Hardware-Aware Quantum Circuit Synthesis

Nathan Jones
najones@clemson.edu
Clemson University
Clemson, SC, USA

Akhilesh Bondapalli
abondap@clemson.edu
Clemson University
Clemson, SC, USA

Toby Cox
tscox@clemson.edu
Clemson University
Clemson, SC, USA

Ian Lewis
iml@clemson.edu
Clemson University
Clemson, SC, USA

Rong Ge
rge@clemson.edu
Clemson University
Clemson, SC, USA

ABSTRACT

Effectively leveraging quantum computing requires generating and manipulating a desired quantum state using a quantum circuit. Quantum circuit synthesis (QCS) is bottlenecked by the exponential complexity of circuit verification via quantum simulation. Diffusion models are promising QCS candidates, because they circumvent quantum simulation during training. Existing diffusion-based QCS models demonstrate success for unconstrained circuits, but prove insufficient for producing hardware topology-constrained circuits–a common restriction for modern quantum machines. This work introduces a novel hardware-aware conditioning framework that enables topology-constrained QCS. Our approach delivers up to 8x higher success rate compared to the baseline for a state-of-the-art hardware-agnostic QCS model, proving the necessity for hardware-aware QCS.

CCS CONCEPTS

• **Hardware → Quantum error correction and fault tolerance**;
• **Computer systems organization → Quantum computing**.

KEYWORDS

Quantum Computing, Hardware-Aware Quantum Computing, Quantum Circuit Synthesis

ACM Reference Format:
Nathan Jones, Akhilesh Bondapalli, Toby Cox, Ian Lewis, and Rong Ge. 2018. Hardware-Aware Quantum Circuit Synthesis. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation emai (Conference acronym 'XX).* ACM, New York, NY, USA, 2 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Although quantum computing offers potential exponential speedups in fields including biology, chemistry, material science, and machine learning, a quantum algorithms broadly require users to manually build complex circuits to generate and manipulate quantum states, hampering accessibility. The task of calculating a quantum circuit to generate a quantum state is termed quantum circuit synthesis (QCS). QCS faces two main challenges: the vast search space of possible gate sequences and the exponential time and space required for quantum simulation. A promising direction in QCS leverages diffusion models due to their unique training procedure that avoids quantum simulation entirely [1–6].

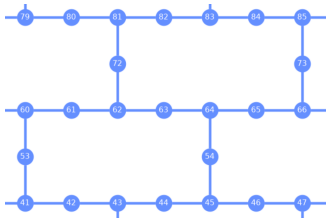

Figure 1: IBM Sherbrooke's hardware topology is far from fully connected: most qubits only connect to two neighbors.

However, existing generative QCS methodologies are hardware-agnostic. They operate under the assumption of an unconstrained qubit architecture with all-to-all connectivity, where a two-qubit gate can apply to any arbitrary pair of qubits. While these models are effective in this abstract setting, their output is incompatible with the physical constraints of contemporary quantum hardware, which often have restrictive topologies as shown in Figure 1. While an arbitrary quantum circuit can be transpiled to run on any hardware topology through the addition of SWAP gates, doing so significantly increases circuit depth and accumulates error.

Alternatively, hardware-aware synthesis promises shallower, higher-fidelity quantum circuits by avoiding transpilation entirely, but poses its own challenges. The number of topologies grows exponentially with the number of qubits, which is far from ideal for training a hardware-aware QCS model. Furthermore, hardware topologies may have large differences in circuit compatibility, which can introduce significant dataset and training imbalances.

This work introduces a novel and scalable conditioning pipeline that provides the benefits of hardware-aware QCS to any diffusion-based QCS approach. Our approach uses graph isomorphism and sub-topology search to eliminate irrelevant sub-topologies to avoid class explosion. We also employ cartesian-product-based dataset balancing to ensure equal training across classes. Our key contributions are as follows:

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

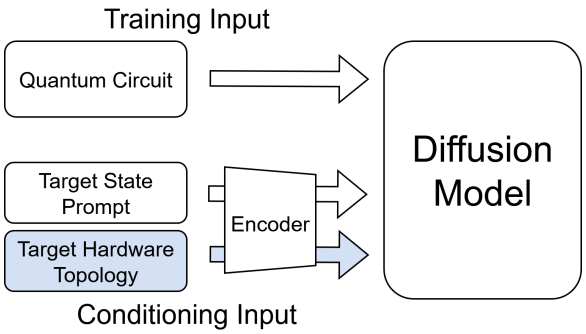

Figure 2: Our work applied to GenQC pipeline. We add the green shaded steps, the rest remains the same. For any hardware-unaware QCS diffusion model, our method unintrusively provides hardware awareness.

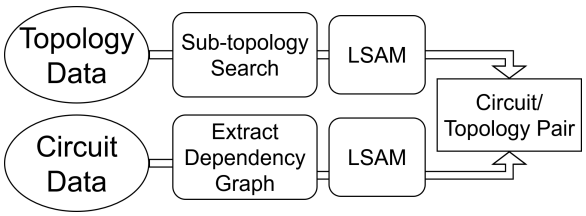

Figure 3: Data generation pipleine. We include only sub-topologies which occur in real semiconductor hardware. Representing both sub-topologies and circuit dependency graphs as LSAM allows us to remove isomorphic duplicate topologies and match circuits to a topology.

(1) We present a novel and scalable conditioning pipeline that augments existing QCS models to enable topology-constrained circuit generation
(2) We demonstrate two topology-encoding methods, analyze their performance, and provide insights on their trade-offs.

Our approach delivers significant improvement over hardware-agnostic approaches when constrained to realistic topologies, achieving up to 8x higher success rate than the original hardware-agnostic model. Our success demonstrates the necessity of hardware-aware conditioning.

2 METHODOLOGY

Our methodology addresses the two main challenges in topology conditioning: class explosion and dataset imbalance.

To ease the exponential growth of topologies with qubit count, we eliminate isomorphic topologies using the Lexigraphically Smallest Adjacency Matrix (LSAM) graph representation. LSAM is a canonical representation, meaning all isomorphic graphs are mapped to an identical representation. The LSAM representation reduces the set of labeled hardware topologies to the set of unlabeled topologies, a $N!$-scale reduction on N qubits.

Additionally, most graphs are not realistic hardware topologies. Quantum hardware schedulers allocate a connected, N-qubit subgraph of the full device hardware topology to a given quantum application. Therefore, we search the full-device hardware topologies of leading superconductor machines for a set of all non-isomorphic size N subgraphs. This set serves as our minimal dataset of realistic hardware topologies.

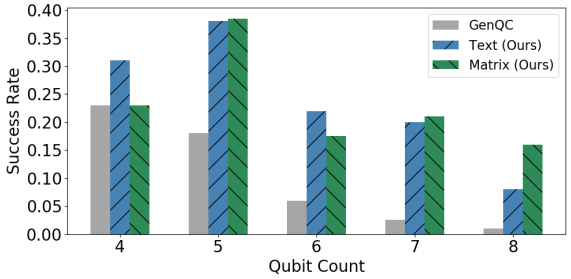

Figure 4: Accuracy of generated circuits constrained by hardware topology. Both topology-conditioned models outperform GenQC's masking approach.

We also analyze the effect of two embedding schemes for conditioning. In the first, we encode LSAM, written as a coordinate list, with a CLIP encoder. In the second, train a matrix encoder alongside the model. After embedding, conditioning inputs are incorporated as defined by the model.

3 RESULTS

We evaluate our approach using QCS model GenQC [4]. The hardware-agnostic GenQC model uses topology-based masking to generate topology-constrained circuits. Figure 4 shows that both of our topology-conditioning methods significantly improve GenQC. Despite minor decline for 6-8 qubits, hardware-aware conditioning prevents the exponential decline of hardware-agnostic GenQC. This indicates that our hardware-topology-conditioning improves scalability for higher qubit-counts.

4 CONCLUSION

This work proposes a novel conditioning framework to enable topology-constrained QCS in previously hardware-agnostic models. We evaluate two conditioning embeddings with a state-of-the-art hardware-agnostic model, and demonstrate that our approach significantly improves scalability and accuracy. Future work will further analyze the model architecture and tunable encoding techniques to extend this work to larger, more complex quantum states.

ACKNOWLEDGEMENT

This work was supported in part by the Clemson University Creative Inquiry + Undergraduate Research Program, the South Carolina Quantum Association, and U.S. National Science Foundation under Grant OAC-2518605.

REFERENCES

[1] Daniel Barta, Darya Martyniuk, Johannes Jung, and Adrian Paschke. 2025. Leveraging Diffusion Models for Parameterized Quantum Circuit Generation. http://arxiv.org/abs/2505.20863
[2] Collin Beaudoin and Swaroop Ghosh. 2025. Q-Fusion: Diffusing Quantum Circuits. http://arxiv.org/abs/2504.20794
[3] Zhiwei Chen and Hao Tang. 2025. UDiTQC: U-Net-Style Diffusion Transformer for Quantum Circuit Synthesis. http://arxiv.org/abs/2501.16380
[4] Florian Fürrutter, Gorka Muñoz-Gil, and Hans J Briegel. 2024. Quantum circuit synthesis with diffusion models. *Nature Machine Intelligence* 6, 5 (2024), 515–524.
[5] Florian Fürrutter, Zohim Chandani, Ikko Hamamura, Hans J. Briegel, and Gorka Muñoz-Gil. 2025. Synthesis of discrete-continuous quantum circuits with multimodal diffusion models. http://arxiv.org/abs/2506.01666
[6] Hang Zou, Martin Rahm, Anton Frisk Kockum, and Simon Olsson. 2025. Generative flow-based warm start of the variational quantum eigensolver. http://arxiv.org/abs/2507.01726