Quantum Approximate Optimization Algorithm on Different Qubit Systems

Seongmin Kim*, In-Saeng Suh*, Eduardo Antonio Coello Perez*, Ryan Landfield*, Michael Sandoval*,
Josh Cunningham*, Tom Beck*, Heidi Nelson-Quillin†, Erica Stump†
*Oak Ridge National Laboratory, Oak Ridge, TN 37830, USA
†IonQ Inc, 4505 Campus Dr, College Park, MD 20740, USA

Solving optimization problems is critical across many research domains, but the high dimensionality of parameter spaces often poses significant challenges. The Quantum Approximate Optimization Algorithm (QAOA) has emerged as a promising approach for accelerating optimization in the Noisy Intermediate-Scale Quantum (NISQ) era by leveraging both classical and quantum computational resources. However, its performance can vary depending on the underlying quantum hardware architecture. In this work, we evaluate the performance of QAOA on different quantum hardware platforms, specifically, superconducting transmon qubits and trapped-ion qubits, targetting real-world optimization problems formulated as fully connected Quadratic Unconstrained Binary Optimization (QUBO) instances. We evaluate both the solution quality and time-to-solution using dense QUBO matrices. Furthermore, we show that large-scale problems, such as a 100-bit QUBO instance, can be effectively tackled by integrating quantum computing with high-performance computing (HPC) resources. This study provides practical insights into the strengths and limitations of different qubit technologies and advances the application of quantum computing in solving real-world optimization problems.

***Index Terms*—Quantum approximate optimization algorithm, superconducting transmon qubit, trapped-ion qubit, quadratic unconstrained binary optimization, performance analysis**

I. Introduction

Quadratic Unconstrained Binary Optimization (QUBO) provides a framework for formulating many real-world combinatorial optimization problems. Due to their NP-hard nature, QUBO problems are often challenging for classical algorithms to solve [1], [2]. Quantum computing (QC) offers a potential advantage by enabling alternative approaches to explore large solution spaces effectively. In particular, hybrid quantum-classical algorithms, such as the Quantum Approximate Optimization Algorithm (QAOA), have attracted significant interest for tackling complex optimization problems [3], [4]. However, QAOA's performance can vary depending on the underlying quantum hardware architecture. Their differences, such as qubit connectivity, gate fidelity, the number of available qubits, and gate operation times, can significantly affect algorithmic performance [5].

In this study, we evaluate the performance of QAOA on two quantum hardware architectures: superconducting transmon qubits and trapped-ion qubits. We focus on solving fully connected QUBO problems derived from real-world applications, and compare solution quality and time-to-solution across different backends. Furthermore, we demonstrate the feasibility of solving large-scale QUBO problems by integrating quantum computing with HPC through problem decomposition.

II. Methodology

A. QUBO Generation: We generate QUBO matrices using machine learning techniques applied to materials science, specifically, designing planar multilayered structures [6], [7]. We employ a factorization machine to capture both linear and pairwise interactions among features, where the input features represent material structural configurations and the labels correspond to performance metrics. The resulting QUBO problems vary in size, ranging from 4 to 100 binary variables. These instances are solved using QAOA implemented on different quantum simulators and hardware backends.

B. QAOA: We construct QAOA circuits using Qiskit (v1.2.4). Quantum backends include Qiskit-Aer (v0.15.1) for noiseless simulation, Qiskit-IonQ (v0.5.8) for trapped-ion hardware and its simulator access, and Qiskit-IBM-Runtime (v0.32.0) for superconducting transmon qubit hardware.

C. Distributed QAOA (DQAOA): Current quantum hardware and simulators are limited in the number of qubits, making it infeasible to solve large-scale QUBO problems. To address this limitation, we adopt a distributed QAOA framework, which decomposes a large QUBO problem into smaller sub-problems that are solved independently [8]. Here, we validate the distributed approach on quantum simulators using a single-core setup to establish baseline performance.

III. Results

The QUBO problems used in our experiments, representing real-world optimization, are fully connected (Fig. 1), making them particularly challenging to solve. We first analyze the QAOA's performance on quantum simulators. A simulator emulating trapped-ion devices exhibits higher approximation ratios, likely due to their use of arbitrary two-qubit gate operations [9]. However, this simulator generally requires longer execution times, especially for larger problems (Fig.

2). This trend is also observed with our evaluation on actual quantum hardware. A device based on superconducting transmon qubits offers faster execution but lower approximation accuracies, while trapped-ion hardware provides better solution quality with longer runtime (Fig. 3). This performance trade-off arises from the architectural characteristics: trapped-ion systems feature fully connected qubit topologies, which reduce or eliminate the need for SWAP gates, but their longer gate durations impact overall speed compared to superconducting architectures [5], [10]. However, the approximation ratio achieved with superconducting qubits is expected to be significantly improved by employing the DQAOA framework (Fig. 3) [8].

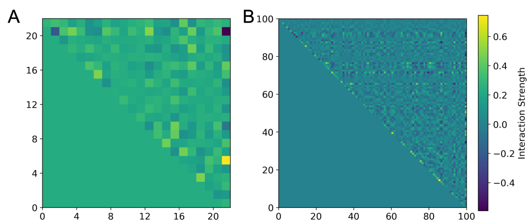

Fig. 1: QUBO problems representing real-world optimization tasks. Problem size: (A) 22, and (B) 100.

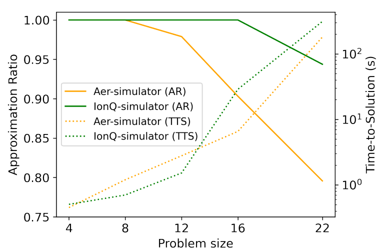

Fig. 2: Approximation ratio and time-to-solution of QAOA on different quantum simulators.

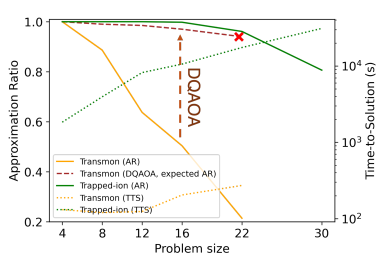

Fig. 3: Approximation ratio and time-to-solution of QAOA on different quantum hardware architecture. The red cross point data is adopted from [8].

To tackle large-scale dense QUBO problems (e.g., a 100-bit problem), which remain infeasible for current quantum devices, we apply DQAOA. As illustrated in Fig. 4, increasing the number of DQAOA iterations improves the approximation ratio. Although this leads to longer runtimes in a single-node setup, the total time-to-solution can be significantly reduced by parallelizing sub-problem execution across multiple HPC nodes equipped with quantum resources [8].

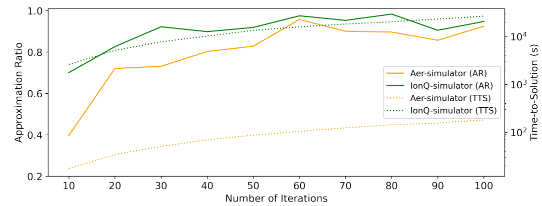

Fig. 4: Approximation ratio and time-to-solution of DQAOA on quantum simulators.

IV. Conclusion

Our results highlight the performance characteristics of QAOA across different quantum hardware platforms: superconducting transmon and trapped-ion qubits. While superconducting qubit devices benefit from faster gate operations, trapped-ion systems offer higher accuracy due to their fully connected architecture. These hardware-specific limitations can be mitigated through the use of DQAOA. Additionally, we show the potential to address large-scale optimization problems using a distributed QAOA framework. These findings offer practical insights into the suitability of different quantum hardware for real-world optimization problems.

V. Acknowledgments

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Science Center.

References

[1] S. Kim *et. al.*, Quantum annealing for combinatorial optimization: a benchmarking study. *npj Quantum Information*, 11:77, 2025.

[2] S. Kim *et. al.*, Distributed Variational Quantum Algorithm with Many-qubit for Optimization Challenges. arXiv:2503.00221, 2025.

[3] D. Lykov *et. al.*, Sampling frequency thresholds for the quantum advantage of the quantum approximate optimization algorithm *npj Quantum Information*, 9(1):73, 2023.

[4] R. Shaydulin *et. al.*, Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem. *Science Advances*, 10(22):eadm6761, 2024.

[5] S. Bilnov *et. al.*, Comparison of cloud-based ion trap and superconducting quantum computer architectures. *AVS Quantum Science*, 3:033801, 2021.

[6] S. Kim *et. al.*, High-Performance Transparent Radiative Cooler Designed by Quantum Computing. *ACS Energy Letters*, 7(12):4134-4141, 2022.

[7] S. Kim *et. al.*, Wide-angle spectral filter for energy-saving windows designed by quantum annealing-enhanced active learning. *Cell Reports Physical Science*, 5(3):101847, 2024.

[8] S. Kim *et. al.*, Distributed quantum approximate optimization algorithm on integrated high-performance computing and quantum computing systems for large-scale optimization. arXiv:2407.20212, 2024.

[9] W. Wu *et. al.*, Quantum computing and simulation with trapped ions: On the path to the future. *Fundamental Research*, 1(2):213-216, 2021.

[10] C.W. Hogle *et. al.*, High-fidelity trapped-ion qubit operations with scalable photonic modulators. *npj Quantum Information*, 9:74, 2023.