie, J. C. P. Miller, D. H. Sadler, and A. J. Thompson. *Bessel Functions: Part II. Functions of Positive Integer Order*. Cambridge University Press, 1960.
36. F. W. J. Olver. Error analysis of miller's recurrence algorithm. *Mathematics of Computation*, 18(85):65–74, 1964.
37. Scott Aaronson. Read the fine print. *Nature Physics*, 11(4):291, 2015.
38. Jeremy Adcock, Euan Allen, Matthew Day, Stefan Frick, Janna Hinchliff, Mack Johnson, Sam Morley-Short, Sam Pallister, Alasdair Price, and Stasja Stanisic. Advances in quantum machine learning. *arXiv preprint arXiv:1512.02900*, 2015.
39. Carlo Ciliberto, Mark Herbster, Alessandro Davide Ialongo, Massimiliano Pontil, Andrea Rocchetto, Simone Severini, and Leonard Wossnig. Quantum machine learning: a classical perspective. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2209):20170551, 2018.
40. Srinivasan Arunachalam, Vlad Gheorghiu, Tomas Jochym-O'Connor, Michele Mosca, and Priyaa Varshinee Srinivasan. On the robustness of bucket brigade quantum ram. *New Journal of Physics*, 17(12):123010, 2015.
41. Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 1990.

Appendix A The relation between $\|H\|_1$ and $\|H\|$

We prove the following proposition.

Proposition 1. *If $A \in \mathbb{C}^{N \times N}$ has at most d non-zero entries in any row, it holds that $\|A\|_1 \leq \sqrt{d}\|A\|$.*

Proof. First observe that $\|A\|^2 \leq \sum_i \lambda_i(A^\dagger A) = \text{Tr}(A^\dagger A) = \|A\|_F^2$, and furthermore we have that $\sum_{ij} |a_{ij}|^2 \leq d \max_{j \in [N]} \sum_i |a_{ij}|^2$. From this we have that $\|A\| \leq \sqrt{d}\|A\|_1$ for a d -sparse A . By [41, Theorem 5.6.18], we have that $\|A\|_1 \leq C_M(1, *)\|A\|$ for $C_M(1, *) = \max_{A \neq 0} \frac{\|A\|}{\|A\|_1}$ and using the above we have $C_M(1, *) \leq \sqrt{d}$. Therefore we find that $\|A\|_1 \leq \sqrt{d}\|A\|$ as desired. \square

It immediately follows that $\|A\|_1 \leq \sqrt{N}\|A\|$ for dense A , i.e., $d = N$.

Appendix B Proofs of technical lemmas

Proof of Lemma 4. The proof outlined here follows closely the proof of Lemma 8 in [16]. Recalling the definition of V_k and V_∞ , we define the weights in V_k by

$$\alpha_m := \frac{J_m(z)}{C_k}, \tag{B.1}$$

where $C_k = \sum_{l=-k}^k J_l(z)$. The normalization here is chosen so that $\sum_m \alpha_m = 1$ which will give the best result [16].

Since

$$\sum_{m=-\infty}^{\infty} J_m(z) = \sum_{m=[-\infty: -k-1; k+1: \infty]} J_m(z) + \sum_{m=-k}^k J_m(z) = 1, \tag{B.2}$$

observe that we have two error sources. The first one comes from the truncation of the series, and the second one comes from the different renormalization of the terms which introduces an

error in the first $|m| \leq k$ terms in the sum. We therefore start by bounding the normalization factor C_k . For the Bessel-functions for all m it holds that $|J_m(z)| \leq \frac{1}{|m!}| \left|\frac{z}{2}\right|^{|m|}$, since $J_{-m}(z) = (-1)^m J_m(z)$ [33, (9.1.5)]. For $|m| \leq k$ we can hence find the following bound on the truncated part

$$\begin{aligned} \sum_{m=[-\infty: -k-1; k+1: \infty]} |J_m(z)| &= 2 \sum_{m=k+1}^{\infty} |J_m(z)| \leq 2 \sum_{m=k+1}^{\infty} \frac{|z/2|^m}{m!} \\ &= 2 \frac{|z/2|^{k+1}}{(k+1)!} \left(1 + \frac{|z/2|}{k+2} + \frac{|z/2|^2}{(k+2)(k+3)} + \dots \right) \\ &< 2 \frac{|z/2|^{k+1}}{(k+1)!} \sum_{m=k+1}^{\infty} \left(\frac{1}{2}\right)^{m-(k+1)} = \frac{4|z/2|^{k+1}}{(k+1)!}. \end{aligned} \tag{B.3}$$

Since $\sum_m J_m(z) = 1$, based on the normalization, we hence find that

$$\sum_{m=-k}^k J_m(z) \geq \left(1 - \frac{4|z/2|^{k+1}}{(k+1)!} \right), \tag{B.4}$$

which is a lower bound on the normalization factor C_k . Since $a_m = \frac{J_m(z)}{C_k}$, the correction is small, which implies that

$$a_m = J_m(z) \left(1 + O\left(\frac{|z/2|^{k+1}}{(k+1)!}\right) \right), \tag{B.5}$$

and we have a multiplicative error based on the renormalization.

Next we want to bound the error in the truncation before we join the two error sources. From Eq. (25) we know that

$$e^{iz\lambda/\Lambda} - 1 = \sum_{m=-\infty}^{\infty} J_m(z)(\mu_{\pm}^m - 1), \tag{B.6}$$

by the normalization of $\sum_m J_m(z)$. From this we can see that we can hence obtain a bound on the truncated $J_m(z)$ as follows.

$$\sum_{m=-k}^k J_m(z)(\mu_{\pm}^m - 1) = e^{iz\lambda/\Lambda} - 1 - \sum_{\substack{m=[-\infty: -(k+1); \\ (k+1): \infty]}} J_m(z)(\mu_{\pm}^m - 1). \tag{B.7}$$

Therefore we can upper bound the left-hand side in terms of the exact value of V_{∞} , i.e. $e^{iz\lambda/\Lambda}$ if we can bound the right-most term in Eq. (B.7). Using furthermore the bound in Eq. (B.5) we obtain

$$\sum_{m=-k}^k a_m(z)(\mu_{\pm}^m - 1) = \left(e^{iz\lambda/\Lambda} - 1 - \sum_{\substack{m=[-\infty: -(k+1); \\ (k+1): \infty]}} J_m(z)(\mu_{\pm}^m - 1) \right) \left(1 + O\left(\frac{|z/2|^{k+1}}{(k+1)!}\right) \right), \tag{B.8}$$

which reduced with $|2^{iz\lambda/\Lambda} - 1| \leq |z\lambda/\Lambda|$ and $|z| \leq k$ to

$$\sum_{m=-k}^k a_m(z)(\mu_{\pm}^m - 1) = e^{iz\lambda/\Lambda} - 1 - O\left(\sum_{\substack{m=[-\infty:-(k+1); \\ (k+1):\infty]} J_m(z)(\mu_{\pm}^m - 1)\right). \quad (\text{B.9})$$

We can then obtain the desired bound $\|V_{\infty} - V_k\|$ by reordering the above equation, and using that $\sum_{m=-k}^k a_m(z) = 1$ such that we have

$$\|V_{\infty} - V_k\| = \left| \sum_{m=-k}^k a_m \mu_{\pm}^m - e^{iz\lambda/\Lambda} \right| = O\left(\sum_{\substack{m=[-\infty:-(k+1); \\ (k+1):\infty]} J_m(z)(\mu_{\pm}^m - 1)\right). \quad (\text{B.10})$$

We hence only need to bound the right-hand side.

For μ_+ we can use that $|\mu_+^m - 1| \leq 2|m\lambda/\Lambda| =: 2|m\nu|$ and obtain the bound $2\left|\frac{\nu}{k}\right|\left|\frac{z}{2}\right|^{k+1}$ [16]. For the μ_- case we need to refine the analysis and will show that the bound remains the same. Let $\nu := \lambda/\Lambda$ as above. First observe that $J_m(z)\mu_-^m + J_{-m}(z)\mu_i^{-m} = J_{-m}(z)\mu_+^{-m} + J_m(z)\mu_+^m$, and it follows that

$$\begin{aligned} & \sum_{m=-\infty}^{-(k+1)} J_m(z)(\mu_-^m - 1) + \sum_{m=k+1}^{\infty} J_m(z)(\mu_-^m - 1) \\ &= \sum_{m=-\infty}^{-(k+1)} J_m(z)(\mu_+^m - 1) + \sum_{m=k+1}^{\infty} J_m(z)(\mu_+^m - 1). \end{aligned} \quad (\text{B.11})$$

Therefore we only need to treat the μ_+ case.

$$\begin{aligned} & \left| \sum_{\substack{m=[-\infty:-(k+1); \\ (k+1):\infty]} J_m(z)(\mu_+^m - 1) \right| \leq 2 \left| \sum_{m=k+1}^{\infty} J_m(z)(\mu_+^m - 1) \right| \\ & \leq 2 \sum_{m=k+1}^{\infty} |J_m(z)| |\mu_+^m - 1| \\ & = 2 \sum_{m=k+1}^{\infty} \frac{1}{|m|!} \left|\frac{z}{2}\right|^{|m|} |\mu_+^m - 1| \\ & \leq 4 \sum_{m=k+1}^{\infty} \frac{1}{|m|!} \left|\frac{z}{2}\right|^{|m|} m|\nu| \\ & < \frac{8|\nu|}{(k+1)!} \left|\frac{z}{2}\right|^{k+1} (k+2). \end{aligned} \quad (\text{B.12})$$

Using this bound we hence obtain from Eq. (B.10),

$$\|V_{\infty} - V_k\| = \left| \sum_{m=-k}^k a_m \mu_{\pm}^m - e^{iz\lambda/\Lambda} \right| \leq O\left(\frac{\lambda}{k! \Lambda} \left|\frac{z}{2}\right|^{k+1}\right) = O\left(\frac{\|H\|(z/2)^{k+1}}{\Lambda k!}\right). \quad (\text{B.13})$$

In order for the above equation being upper-bounded by ϵ , it suffices to choose some k that is upper bounded as claimed. \square

Proof of Lemma 5. Let $s = \sum_{j=0}^{m-1} |\alpha_j|$. We first define the unitary operator B to prepare the coefficients:

$$B|0\rangle|0\rangle = \left(\sqrt{\frac{s}{2}}|0\rangle + \sqrt{1 - \frac{s}{2}}|1\rangle \right) \otimes \frac{1}{\sqrt{s}} \sum_{j=0}^{m-1} \sqrt{\alpha_j}|j\rangle. \quad (\text{B.14})$$

Define the unitary operator W as $W = (B^\dagger \otimes I)(I \otimes \text{multi-}U)(B \otimes I)$. We claim that W performs the desired mapping, as

$$\begin{aligned} W|0\rangle|0\rangle|\psi\rangle &= (B^\dagger \otimes I)(I \otimes \text{multi-}U)(B \otimes I)|0\rangle|0\rangle|\psi\rangle \\ &= \frac{1}{\sqrt{2}}(B^\dagger \otimes I)|0\rangle \sum_{j=0}^{m-1} \sqrt{\alpha_j}|j\rangle U_j|\psi\rangle + \sqrt{\frac{2-s}{2s}}(B^\dagger \otimes I)|1\rangle \sum_{j=0}^{m-1} \sqrt{\alpha_j}|j\rangle U_j|\psi\rangle \\ &= \frac{1}{2}|0\rangle|0\rangle \sum_{j=0}^{m-1} \alpha_j U_j|\psi\rangle + \sqrt{\gamma}|\Phi^\perp\rangle, \end{aligned} \quad (\text{B.15})$$

where $|\Phi^\perp\rangle$ is a state satisfying $(|0\rangle\langle 0| \otimes |0\rangle\langle 0| \otimes I)|\Phi^\perp\rangle = 0$, and γ is some normalization factor.

The number of applications of $\text{multi-}U$ is constant, as in the definition of W . To implement the unitary operator B , $O(m)$ 1- and 2-qubit gates suffice. \square

Proof of Lemma 6. We have

$$\begin{aligned} -\widetilde{W}(I - 2P)\widetilde{W}^\dagger(I - 2P)\widetilde{W}|0\rangle|0\rangle|\psi\rangle &= (\widetilde{W} + 2P\widetilde{W} - 4\widetilde{W}P\widetilde{W}^\dagger P\widetilde{W})|0\rangle|0\rangle|\psi\rangle \\ &= (\widetilde{W} + 2P\widetilde{W} - 4\widetilde{W}P\widetilde{W}^\dagger P P\widetilde{W}P)|0\rangle|0\rangle|\psi\rangle \\ &= \widetilde{W}|0\rangle|0\rangle|\psi\rangle + |0\rangle|0\rangle\widetilde{V}|\psi\rangle - \widetilde{W}(|0\rangle|0\rangle\widetilde{V}^\dagger\widetilde{V}|\psi\rangle) \end{aligned} \quad (\text{B.16})$$

Because $\|\widetilde{V} - V\| \leq \epsilon$ and V is a unitary operator, we have $\|\widetilde{V}^\dagger\widetilde{V} - I\| = O(\epsilon)$. Therefore, we have

$$\|-\widetilde{W}(I - 2P)\widetilde{W}^\dagger(I - 2P)\widetilde{W}|0\rangle|0\rangle|\psi\rangle - |0\rangle|0\rangle\widetilde{V}|\psi\rangle\| = O(\epsilon). \quad (\text{B.17})$$

Thus

$$\|-\widetilde{W}(I - 2P)\widetilde{W}^\dagger(I - 2P)\widetilde{W}|0\rangle|0\rangle|\psi\rangle - |0\rangle|0\rangle V|\psi\rangle\| = O(\epsilon). \quad (\text{B.18})$$

\square