Quantum State Preparation with Optimal Circuit Depth: Implementations and Applications

Xiao-Ming Zhang, Tongyang Li, and Xiao Yuan*
Center on Frontiers of Computing Studies, Peking University, Beijing 100871, China
and School of Computer Science, Peking University, Beijing 100871, China

(Received 30 January 2022; revised 1 August 2022; accepted 1 November 2022; published 30 November 2022)

Quantum state preparation is an important subroutine for quantum computing. We show that any n-qubit quantum state can be prepared with a $\Theta(n)$-depth circuit using only single- and two-qubit gates, although with a cost of an exponential amount of ancillary qubits. On the other hand, for sparse quantum states with $d \geqslant 2$ nonzero entries, we can reduce the circuit depth to $\Theta(\log(nd))$ with $O(nd \log d)$ ancillary qubits. The algorithm for sparse states is exponentially faster than best-known results and the number of ancillary qubits is nearly optimal and only increases polynomially with the system size. We discuss applications of the results in different quantum computing tasks, such as Hamiltonian simulation, solving linear systems of equations, and realizing quantum random access memories, and find cases with exponential reductions of the circuit depth for all these three tasks. In particular, using our algorithm, we find a family of linear system solving problems enjoying exponential speedups, even compared to the best-known quantum and classical dequantization algorithms.

DOI: 10.1103/PhysRevLett.129.230504

The speed limit of quantum state preparation is a question with fundamental and practical interests, determining the efficiency of inputting classical data into a quantum computer, and playing as a critical subroutine for many quantum algorithms, such as in machine learning [1–3] and Hamiltonian simulations [4,5]. Without ancillary qubits, an exponential circuit depth is inevitable to prepare an arbitrary quantum state [6–16] and the optimal result $\Theta(2^n/n)$ was recently obtained by Sun *et al.* [17]. Leveraging ancillary qubits, the circuit depth could be reduced to be subexponential scaling [17–23], yet in the worse case with an exponential number of ancillas. Very recently, the optimal circuit depth $\Theta(n)$ was achieved by Refs. [17,21] with $O(2^n)$ [17] and $\tilde{O}(2^n)$ [21] ancillary qubits.

Despite the previous results in minimizing circuit depth, the subexponential circuit depth is only achieved at the cost of exponential space complexity. Moreover, when considering applications in the field of quantum machine learning, strong data structure assumptions leave space for quantum-inspired classical algorithms. With a classical data structure enabling l^2 sampling, there are classical algorithms with polylogarithmic runtime dequantizing the quantum algorithms for recommendation systems [24], solving linear systems [25,26], semidefinite programs [27], etc. These results show that space resources should not be neglected when discussing the quantum exponential advantages.

In practice, the data may behave with a certain structure. Indeed, if the one imposes certain restrictions on the target quantum states, the circuit depth and the ancillary qubit number might be further reduced [28–34]. A typical scenario that has both theoretical and practical relevance is the sparse data structure, such as sparse classical data, Hamiltonians of physics systems, etc. Using a constant number of ancillary qubits, arbitrary d-sparse quantum states (with d nonzero entries) can be prepared using a circuit depth of $O(dn)$ [28–31]. However, it was unclear if the sparse preparation procedure could be further sped up with more, but polynomial, ancillary qubits. The fundamental speed limit of sparse state preparation is still an open question, which is important for studying the ultimate power of quantum machine learning algorithms.

In this Letter, we study the speed limit of quantum state preparation. We first develop a deterministic algorithm (independent of Refs. [17,21]) for preparing an arbitrary quantum state with optimal circuit depth $\Theta(n)$ and $O(2^n)$ ancillary qubits. The scheme requires a much more sparse connectivity than Ref. [17], as each qubit connects to a constant number of other qubits. We next introduce an algorithm for d-sparse quantum states ($d \geqslant 2$) that achieves the optimal circuit depth $\Theta(\log(nd))$, exponentially faster than the best-known results [28–30]. The sparse state preparation requires $O(nd \log d)$ ancillary qubits, which is also nearly optimal. Based on the results, we find a family of linear system tasks that can be solved with the circuit depth and the number of ancillary qubits being $O(\mathrm{poly}(n))$, and hence show an exponential improvement compared to the best known quantum and classical dequantization algorithms. We also show how our techniques can be applied to improving Hamiltonian simulations and quantum random access memories (QRAMs).

Access model.—A general n-qubit state can be expressed as

$$|\psi\rangle = \sum_{k=0}^{N-1} a_k |k\rangle, \tag{1}$$

with $N = 2^n$, $a_k \in \mathbb{C}$, $\sum_{k=0}^{N-1} |a_k|^2 = 1$, and $|k\rangle \equiv |k_n k_{n-1} \cdots k_1\rangle$ being the basis with bits k_j for $j = 1, 2, \ldots, n$. Before discussing our state preparation protocol, we first introduce how our quantum circuit accesses the classical description of a target quantum state. Let $b_{n,k} \equiv |a_k|$, $b_{l,j} \equiv \sqrt{|b_{l+1,2j}|^2 + |b_{l+1,2j+1}|^2}$ for $0 \leqslant l \leqslant n-1$, $\theta_{l,j} = \arccos(b_{l,2j}/b_{l-1,j})$ for $b_{l-1,j} \neq 0$, and $\theta_{l,j} = 0$ for $b_{l-1,j} = 0$. We require classical preprocessing to calculate $\theta_{l,j}$, and $\arg(a_k)$. Here, $b_{l,j}$ are recursively defined so that we can encode the amplitudes in a treelike fashion allowing parallelization. This recursive definition is not required for phase $\arg(a_k)$, because after encoding the amplitude, the phase can be encoded with a single layer of phase gates (see Sec. I of Ref. [35] for details). The preprocessing takes time $O(N)$ by sequential calculations, or $O(\log N)$ by parallel calculations with $O(N)$ space complexity. These complexities are optimal because reading and writing N values already require $\Omega(N)$ resource.

For sparse quantum state with d nonzero elements, the quantum state can be expressed as

$$|\psi\rangle = \sum_{k=0}^{d-1} \psi_k |q_k\rangle, \tag{2}$$

where $\psi_k \in \mathbb{C}$ and q_k is the index (with n digits) of the kth nonzero entries. We assume $d = 2^{\tilde{n}}$ with integer $\tilde{n}$, which can be always satisfied by appending $|q_k\rangle$ with zero amplitude. Similarly, we let $b'_{n,k} \equiv |\psi_k|$, $b'_{l,j} \equiv \sqrt{|b'_{l+1,2j}|^2 + |b'_{l+1,2j+1}|^2}$ for $0 \leqslant l \leqslant \tilde{n}-1$, and $\theta'_{l,j} \equiv \arccos(b'_{l,2j}/b'_{l-1,j})$, and require classical preprocessing to calculate $\theta'_{l,j}$, $\arg(\psi_k)$. The value of q_k should also be encoded to the circuit. The preprocessing time is $O(nd)$ for sequential calculation, or $O(\log(nd))$ for parallel calculation with $O(nd)$ space complexity.

The calculated angles and the labels of nonzero basis $|q_k\rangle$ for sparse states can then be directly mapped to the parameters of the quantum circuit, so the time complexity for generating quantum circuits are identical to the preprocessing time. Note that the preprocessing only needs to be performed once for preparing arbitrary copies of state. Here and after, we assume that the classical preprocessing has been completed.

Quantum state preparation.—Without loss of generality, the task of quantum state preparation is to prepare $|\psi\rangle$ from an initial product state $|0\rangle^{\otimes n}$ using single- and two-qubit gates. The qubit layout of our protocol is illustrated in Figs. 1(a) and 1(b). There is an $(n+1)$-layer binary tree of qubits, H. The lth layer of H is denoted as H_l, and its jth qubit is denoted as $H_{l,j}$. The lth layer of H is connected to the leaf layer of another binary tree V_l with $(l+1)$ layers. In this layout, each qubit connects to at most constant number of the other qubits, while Ref. [17] assumes that two-qubit gates can be applied on any two qubits. With the qubit layout above and the access model introduced previously, we have the following result.

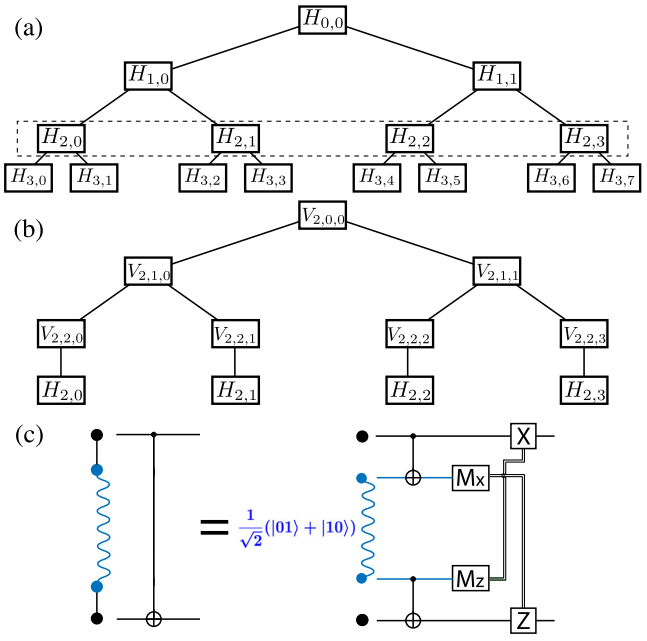

FIG. 1. (a) Layout of binary tree H. Each block represents a qubit. (b) Layout of binary tree V_2, which connects to the second layer of H with dashed box, i.e., H_2. Here, $V_{2,\text{root}}$ is $V_{2,2,0}$. In (a) and (b), CNOT gates are only applied at qubit pairs connected by solid lines. (c) CNOT gate between two distant qubits (black circles) based on pre-shared Bell states (blue circles). $M_{x,z}$ and X, Z represent measurements and Pauli gates [40].

Theorem 1: (Arbitrary quantum state preparation) With only single- and two-qubit gates, an arbitrary n-qubit quantum state can be deterministically prepared with a circuit depth $\Theta(n)$ and $O(N)$ ancillary qubits.

Our method saturates the circuit depth lower bound $\Omega(n)$ [17,19]. Below we sketch our protocol and refer to Sec. I of [35] for the formal description.

The root of H is initialized as $|1\rangle$ and all other qubits are initialized as $|0\rangle$. The protocol contains 5 stages. In stage 1, with a $O(n)$ layer of CNOT and single qubit gates, H is prepared as

$$|\psi_{\text{stage 1}}\rangle = \sum_{k=0}^{2^n-1} a_k |1\rangle_{H_0} \bigotimes_{l=1}^{n} |(k,l)\rangle'_{H_l}, \tag{3}$$

where $|1\rangle_{H_0}$ the state of H_0, and $|(k,l)\rangle'_{H_l}$ is the state of H_l. Here $(k,l) \equiv k_n k_{n-1} \cdots k_{n-l+1}$ represents the last l digits of

k, and $|(k,l)\rangle' \equiv |0\rangle^{\otimes(k,l)}|1\rangle|0\rangle^{\otimes 2^l-(k,l)-1}$. At each layer, there is only one qubit activated (at state $|1\rangle$) while the rest of the qubits are at state $|0\rangle$. The amplitude of the basis $|1\rangle_{H_0}\otimes_{l=1}^{n}|(k,l)\rangle'_{H_l}$ at Eq. (3) is identical to the amplitude of $|k\rangle$ at Eq. (1). So the remaining task is all about basis transformation.

In stage 2, for each l, we map the state of H_l to the state of the root of V_l. More specifically, we perform $|(k,l)\rangle'_{H_l}|0\rangle_{V_{l,\mathrm{root}}} \rightarrow |(k,l)\rangle'_{H_l}|k_{n-l+1}\rangle_{V_{l,\mathrm{root}}}$, where $V_{l,\mathrm{root}} = V_{l,0,0}$. With a total circuit depth $O(n)$, we obtain

$$|\psi_{\mathrm{stage\,2}}\rangle = \sum_{k=0}^{2^n-1} a_k|1\rangle_{H_0}\bigotimes_{l=1}^{n}|(k,l)\rangle'_{H_l}|k_{n-l+1}\rangle_{V_{l,\mathrm{root}}}. \quad (4)$$

In the remaining of the algorithm (stage 3 to 5), our goal is to uncompute H. This can be realized by flipping each qubit of H_l conditioned on the states of its parent and $V_{l,\mathrm{root}}$. By utilizing the binary trees V_l, the uncomputation can be done with $O(n)$ circuit depth [35]. We can then trace out all qubits except for the roots of V_l, and the state becomes $\sum_{k=0}^{2^n-1}\alpha_k\bigotimes_{l=1}^{n}|k_{n-l+1}\rangle_{V_{l,\mathrm{root}}}$, which is equivalent to Eq. (1).

We also show in [35] that our scheme can be approximated to accuracy ε using Clifford + T gates with depth $O(n\log(n/\varepsilon))$. This decomposition is important for fault-tolerant implementation based on surface code [41].

Nonlocal entangling gate implementation.—Similar to QRAMs, encoding exponential data to a quantum state may require spatially nonlocal gates. The nonlocal gate can be realized by quantum network with preshared Bell states [40,42]. As shown in Fig. 1(c), the scheme requires a pair of ancillary qubits at Bell state $1/\sqrt{2}(|01\rangle + |10\rangle)$. Each ancillary qubit couples to either the control or target qubit. Effective CNOT can then be realized with local operations and classical communication. The protocol has been demonstrated in superconducting qubit [40] and trapped-ion systems [43]. Alternatively, nonlocal gates can also be realized with spin-photon network [44]. In fault-tolerant settings, surface code based on teleportated CNOT gate above has also been proposed in [45], which can be straightforwardly applied to our scheme. In [35], we further show that our scheme can be implemented even in a nearest-neighbor coupled two-dimensional qubit array, only at the cost of a mild increase of the ancillary qubit number to $O(n^2N)$.

Sparse quantum state preparation.—The protocol can be further improved if the target states are sparse. We first introduce two subroutines that are useful for sparse state preparation and then discuss several other applications. Both subroutines work on a quantum system containing an *index register* and a *word register*, which are systems with a certain number of qubits.

The first subroutine is the product unitary memory (PUM), which can be considered as a generalization of QRAM protocol in [46,47]. We define an n-word product unitary function $\hat{U}(k) \equiv \bigotimes_{l=n}^{1}\hat{U}_l(k)$ with $\hat{U}_l(k) \in \mathrm{SU}(2)$ and $k \in \{0, 1\cdots, d-1\}$. We define the selector unitary of $\hat{U}$ as $\mathrm{select}(\hat{U})$, which satisfies $\mathrm{select}(\hat{U})|k\rangle|z\rangle = |k\rangle\hat{U}(k)|z\rangle$. Here, $|k\rangle$ is the basis of the $\lceil\log_2 d\rceil$-qubits index register, and $|z\rangle$ is the basis of the n-qubits word register. The selector unitary can also be represented as $\mathrm{select}(\hat{U}) \equiv \sum_{k=0}^{d-1}|k\rangle\langle k|\otimes\hat{U}(k)$. We have the following result (see Sec. III A of [35] for details).

Lemma 1: (PUM) Given an arbitrary $\lceil\log_2 d\rceil$-index, n-word product unitary function $\hat{U}(k)$, $\mathrm{select}(\hat{U})$ can be realized with circuit depth $O(\log(nd))$ and $O(nd)$ ancillary qubits using only single- and two-qubit gates.

The second subroutine is the sparse Boolean memory (SBM). We consider an n-index, $\tilde{n}$-word Boolean function $f:\{0,1\}^n \rightarrow \{0,1\}^{\tilde{n}}$. Let $\mathcal{S}_f = \{k|f(k)\neq 0\ldots 0\}$ containing all input indexes with nonzero output. We say that f is s sparse if $\mathcal{S}_f$ has no more than s elements. Its corresponding sparse Boolean function selector satisfies $\mathrm{select}(f)|k\rangle|z\rangle = |k\rangle|z\oplus f(k)\rangle$, where $\oplus$ represents bitwise XOR. Let $f_l(k)$ be the lth digit of $f(k)$, $\mathrm{select}(f)$ can also be expressed as

$$\mathrm{select}(f) \equiv \sum_{k=0}^{2^n-1}|k\rangle\langle k|\bigotimes_{l=\tilde{n}}^{1}(f_l(k)\hat{X} + \bar{f}_l(k)\hat{\mathbb{I}}_1), \quad (5)$$

where $\bar{f}_l(k)$ is the NOT of $f_l(k)$, $\hat{X}$ is the Pauli-X operator, and $\hat{\mathbb{I}}_m$ represents the m-qubit identity. We have the following result (see Sec. III A of [35] for details).

Lemma 2: (SBM) Given an arbitrary n-index, $\tilde{n}$-word, s-sparse Boolean function f, $\mathrm{select}(f)$ in Eq. (5) can be realized with a quantum circuit with circuit depth $O(\log(ns\tilde{n}))$ and $O(ns\tilde{n})$ ancillary qubits using only single- and two-qubit gates.

Based on Lemmas 1,2, and the access model discussed previously, we are now ready for our sparse state preparation protocol. Our result is as follows.

Theorem 2: (Sparse state preparation) With only single- and two-qubit gates, arbitrary n-qubit, d-sparse ($d \geqslant 2$) quantum states can be deterministically prepared with a circuit depth $\Theta(\log(nd))$ and $O(nd\log d)$ ancillary qubits.

As we prove in Lemma 3 at Sec. II of [35] that the circuit depth is lower bounded by $\Omega(\log(nd))$, our protocol also achieves the optimal circuit depth for sparse states.

Below, we show how our protocol works for preparing $\sum_{k=0}^{d-1}\psi_k|q_k\rangle$ in Eq. (2). We introduce registers A and B, consisting of $\tilde{n} = \lceil\log_2 d\rceil$ and n qubits, respectively. All qubits are initialized to $|0\rangle$. Then, we prepare register A to state $\sum_{k=0}^{d-1}\psi_k|k\rangle_A$, which uses $O(\log d)$ circuit depth and $O(d)$ ancillary qubits according to Theorem 1. Next, we introduce an n-word product unitary function $\hat{U}_{\mathrm{prep}}(k) = \otimes_{j=n}^{1}(q_{k,j}\hat{X} + \bar{q}_{k,j}\hat{\mathbb{I}}_1)$, where $q_{k,j}$ is the jth digit

of q_k. We query select $(\hat{U}_{\text{prep}})$ with register A as index register and register B as word register, and the output state is $\sum_{k=0}^{d-1}\psi_k|k\rangle_A|q_k\rangle_B$. According to Lemma 1, this step can be realized with $O(\log(nd))$ circuit depth and $O(nd)$ ancillary qubits. The remaining procedure is to uncompute register A. To do so, we introduce another n-index, $\tilde{n}$-word Boolean function f_{prep}. Let $\mathcal{Q}$ be a set containing all nonzero entries q_k of the target state. The definition of f_{prep} is that $f_{\text{prep}}(q_k)=k$ for $q_k\in\mathcal{Q}$ and $f_{\text{prep}}(q)=0$ for $q\notin\mathcal{Q}$ (i.e., $\mathcal{S}_{f_{\text{prep}}}=\mathcal{Q}$). We query select$(f_{\text{prep}})$ with register B as index register and register A as word register, after which the state becomes $\sum_{k=0}^{d-1}\psi_k|0\ldots0\rangle_A|q_k\rangle_B$. The target state is obtained by tracing out register A.

Because f_{prep} is d sparse, according to Lemma 2, this step has circuit depth $O(\log(nd))$ and space complexity $O(nd\log d)$. So the total circuit depth and space complexity of sparse state preparation is $\Theta(\log(nd))$ and $O(nd\log d)$. Moreover, as mentioned previously, it takes classical runtime $O(nd)$ to generate the quantum circuit, which can be reduced to $O(\log(nd))$ for parallel calculation with $O(nd)$ space complexity.

Theorem 2 also provides a method for approximating nonsparse states with sparse states. We denote $a_j^{\max}$ as the jth largest value of $|a_k|$. Suppose $\sum_{j=1}^{d}|a_j^{\max}|^2=1-\varepsilon$, we can then set all amplitudes $|a_k|<a_d^{\max}$ to zero and normalize the sparse state. According to Theorem 2, the quantum state can be approximated to fidelity $F=1-\varepsilon$ with circuit depth $O(\log(nd))$ (see Sec. IX of [35] for details).

We next discuss applications of our results.

Hamiltonian simulation.—A Hamiltonian $\hat{H}$ can generally be expressed as a linear combination of products of single qubit unitaries (such as Pauli strings)

$$\hat{H}=\sum_{p=0}^{P-1}\alpha_p\hat{V}(p), \tag{6}$$

for some $\alpha_p>0$, $\hat{V}(p)=\otimes_{l=1}^{n}\hat{V}_l(p)$ and $\hat{V}_l(p)\in\mathrm{SU}(2)$. Simulation of e^{-iHt} with optimal complexity with respect to the accuracy can be achieved with block encodings [5,48]. We say that $\hat{U}$ is a block encoding [48] of $\hat{H}$ if $(\langle 0|^{\otimes a}\otimes\mathbb{I}_n)\hat{U}(|0\rangle^{\otimes a}\otimes\mathbb{I}_n)\propto\hat{H}$ for some integer a. One common construction way of block-encoding is based on linear combination of unitaries [5]. We define $\hat{G}$ as a quantum state preparation operator satisfying $\hat{G}|0\rangle=|G\rangle\equiv\sum_p\sqrt{\alpha_p/\alpha}|p\rangle$ with $\alpha\equiv\sum_p\alpha_p$. It can be verified that $(\hat{G}^\dagger\otimes\mathbb{I}_n)\text{select}(\hat{V})(\hat{G}\otimes\mathbb{I}_n)$ is a block encoding of $\hat{H}$. Conventional ways to implement select$(\hat{V})$ and block encoding requires a circuit depth $O(nP)$ [4].

In contrast, according to Theorem 1 and Lemma 1, $\hat{G},\hat{G}^\dagger$ can be realized with circuit depth $O(\log P)$, and select$(\hat{V})$ can be realized with circuit depth $O(\log(Pn))$. So the block encoding can be constructed with circuit depth $O(\log(nP))$. Combining qubitization [5] with our fast construction of block encoding, we have the following result (see Sec. IV of [35] for details), which reduce the circuit depth exponentially to respect to nP.

Theorem 3: (Hamiltonian simulation by qubitization) Let $\hat{H}$ be a Hermitian operator expressed as Eq. (6). Using only single- and two-qubit gates, the evolution $e^{-i\hat{H}t}$ can be simulated to precision ε with circuit depth $O(\log(nP)(\alpha t+\log(1/\varepsilon)))$ and $O(nP)$ qubits.

We note that another version of parallel Hamiltonian simulation has been proposed in Ref. [20], achieving doubly logarithmic circuit depth with respect to the precision $O(\log^3\log(1/\varepsilon))$. The algorithm is based on a state preparation method with cubic circuit depth. In Sec. V of [35], we show that the circuit depth can be further reduced to $O(\log^2\log(1/\varepsilon))$ based on Theorem 1.

Solving linear systems.—Given an invertible matrix $H\in R^{2^n\times 2^n}$ and vector $b\in R^{2^n}$, quantum algorithms of linear systems aim at generating an approximation of quantum state $|x\rangle$ proportional to $H^{-1}\cdot b$. For sparse H, $|x\rangle$ can be obtained with a circuit depth $O(\text{poly}(n))$ [1,48,49]. The results has also been generalized to nonsparse cases [50]. However, these quantum algorithms assume the query of quantum state preparation and Hamiltonian simulation oracles. In general, to achieve poly-logarithmic circuit depth, data structure with space complexity $O(2^n)$ is required, leaving room for classical dequantization algorithms [24–26]. Specifically, based on an analog data structure with $O[\text{nnz}(H)n]$ space complexity [nnz$(\cdot)$ refers to the number of nonzero entries], classical dequantization algorithms [25,26] can sample from the distribution of the measurement outcomes of $|x\rangle$ with a circuit depth $O(\text{poly}(n))$.

Therefore, whether we could more efficiently solve linear systems heavily relies on efficiency of the quantum state preparation and Hamiltonian simulation oracles. Here, considering sparse matrices of Eq. (6), we show an exponential advantage of quantum computing algorithms based on Theorems 2,3.

Theorem 4: (Solving linear system) Let $\hat{H}$ be a Hermitian expressed as Eq. (6) with condition number κ. Let $|b\rangle$ be a $O(\text{poly}(n))$-sparse quantum state. With only single- and two-qubit gates, the quantum state $|x\rangle$ proportional to $H^{-1}|b\rangle$ can be approximately prepared to precision ε using $\tilde{O}(\text{poly}(\log(nP),\alpha,\kappa))$ circuit depth and $O(\text{poly}(n,P))$ qubits, where $\tilde{O}$ neglects the logarithmic dependence on κ, $1/\varepsilon$.

With $P,\alpha=O(\text{poly}(n))$, our method has $\tilde{O}(\text{poly}(n,\kappa))$ circuit depth and $O(\text{poly}(n))$ space complexity. In comparison, the data structures by classical dequantization algorithms [25,26] have $O(\text{nnz}(\hat{H})\log N)=O(NPn)$ space complexities, which is exponentially larger. Furthermore, assuming that $P,\alpha=O(1)$ and $|b\rangle$ is $O(1)$-sparse, the circuit depth of our parallel method is

further reduced to $\tilde{O}(\log(n)\mathrm{poly}(\kappa))$, which is also exponentially lower than the depth $\tilde{O}(n\mathrm{poly}(\kappa))$ of the quantum algorithms using sequential select$(\hat{H})$ [4,51]; Yet, both of them have $O(n)$ space complexity. More details are provided in Sec. VI of [35].

QRAMs.—At last, we show applications for QRAMs. Given a binary dataset $\mathcal{D} \equiv \{D_k\}_{k=0}^{2^n-1} \in \{0,1\}^n$, QRAMs are memory architectures enabling the transformation

$$\mathrm{QRAM}(\mathcal{D}) \sum_{k=0}^{2^n-1} \psi_k |k\rangle|0\rangle = \sum_{k=0}^{2^n-1} \psi_k |k\rangle|D_k\rangle. \tag{7}$$

Efficient implementation of QRAM is important for many applications, especially for quantum machine learning [52]. Conventional methods have $O(n)$ circuit depth using $O(N)$ ancillary qubits [46,47,53,54]. If $\mathcal{D}$ is d sparse (with at most d nonzero elements), the space complexity can be significantly reduced using quantum read-only access memory (QROM) [55]. Alternatively, Eq. (7) can be realized by performing n-Toffoli gates for d times. However, these methods have circuit depth linear with d, which is not yet optimal.

On the other hand, by defining $f_{\mathrm{qram}}(k) \equiv D_k$, Eq. (7) can be satisfied by select (f_{qram}). According to Lemma 2, we can obtain the following result.

Theorem 5: (Sparse QRAM) With only single- and two-qubit gates, arbitrary QRAM$(\mathcal{D})$ in Eq. (7) with d-sparse $\mathcal{D}$ can be implemented with circuit depth $O(\log(nd))$ and $O(nd)$ ancillary qubits.

Our protocol thus has an exponentially lower circuit depth compared to existing ones [46,47,53–55], while the space complexity remains polynomial.

Moreover, we can construct a nonsparse QRAM for continuous amplitude, i.e., $|D_k\rangle \in \mathbb{C}^2$, based on Lemma 1 (see Sec. VII of [35]). Our method requires $O(n)$ circuit depth and $O(N)$ ancillary qubits. The connectivity is identical to those for binary QRAMs [46,47,53,54], which is more sparse than other continuous QRAM schemes developed recently [22,23] assuming all-to-all connectivity.

Discussions.—We have achieved optimal circuit depth for general and sparse quantum state preparation. While Theorems 1,2 assume that only single- and two-qubit gates are allowed, our results can be generalized to constant weight operations, i.e., operations applied to a constant number of qubits. It therefore represents a fundamental limit of quantum information processing imposed by constant-weight operations. Future direction includes finding optimal space-time trade-offs for sparse state preparation, and exploring more practical applications.

We thank Bujiao Wu for insightful discussions. This work is supported by the National Natural Science Foundation of China Grant No. 12175003, and Emerging Engineering Interdisciplinary-Young Scholars Project, Peking University, the Fundamental Research Funds for the Central Universities.

*xiaoyuan@pku.edu.cn

[1] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, Phys. Rev. Lett. **103**, 150502 (2009).
[2] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum principal component analysis, Nat. Phys. **10**, 631 (2014).
[3] I. Kerenidis and A. Prakash, Quantum recommendation systems, arXiv:1603.08675.
[4] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, Proc. Natl. Acad. Sci. U.S.A. **115**, 9456 (2018).
[5] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, Quantum **3**, 163 (2019).
[6] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, Phys. Rev. A **52**, 3457 (1995).
[7] E. Knill, Approximation by quantum circuits, arXiv:quant-ph/9508006.
[8] J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, Efficient Decomposition of Quantum Gates, Phys. Rev. Lett. **92**, 177902 (2004).
[9] V. V. Shende, I. L. Markov, and S. S. Bullock, Minimal universal two-qubit controlled-not-based circuits, Phys. Rev. A **69**, 062321 (2004).
[10] V. Shende, S. Bullock, and I. Markov, Synthesis of quantum-logic circuits, IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. **25**, 1000 (2006).
[11] G.-L. Long and Y. Sun, Efficient scheme for initializing a quantum register with an arbitrary superposed state, Phys. Rev. A **64**, 014303 (2001).
[12] L. Grover and T. Rudolph, Creating superpositions that correspond to efficiently integrable probability distributions, arXiv:quant-ph/0208112.
[13] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, Transformation of quantum states using uniformly controlled rotations, Quantum Inf. Comput. **5**, 467 (2005).
[14] V. Bergholm, J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, Quantum circuits with uniformly controlled one-qubit gates, Phys. Rev. A **71**, 052330 (2005).
[15] M. Plesch and Č. Brukner, Quantum-state preparation with universal gate decompositions, Phys. Rev. A **83**, 032302 (2011).
[16] R. Iten, R. Colbeck, I. Kukuljan, J. Home, and M. Christandl, Quantum circuits for isometries, Phys. Rev. A **93**, 032318 (2016).
[17] X. Sun, G. Tian, S. Yang, P. Yuan, and S. Zhang, Asymptotically optimal circuit depth for quantum state preparation and general unitary synthesis, arXiv:2108.06150v2.
[18] G. H. Low, V. Kliuchnikov, and L. Schaeffer, Trading t-gates for dirty qubits in state preparation and unitary synthesis, arXiv:1812.00954.
[19] X.-M. Zhang, M.-H. Yung, and X. Yuan, Low-depth quantum state preparation, Phys. Rev. Res. **3**, 043200 (2021).
[20] Z. Zhang, Q. Wang, and M. Ying, Parallel quantum algorithm for Hamiltonian simulation, arXiv:2105.11889.
[21] G. Rosenthal, Query and depth upper bounds for quantum unitaries via Grover search, arXiv:2111.07992.

[22] P. Yuan and S. Zhang, Optimal qram and improved unitary synthesis by quantum circuits with any number of ancillary qubits, arXiv:2202.11302.

[23] B. D. Clader, A. M. Dalzell, N. Stamatopoulos, G. Salton, M. Berta, and W. J. Zeng, Quantum resources required to block-encode a matrix of classical data, arXiv:2206.03505.

[24] E. Tang, A quantum-inspired classical algorithm for recommendation systems, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (ACM Press, New York, 2019).

[25] N.-H. Chia, A. Gilyén, T. Li, H.-H. Lin, E. Tang, and C. Wang, Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning, in *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing* (ACM Press, New York, 2020).

[26] A. Gilyén, Z. Song, and E. Tang, An improved quantum-inspired algorithm for linear regression, Quantum **6**, 754 (2022).

[27] N.-H. Chia, T. Li, H.-H. Lin, and C. Wang, Quantum-inspired sublinear algorithm for solving low-rank semidefinite programming, in *45th International Symposium on Mathematical Foundations of Computer Science* (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, 2020).

[28] E. Malvetti, R. Iten, and R. Colbeck, Quantum circuits for sparse isometries, Quantum **5**, 412 (2021).

[29] T. M. de Veras, L. D. da Silva, and A. J. da Silva, Double sparse quantum state preparation, Quantum Inf. Process. **21**, 204 (2022).

[30] N. Gleinig and T. Hoefler, An efficient algorithm for sparse quantum state preparation, in *Proceedings of the 58th ACM/IEEE Design Automation Conference* (IEEE, New York, 2021).

[31] T. M. Veras, I. C. De Araujo, K. D. Park, and A. J. Dasilva, Circuit-based quantum random access memory for classical data with continuous amplitudes, IEEE Trans. Comput. **70**, 2125 (2021).

[32] G. Marin-Sanchez, J. Gonzalez-Conde, and M. Sanz, Quantum algorithms for approximate function loading, arXiv:2111.07933.

[33] I. F. Araujo, C. Blank, and A. J. da Silva, Entanglement as a complexity measure for quantum state preparation, arXiv:2111.03132.

[34] A. G. Rattew and B. Koczor, Preparing arbitrary continuous functions in quantum registers with logarithmic complexity, arXiv:2205.00519.

[35] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.129.230504 for details, which includes Refs. [36–39].

[36] D. K. Park, F. Petruccione, and J.-K. K. Rhee, Circuit-based quantum random access memory for classical data, Sci. Rep. **9**, 1 (2019).

[37] P. Selinger, Efficient Clifford + T approximation of single-qubit operators, arXiv:1212.6253.

[38] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, Surface code quantum computing by lattice surgery, New J. Phys. **14**, 123011 (2012).

[39] D. Litinski, A game of surface codes: Large-scale quantum computing with lattice surgery, Quantum **3**, 128 (2019).

[40] K. S. Chou, J. Z. Blumoff, C. S. Wang, P. C. Reinhold, C. J. Axline, Y. Y. Gao, L. Frunzio, M. Devoret, L. Jiang, and R. Schoelkopf, Deterministic teleportation of a quantum gate between two logical qubits, Nature (London) **561**, 368 (2018).

[41] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A **86**, 032324 (2012).

[42] D. Gottesman and I. L. Chuang, Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations, Nature (London) **402**, 390 (1999).

[43] Y. Wan, D. Kienzler, S. D. Erickson, K. H. Mayer, T. R. Tan, J. J. Wu, H. M. Vasconcelos, S. Glancy, E. Knill, D. J. Wineland *et al.*, Quantum gate teleportation between separated qubits in a trapped-ion processor, Science **364**, 875 (2019).

[44] K. C. Chen, W. Dai, C. Errando-Herranz, S. Lloyd, and D. Englund, Scalable and high-fidelity quantum random access memory in spin-photon networks, PRX Quantum **2**, 030319 (2021).

[45] Q. Xu, A. Seif, H. Yan, N. Mannucci, B. O. Sane, R. Van Meter, A. N. Cleland, and L. Jiang, Distributed quantum error correction for chip-level catastrophic errors, arXiv:2203.16488.

[46] C. T. Hann, C.-L. Zou, Y. Zhang, Y. Chu, R. J. Schoelkopf, S. M. Girvin, and L. Jiang, Hardware-Efficient Quantum Random Access Memory with Hybrid Quantum Acoustic Systems, Phys. Rev. Lett. **123**, 250501 (2019).

[47] C. T. Hann, G. Lee, S. M. Girvin, and L. Jiang, Resilience of quantum random access memory to generic noise, PRX Quantum **2**, 020311 (2021).

[48] S. Chakraborty, A. Gilyén, and S. 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)* (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, 2019).

[49] A. M. Childs, R. Kothari, and R. D. Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, SIAM J. Comput. **46**, 1920 (2017).

[50] L. Wossnig, Z. Zhao, and A. Prakash, Quantum Linear System Algorithm for Dense Matrices, Phys. Rev. Lett. **120**, 050502 (2018).

[51] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian Dynamics with a Truncated Taylor Series, Phys. Rev. Lett. **114**, 090502 (2015).

[52] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, Nature (London) **549**, 195 (2017).

[53] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Random Access Memory, Phys. Rev. Lett. **100**, 160501 (2008).

[54] V. Giovannetti, S. Lloyd, and L. Maccone, Architectures for tpdela quantum random access memory, Phys. Rev. A **78**, 052310 (2008).

[55] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding Electronic Spectra in Quantum Circuits with Linear t Complexity, Phys. Rev. X **8**, 041015 (2018).