

POSTER TITLE

Quantum Simulators and Applications on
Quantum Framework

POSTER AUTHORS

Srikar Chundury	schundu3@ncsu.edu
Zhihao Xu	z xu8@nd.edu
Amir Shehata	shehataa@ornl.gov
Seongmin Kim	kims@ornl.gov
Frank Mueller	fmuelle@ncsu.edu
In-Saeng Suh	suhi@ornl.gov

POSTER ABSTRACT

Simulating quantum circuits is essential for validating quantum algorithms. However, no single simulator consistently performs best — efficiency depends on circuit structure, entanglement, and depth. In this work, we integrate Qiskit-Aer (state-vector and matrix product state) and QTensor, a tree-tensor-network based simulator, into the Quantum Framework (QFw), a modular platform that supports multiple quantum backends via a unified interface. We also enable distributed quantum approximate optimization algorithm (DQAOA) application compatibility with QFw, allowing sub-problems to be solved in parallel at scale. We then benchmark DQAOA and TFIM (transverse field Ising model) circuits across supported simulators, showing how performance varies significantly with problem type. All simulations are deployed on the Frontier supercomputer using QFw’s MPI-based orchestration for distributed, multi-node execution. These results underscore the need for simulator-agnostic infrastructure to enable systematic evaluation and high-performance scaling of quantum workloads. QFw provides a practical and extensible path toward reproducible quantum algorithm development across diverse application domains.

POSTER RELEVANCE

- Quantum Software Engineering
- Quantum Computing

4), leveraging QFw’s extensible design to manage backend-specific configuration, execution, and result retrieval while remaining transparent to frontend application code. To support scalable quantum optimization, we enabled DQAOA application compatibility with QFw. Our workflow (Fig. 1) begins with a dual-group heterogeneous SLURM allocation (step-1), one group for DQAOA and another for QFw. In hetgroup-1, QFw creates a unique PRTE-DVM URI and starts QPM services (step-2). In hetgroup-0, the Qiskit-Python-based DQAOA application starts, taking the quadratic unconstrained binary optimization (QUBO) matrix as input (step-5). While this happens, the QFwBackend establishes secure connections to the QPM services (step-4). The large QUBO is then decomposed into multiple subQUBOs using random or impact factor directed (IFD) decomposition methods (step-6), and Python’s threading package issues concurrent subQUBO solve calls. We opted for threading because our workload is I/O bound, as each call makes an asynchronous RPC to QFw (step-7) to trigger circuit execution. These RPC calls involve network communication to QPM services, which submit jobs to PRTE (step-8) and launch MPI-based circuit executions (step-9) across allocated HPC cores and nodes. Sub-solution results are collected and aggregated (step-10) before the optimization cycle repeats. After the application converges, the QPM services are shut down, the DVM URI is deleted and SLURM resources are deallocated (step-13). This design enables asynchronous backend invocation using HPC parallelism, solving DQAOA efficiently at scale.

C. Tests and Results: We benchmarked multiple simulators (Qiskit-Aer, NWQ-Sim [7], TN-QVM [8], and QTensor) for TFIM, QAOA, and DQAOA circuits. Figure 2a shows TFIM scaling simulation performance, where Qiskit-Aer’s matrix product state (MPS) method efficiently solves up to 33-qubit TFIM circuits, while other simulators do not scale to such problem sizes. We observe that NWQ-Sim also solves TFIM-32, although with higher runtime compared to Qiskit-Aer’s MPS. For QAOA benchmarks, Figure 2b presents scaled optimization runs. Some data is missing because simulators took longer than two hours to solve larger QUBO problems, when execution was cut off.

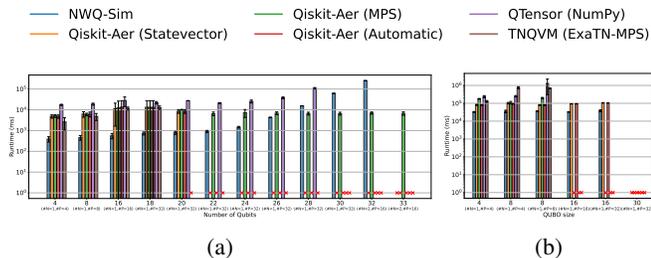


Fig. 2: Simulation performance results: (a) TFIM, and (b) QAOA. Above, #N is number of nodes, and #P is number of processes per node.

As seen in Figures 3a and 3b, DQAOA can benefit from scaling each circuit execution within subQAOAs using QFw. This provides developers with more detail about the behavior of such an algorithm at scale. We also noted that with smaller

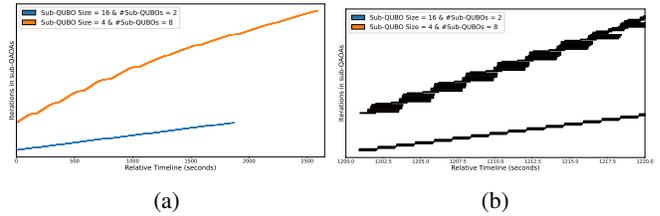


Fig. 3: DQAOA timing analysis with NWQ-Sim backend: (a) full timing results for various subQUBO configurations, and (b) zoomed view showing concurrent subQAOAs.

subQUBOs, each subQAOA took more iterations for completion compared to larger subQUBOs. These results demonstrate that simulator performance varies significantly based on circuit type, entanglement, and problem size, underscoring the importance of simulator-agnostic frameworks such as QFw for scalable and reproducible quantum algorithm development on HPC systems.

E. Future Work: QFw could benefit from automatic QPM selection based on circuit metadata to select optimal simulators and execute them at scale. Also, we intend to explore hybrid HPC/QC execution of DQAOA, where subQUBOs are dynamically mapped across simulators and actual hardware (as shown via dotted lines in Figure 1) for performance portable quantum workflows.

Acknowledgment: This work was supported in part by NSF awards PHY-1818914, MPS-2120757, CCF-2217020, CISE-2316201, and PHY-2325080 as well as DOE DE-SC0025384. 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.

REFERENCES

- [1] A. Shehata, T. Naughton, and I.-S. Suh, “A framework for integrating quantum simulation and high performance computing,” 2024. [Online]. Available: <https://arxiv.org/abs/2408.08098>
- [2] S. Chundury, A. Shehata, T. Naughton III, S. Kim, F. Mueller, and I.-S. Suh, “QFw: A quantum framework for large-scale hpc ecosystems.” Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States), 11 2024. [Online]. Available: <https://www.osti.gov/biblio/2498439>
- [3] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta, “Quantum computing with qiskit,” 2024. [Online]. Available: <https://doi.org/10.48550/arXiv.2405.08810>
- [4] D. Lykov, A. Chen, H. Chen, K. Keipert, Z. Zhang, T. Gibbs, and Y. Alexeev, “Performance evaluation and acceleration of the qtensor quantum circuit simulator on gpu,” 2021. [Online]. Available: <http://dx.doi.org/10.1109/QCS54837.2021.00007>
- [5] Z. Xu, S. Chundury, S. Kim, A. Shehata, X. Li, A. Li, T. Luo, F. Mueller, and I.-S. Suh, “GPU-Accelerated Distributed QAOA on Large-scale HPC Ecosystems,” 6 2025. [Online]. Available: <https://arxiv.org/abs/2506.10531>
- [6] S. Kim, T. Luo, E. Lee, and I.-S. Suh, “Distributed quantum approximate optimization algorithm on integrated high-performance computing and quantum computing systems for large-scale optimization,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.20212>
- [7] A. Li, B. Fang, C. E. Granade, G. Prawiroatmodjo, B. Heim, M. Rötteler, and S. Krishnamoorthy, “rm SV-Sim: Scalable pgas-based state vector simulation of quantum circuits,” 2021.
- [8] A. J. McCaskey, “Tensor Network Quantum Virtual Machine (TNQVM),” 11 2016. [Online]. Available: <https://www.osti.gov/servlets/purl/1340180>

Motivation

- Quantum circuit simulation is essential for validating quantum algorithms
- No single simulator is universally optimal
 - Performance depends on problem structure, entanglement, circuit depth, number of connected components, and maybe more (?)
- Systematic benchmarking is required to identify the most efficient simulator for each workload/circuit
- Therefore, scalable, simulator-agnostic frameworks are needed to identify quantum advantage candidates

Background

Quantum Framework (QFw) [1-3]

- QFw is a modular HPC quantum framework developed at ORNL to integrate quantum simulators and hardware in a unified API → see component workflow in Fig. 1
- Supports backend-agnostic quantum application development
- Built atop PRTE-DVM to allow scalable multinode deployment with fault tolerance and rapid job spawning

Application Frontend Invocation

Users write quantum applications by using frontends such as Qiskit or Pennylane. These applications construct quantum circuits, optimization problems, or time evolution workflows that need to be simulated or executed efficiently on HPC resources.

Routing via QFwBackend (Python Layer)

The QFwBackend is a lightweight Python interface between user frontends and QFw. It translates standard circuit definitions and execution calls into QFw API requests, thereby enabling backend-agnostic operation without requiring application code changes.

Platform Management with QPM (Quantum Platform Manager)

The QPM allocates computational resources across available nodes. This decouples the choice of simulator from the application logic to facilitate seamless benchmarking and deployment.

Scheduling and Control by QRC (Quantum Resource Controller)

The QRC orchestrates task scheduling across nodes and processes to ensure efficient use of HPC resources. It also coordinates execution order and monitors job status.

Simulator Execution

The selected backend simulators (e.g., NWQ-Sim, QTensor) execute the quantum circuits. Each simulator has different performance characteristics depending on circuit structure, entanglement, and depth.

Result Aggregation and Return

Finally, QFw seamlessly aggregates simulation results and returns them to the application frontend. This allows users to retrieve expectation values, measurement samples, or final state vectors for further processing without managing backend-specific data handling.

Fig. 1. QFw components.

- QFw uses the Distributed Execution Framework (DEFw) to enable scalable, low-latency RPC communication between QPM, QRC, and backend simulators
- DEFw → Infrastructure that abstracts inter-component communication via PRTE-DVM without being part of the application or QPM layers directly

Quantum Algorithms (Relevant for this Work)

- TFIM (Transverse-Field Ising Model)
 - Models quantum spin systems with both interaction (Ising) and transverse field terms
 - Used for benchmarking simulator performance across varying problem sizes and connectivity types
- QAOA (Quantum Approximate Optimization Algorithm) [4]
 - A gate-based variational quantum algorithm designed to solve combinatorial optimization problems such as QUBO
 - Used here for metamaterial design optimization by leveraging hybrid HPC-QC workflows
- DQAOA (Distributed QAOA) [1, 9]
 - Extends QAOA by partitioning large problems into subproblems (subQUBOs) that can be solved in parallel across HPC nodes, thereby enabling scalability beyond single simulator memory limits

This Work

- Integrated Qiskit-Aer [7] QPM into QFw
 - Methods → state-vector, matrix-product-state, and automatic
- Integrated QTensor [8] QPM into QFw
 - Support for tree-tensor-network (TTN) based simulation to efficiently handle circuits with low entanglement and high depth
- DQAOA [9] application compatibility with QFw
 - Allows partitioned subQUBOs to be executed in parallel by using QFw's PRTE
- Evaluated circuit simulation performance characteristics
 - Benchmarked multiple simulators via QFw's QPMs-TN-QVM [6], NWQ-Sim [5], Qiskit-Aer, and QTensor
 - TFIM evolution (up to 33 qubits)
 - DQAOA optimization (QUBO size 30)
- Performed scaling tests on Frontier supercomputer with larger subQUBOs for DQAOA without requiring any additional application-level code modifications
 - Integrates seamlessly with SLURM job workflows → Fig. 2
 - QFw abstracts resource management and backend execution

- Ran DQAOA-30 via QFw + NWQ-Sim for different configurations of subQUBOs for 10 cycles using the IFD decomposition method
 - SLURM allocation of 2 Frontier nodes (1 for DQAOA and 1 for QFw). Each process is tied to 1 core via PRTE
 - QFw enables easy scaling → Fig. 5
 - Each circuit execution automatically uses MPI
 - At any point in the execution, #subQUBOs would be running concurrently → Fig. 6

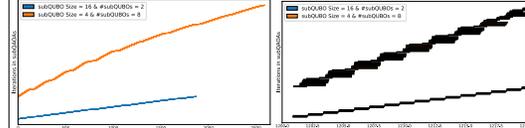


Fig. 5. DQAOA with QFw+NWQ-Sim.

Fig. 6. Region of interest from Fig. 5.

- DQAOA solved by using NWQ-Sim demonstrates—at scale—the performance behavior of a distributed quantum algorithm
 - #Iterations of smaller subQAOAs >> larger subQAOAs
 - COBYLA optimizer is thread-unsafe
 - Running it via mpi4py would solve this problem but remains as future work
 - Used SPSA optimizer for concurrent subQAOA execution across threads

DQAOA + QFw + SLURM

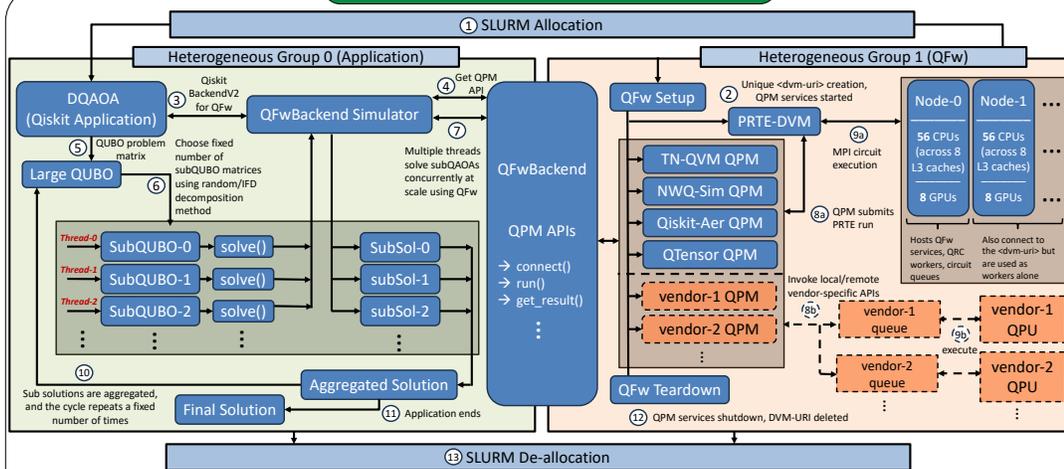


Fig. 2. Solving DQAOA using QFw.

- In Fig. 2, steps ① to ⑬ inform the workflow of how DQAOA is run using QFw
 - Green → Hetgroup 0, Orange → Hetgroup 1
 - Each parameterized ansatz from each subQAOA runs at scale using MPI via QFw (PRTE DVM-URI)

Preliminary Results

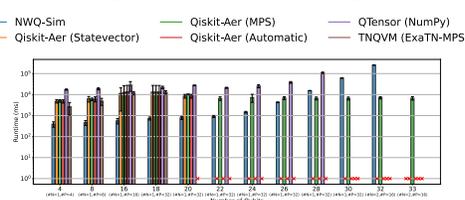


Fig. 3. TFIM scale tests.

- Single-node Qiskit-Aer's MPS can solve 33-qubit TFIM, although other simulators do not scale → Fig. 3
- Multinode NWQ-Sim was able to solve TFIM-32, although it was slower than Qiskit-Aer's MPS

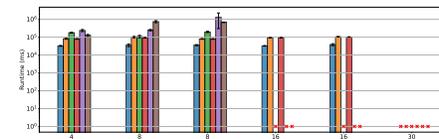


Fig. 4. QAOA with QFw.

- Tensor-Network simulators took over 2 hours to solve QAOA (QUBO >= 16) → Fig. 4

Future Work

- Hybrid HPC/QC execution of DQAOA (e.g., different subQUBOs could be solved using different QPMs)
 - Some on actual hardware as well
- QPMs can be auto-selected for solving a particular quantum circuit depending on its metadata (width, depth, number of connected components)
 - QFw could make that selection automatically

References

- Xu, Z., Chundury, S., Kim, S., Shehata, A., Li, X., Li, A., Luo, T., Mueller, F., and Suh, I.-S. (2025). GPU-Accelerated Distributed QAOA on Large-scale HPC Ecosystems. <https://arxiv.org/abs/2506.10531>
- Shehata, A., Groszkowski, P., Naughton, T. J., Gopalakrishnan Meena, M., Wong, E. Y., Chaves Claudino, D., Ferreira Da Silva, A., and Beck, T. L. (2025). Bridging paradigms: Designing for HPC-Quantum convergence. *Future Generation Computer Systems*, 174, 107980.
- Chundury, S., Shehata, A., Naughton, T. J., Kim, S., Mueller, F., and Suh, I.-S. (2024). QFw: A Quantum Framework for Large-scale HPC Ecosystems. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2024. <https://www.osti.gov/biblio/2498439>
- Kim, S., and Suh, I.-S. (2024). Performance Analysis of an Optimization Algorithm for Metamaterial Design on the Integrated High-Performance Computing and Quantum Systems. *Proceedings of the IEEE International Conference on Quantum Computing and Engineering (QCE 2024)*.
- Li, A., Fang, B., Granade, C. E., Prawiroatmodjo, G., Heim, B., Rötteler, M., and Krishnamoorthy, S. (2021). SV-Sim: Scalable PGAS-based State Vector Simulation of Quantum Circuits. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2021.
- McCaskey, A. J. *Tensor Network Quantum Virtual Machine (TNQVM)*. *Computer software*. <https://www.osti.gov/servelets/purl/1340180>. Vers. 00. USDOE. 18 Nov 2016. Web.
- Javadi-Abhari, A., Treinish, M., Kruslich, K., Wood, C. J., Lishman, J., Gacon, J., Martiel, S., Nation, P. D., Bishop, L. S., Cross, A. W., Johnson, B. R., and Gambetta, J. M. (2024). Quantum computing with Qiskit. <https://doi.org/10.48550/arXiv.2405.08810>
- Lykov, D., Chen, A., Chen, H., Kasper, K., Zhang, Z., Gibbs, T., and Alexeev, Y. (2021). Performance Evaluation and Acceleration of the QTensor Quantum Circuit Simulator on GPUs. *Proceedings of the IEEE/ACM 2nd International Workshop on Quantum Computing Software (QCS 2021)*. <http://dx.doi.org/10.1109/QCS4837.2021.00007>
- Kim, S., Luo, T., Lee, E., and Suh, I.-S. (2024). Distributed Quantum Approximate Optimization Algorithm on Integrated High-Performance Computing and Quantum Computing Systems for Large-Scale Optimization. <https://arxiv.org/abs/2407.20212>