MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

Near-Lower-Bound Approximate Quantum State Preparation with Hardware-Efficient Circuits

Koike-Akino, Toshiaki

TR2026-131 September 17, 2026

Abstract

Preparing arbitrary quantum states is a fundamental primitive for quantum algorithms. Existing exact state-preparation algorithms require substantially more two-qubit gates than the information-theoretic lower bound. In this work, we investigate approximate quantum state preparation using variational hardwareefficient circuits. We optimize several entangling topologies for Haar-random target states and evaluate the fidelity as a function of the number of CNOT gates. Surprisingly, we observe a sharp transition from poor approximation to high-fidelity state preparation when the CNOT count approaches the lower bound. Moreover, a brickwork topology consistently outperforms chain and ring circuits under the same CNOT budget. We further show that a reduced 2-parameter local rotation performs almost identically to the 3-parameter full Euler rotation for brickwork circuits. These observations suggest that approximate variational state preparation can nearly attain the information-theoretic minimum entangling complexity while using hardware-efficient circuits.

IEEE International Conference on Quantum Computing and Engineering (QCE) 2026

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

Near-Lower-Bound Approximate Quantum State Preparation with Hardware-Efficient Circuits

Toshiaki Koike-Akino
Mitsubishi Electric Research Laboratories (MERL), 201 Broadway, Cambridge, MA 02139, USA
Email: koike@merl.com.

***Abstract*—Preparing arbitrary quantum states is a fundamental primitive for quantum algorithms. Existing exact state-preparation algorithms require substantially more two-qubit gates than the information-theoretic lower bound. In this work, we investigate approximate quantum state preparation using variational hardware-efficient circuits. We optimize several entangling topologies for Haar-random target states and evaluate the fidelity as a function of the number of CNOT gates. Surprisingly, we observe a sharp transition from poor approximation to high-fidelity state preparation when the CNOT count approaches the lower bound. Moreover, a brickwork topology consistently outperforms chain and ring circuits under the same CNOT budget. We further show that a reduced 2-parameter local rotation performs almost identically to the 3-parameter full Euler rotation for brickwork circuits. These observations suggest that approximate variational state preparation can nearly attain the information-theoretic minimum entangling complexity while using hardware-efficient circuits.**

***Index Terms*—Quantum state preparation, parameterized quantum circuits, variational algorithm, hardware-efficient ansatz.**

I. Introduction

QUANTUM state preparation is an essential building block for quantum simulation, quantum machine learning, and quantum algorithms. An exact synthesis method based on Möttönen decomposition [1] scales as $\mathcal{O}[2^{n+1}]$ CNOT gates, which is roughly 4-times more than the degree-of-freedom (DoF) lower bound $k_{\rm LB}$ of $\mathcal{O}[2^{n-1}]$ [2] for arbitrary pure-state preparation, where n is the number of qubits. Subsequent improvements by Plesch–Brukner (PB) [3] and, more recently, Li–Zhang–Zhang (LZZ) [4] significantly reduced to $\mathcal{O}[(23/24)2^n]$ and $\mathcal{O}[(11/12)2^n]$, respectively, but still remain nearly twice above the lower bound.

In contrast to exact synthesis, approximate state preparation has received considerably less attention. An open question is whether variational circuits can approach the fundamental lower bound on entangling complexity while maintaining high fidelity. Our main finding is that optimized circuits exhibit a sharp expressibility transition when the CNOT count approaches the DoF lower bound. We perform an empirical study of approximate Haar-random state preparation using hardware-efficient ansätze. Rather than proposing a new optimization algorithm, we study the intrinsic relationship between approximation fidelity and restricted CNOT budget. We demonstrate that a brickwork entangling topology achieves high fidelity at near-lower-bound CNOT counts when circuit parameters are optimized through a standard variational algorithm. In addition, the fidelity can be kept high even with reduced-parameter rotation gates from full Euler rotations, whereas it does not apply to chain/ring entangling topologies.

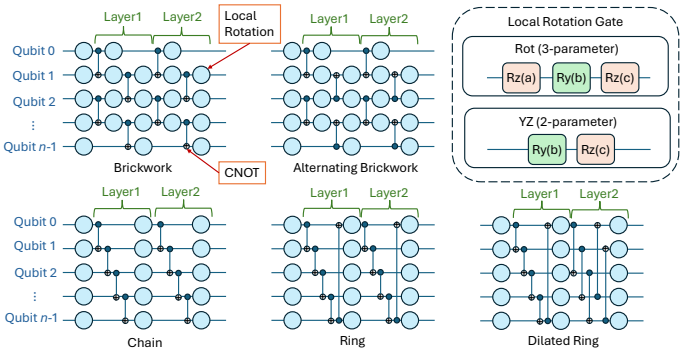

Fig. 1. Quantum ansatz variants.

Unlike the statistical expressibility analysis [5], which evaluates randomly sampled parameterized quantum circuits (PQCs), we study optimized approximate quantum state preparation. We directly measure the achievable fidelity for Haar-random target states under a constrained CNOT budget, revealing an empirical transition near the lower bound and providing a practical perspective on PQC expressibility. Our key observation is that optimized PQCs exhibit an empirical expressibility transition at around $k_{\rm LB}$, whereas existing methods remain around $2k_{\rm LB}$.

II. Approximate Quantum State Preparation

Target quantum states $|\phi\rangle$ are randomly sampled from the Haar distribution by generating complex Gaussian vectors followed by normalization. It has $2^{n+1}-2$ DoF, and its theoretical lower bound for the CNOT counts is given as [3]:

$$k_{\rm LB} = \left\lceil 2^{n-1} - (n+1)/2 \right\rceil . \tag{1}$$

A standard Möttönen decomposition [1] is roughly 4-times away from this lower bound, and recent work [3], [4] still need about 2-times more CNOTs. Although exact achievability remains unknown, we show that variational PQCs can achieve high fidelity at a near-lower-bound CNOT budget.

We consider hardware-efficient circuits consisting of repeated layers of parameterized single-qubit rotations followed by entangling CNOT layers. Fig. 1 shows several entangling topologies considered in this paper: brickwork, alternating brickwork, chain, ring, and dilated ring. Ring entangling requires n CNOTs per layer, whereas the remining topologies need $n-1$ CNOTs per layer. Local rotations are compared between 3-parameter Rot ($R_Z R_Y R_Z$) and 2-parameter YZ ($R_Y R_Z$) gates. Table I lists the required number of CNOTs and rotation gates for different ansatz topologies. The brickwork and chain entangling have slightly fewer CNOTs per layer than ring, whereas the

TABLE I
GATE COUNTS FOR QUANTUM ANSATZ WITH L ENTANGLING LAYERS

Ansatz	CNOTs	Rotations
Brickwork (standard/alternating)	$(n-1)L$	$n+2(n-1)L$
Chain	$(n-1)L$	$n(L+1)$
Ring (standard/dilated)	nL	$n(L+1)$

brickwork has roughly double rotation gates. Brickwork alternates even and odd CNOT layers, inserting a local rotation after each CNOT, thereby doubling the number of local parameters compared with chain and ring circuits. Nevertheless, brickwork and chain may be more hardware friendly as CNOTs are arranged adjacent qubits, whereas standard/dilated ring need conditional flip across far qubits.

Letting $U(\theta)$ be the ansatz operator given variational rotation angles θ, the quantum state output of the PQC is denoted as $|\psi\rangle = U(\theta)|0\rangle^{\otimes n}$. Circuit parameters θ are optimized using gradient-based optimization for each target state. Specifically, we aim at maximizing the fidelity $F = |\langle\phi|\psi\rangle|^2$, or minimizing the infidelity $1-F$. We use a standard optimization based on AdamW [6] with one-cycle learning rate scheduling [7].

III. NUMERICAL RESULTS

Fig. 2 shows the approximation infidelity versus normalized CNOT count relative to the lower bound k_{LB} at $n = 8$ qubits. Results report the median over multiple random target states and optimization initializations, at 10,000 epochs. All brickwork circuits exhibit a sharp transition near $k/k_{\mathrm{LB}} \approx 1$, where the infidelity rapidly decreases by several orders of magnitude. This suggests that approximate state preparation becomes practical as soon as the information-theoretic minimum entangling resources become available. In addition, brickwork consistently outperforms the other topologies. Although the ring circuits possess greater connectivity, they require substantially larger CNOT depth before another local rotation layer can be applied. Brickwork therefore achieves more frequent alternation between local rotations and entangling operations under the same CNOT budget. Interestingly, replacing the conventional 3-parameter Euler rotation by the reduced 2-parameter gate has negligible impact for brickwork circuits, while causing noticeable degradation for chain/ring architectures. This empirically supports the DoF argument underlying the state-preparation lower bound and indicates that brickwork efficiently utilizes local parameters. Fig. 3 shows the required number of CNOTs to achieve a high fidelity of $\geq 1-10^{-4}$ over n qubits. We observe that brickwork offers the small CNOT counts near the lower-bound k_{LB} in practice.

IV. CONCLUSION

We presented an empirical study of approximate Haar-random quantum state preparation using hardware-efficient variational circuits. We showed that i) approximate state preparation exhibits a sharp transition near the information-theoretic CNOT lower bound; ii) brickwork circuits provide the most efficient entangling topology among several baseline architectures; iii) 2-parameter local rotations are sufficient for brickwork circuits with no performance loss. These observations provide empirical evidence that approximate variational state preparation

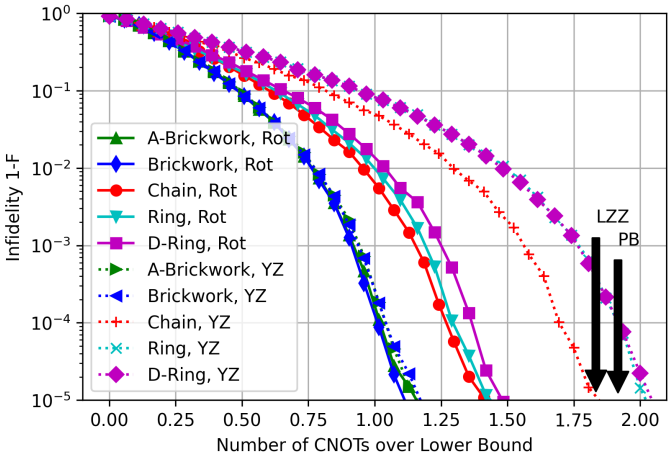

Fig. 2. Infidelity $1-F$ over the number of CNOTs relative to the theoretical lower bound k_{LB} at $n = 8$ qubits.

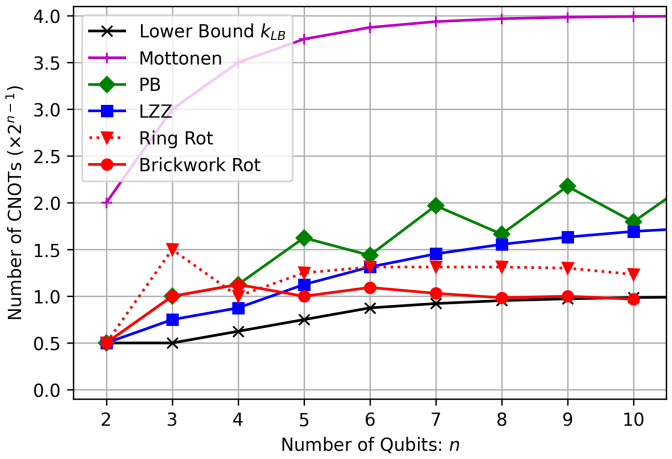

Fig. 3. Required number of CNOTs ($\times 2^{n-1}$) across n qubits for $F \geq 1-10^{-4}$.

can nearly attain the fundamental lower bound on entangling complexity, motivating future theoretical analysis and adaptive circuit construction methods.

REFERENCES

[1] M. Mottonen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, "Transformation of quantum states using uniformly controlled rotations," *arXiv preprint quant-ph/0407010*, 2004.

[2] V. V. Shende, S. S. Bullock, and I. L. Markov, "Synthesis of quantum logic circuits," in *Proceedings of the 2005 Asia and South Pacific Design Automation Conference*, 2005, pp. 272–275.

[3] M. Plesch and Č. Brukner, "Quantum-state preparation with universal gate decompositions," *Physical Review A—Atomic, Molecular, and Optical Physics*, vol. 83, no. 3, p. 032302, 2011.

[4] Z. Li, G. Zhang, and X.-M. Zhang, "Reducing C-NOT counts for state preparation and block encoding via diagonal matrix migration," *arXiv preprint arXiv:2603.16492*, 2026.

[5] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms," *Advanced Quantum Technologies*, vol. 2, no. 12, p. 1900070, 2019.

[6] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.

[7] L. N. Smith and N. Topin, "Super-convergence: Very fast training of neural networks using large learning rates," in *Artificial intelligence and machine learning for multi-domain operations applications*, vol. 11006. SPIE, 2019, pp. 369–386.