A Scalability Study of Quantum Algorithms for Dimensionality Reduction of Multidimensional Data

Kareem El-Araby[1], Thom Popovici[2], Alvir Nobel[1], Sunday Ogundele[1], Katherine Klymko[2], Daan Camps[2], Anastasiia Butko[2], and Esam El-Araby[1]
[1]The University of Kansas (KU), Lawrence, KS, USA
{kareem_el_araby,islam.alvir,sunday,esam}@ku.edu
[2]Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, USA
{dtpopovici,kklymko,dcamps,abutko}@lbl.gov

Abstract

Quantum computing promises exponential speedups over classical computing by leveraging quantum-mechanical properties like superposition and entanglement. As quantum algorithms grow in complexity, classical simulation remains essential for evaluating correctness, scalability, and resource demands. This work focuses on studying the scalability of structured quantum algorithms such as the Quantum-Haar-Transform (QHT) usually used for reducing data dimensionality in signal/image processing and remote-sensing hyperspectral-imagery. We simulate QHT circuits on High-Performance-Computing (HPC) systems by constructing unitary models that mirror the transform's hierarchical decomposition. Simulations track performance metrics such as circuit width, circuit depth, and execution time. Our results provide insight into the practical implementation of structured quantum circuits and serve as a reference for validating algorithmic correctness and guiding future quantum algorithm design.

Keywords

Quantum-Haar-Transform, High-Performance-Computing, quantum-simulation, weak-scaling, Perlmutter supercomputer

1 Introduction and Background

Quantum computing promises a computational paradigm shift, offering exponential speedups over classical algorithms for a wide range of problems by leveraging the quantum-mechanical properties of superposition and entanglement [1]. Foundational algorithms such as Shor's [2] and Grover's [3] demonstrate advantages of quantum computing through parallelism and amplitude amplification. However, hardware constraints, including qubit count, short coherence times, and gate errors/noise, require classical high-performance-computing (HPC) simulation to assess algorithm performance and correctness. Moreover, the exponential growth of the quantum state-space renders simulations exceeding 30–34 qubits impractical, necessitating HPC resources [4–8].

The Haar transform is widely used for dimensionality reduction in applications such as signal-processing and hyperspectral-image analysis to separate data into low- and high-frequency components. We study its quantum analogue, the Quantum-Haar-Transform (QHT) [9], whose circuits (see Figure 1) feature logarithmic depth and polynomial gate complexity, making them suitable for Noisy-Intermediate-Scale-Quantum (NISQ) devices. Our simulations on the Perlmutter supercomputer reveal scaling trends in qubit count, decomposition depth, memory constraints, and runtime, validating the algorithm's structure and guiding future hardware implementations.

2 Proposed Methodology

The Quantum-Haar-Transform (QHT) was implemented as a quantum circuit in which each decomposition level was constructed using Hadamard and swap operations (see Figure 1). At each hierarchical level, the quantum state was split into low- and high-frequency components, and these components were subsequently recombined using entangling operations that preserved reversibility. For multi-level decomposition, we adopted the packet decomposition model [10], iteratively applying the same circuit at each level and postponing measurements until the final stage. We leveraged state-vector evolution instead of building the full $2^n \times 2^n$ unitary for n qubits, enabling us to simulate circuits up to 34 qubits. The circuit was implemented in C++, with custom simulation kernels designed to run in a distributed-memory environment. MPI was used to handle communication between nodes, and CUDA-enabled GPU acceleration for the compute-intensive operations. The overall workflow consisted of three stages: preprocessing to generate the gate sequence, a runtime construct that applied each operation, and postprocessing to record and analyze the results.

3 Experimental Work

3.1 Experimental Setup

All the simulations were run on the Perlmutter HPC system at the National-Energy-Research-Scientific-Computing-Center (NERSC) offering heterogeneous architecture optimized for both memory-bound and compute-intensive workloads. Each node on Perlmutter is equipped with four NVIDIA A100-Tensor-Core GPUs and dual AMD EPYC-7763 CPUs.

The QHT simulation was performed on 3D-data representations encoded into quantum state-vectors. We simulated 30 to 34 qubits using 4 to 64 GPUs. We conducted a weak-scaling study under conditions where both circuit-size and available compute power increased proportionally. The implementation included an optimization pass that balanced computational workload among GPUs. Data movement, however, remained relatively static and is identified as an area for future improvement.

SC25,

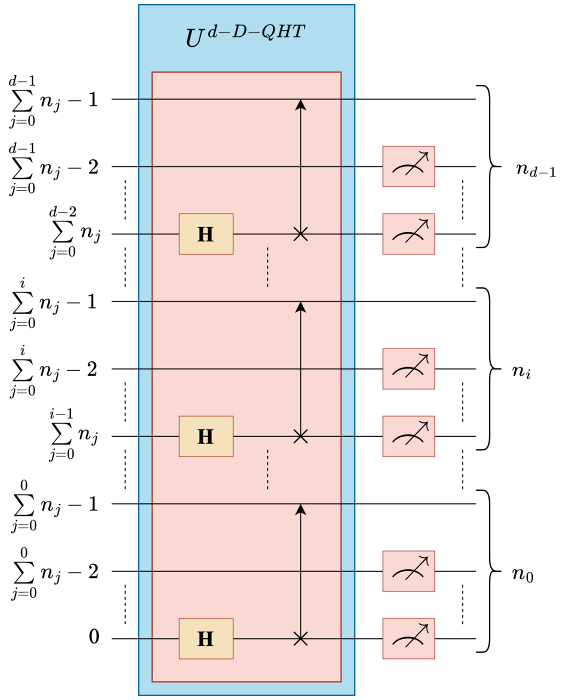

Figure 1: Depth optimized d-dimenstional QHT circuit

3.2 Experimental Results

The experiments demonstrate that our QHT simulation scales efficiently with respect to circuit width (number of qubits). Execution time remained nearly constant as we increased the number of qubits from 30 to 34 while also increasing the number of GPUs proportionally. This behavior suggests an effective time complexity of $O(1)$ under weak-scaling conditions. Conversely, when we increased the number of decomposition levels L in the QHT circuit, effectively increasing its depth, the simulation time grew linearly, which aligns with theoretical expectations of $O(\mathrm{L})$ scaling (see Figure 2).

4 Conclusion

In this work, we provide an analysis of the scalability and performance of the Quantum-Haar-Transform (QHT) when simulated using classical high-performance-computing (HPC) resources (Perlmutter supercomputer). The simulation strategy employed enables testing of circuits well beyond the size limitations of current quantum hardware. Our findings confirm that QHT can be implemented in a scalable fashion and that its behavior matches theoretical predictions for both circuit width and depth. We plan to refine the simulation engine to reduce communication latency and extend support to other quantum algorithms with structured circuit representations as well as potentially scaling on quantum hardware.

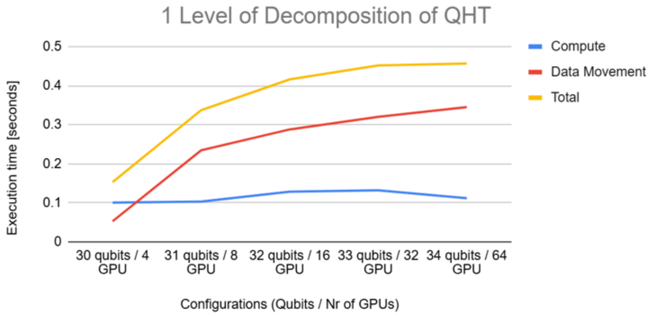

(a) 1-level of decomposition

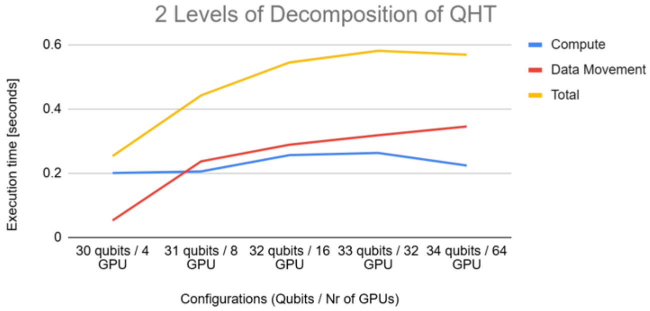

(b) 2-levels of decomposition

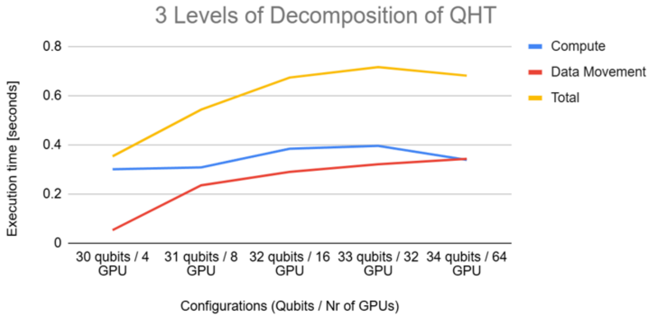

(c) 3-levels of decomposition

Figure 2: Total execution time vs. number of qubits (circuit width) with proportionally increasing GPUs (weak scaling)

Acknowledgment

This work was supported in part by the U.S. Department-of-Energy, Office-of-Science, Office-of-Workforce-Development-for-Teachers-and-Scientists (WDTS) under the Science-Undergraduate-Laboratory-Internship (SULI) program. This research used resources of the National-Energy-Research-Scientific-Computing-Center (NERSC), a Department-of-Energy Office-of-Science-User-Facility under Contract DE-AC02-05CH11231.

References

[1] Tanishq Nanda, B. K. Dhanalakshmi, and B. N. Lakshmi. An explorative cross sectional comprehensive survey on quantum computing. *SN Comput. Sci.*, 5, 2024.

[2] P. W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, Santa Fe, NM, USA, November 1994. IEEE Computer Society.

[3] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 212–219, Philadelphia, PA, USA, 1996. ACM.

[4] S. Heng, T. Kim, and Y. Han. Exploiting gpu-based parallelism for quantum computer simulation: A survey. *IEIE Transactions on Smart Processing & Computing*, 9(6):468–476, 2020.

[5] S. Efthymiou, S. Ramos-Calderer, C. Bravo-Prieto, A. Pérez-Salinas, D. García-Martín, A. Garcia-Saez, J.I. Latorre, and S. Carrazza. Qibo: a framework for quantum simulation with hardware acceleration. *Quantum Science and Technology*, 7(1):015018, 2021.

[6] H. Horii and C. Wood. Efficient techniques to gpu accelerations of multi-shot quantum computing simulations. arXiv preprint arXiv:2308.03399, 2023.

[7] T. Jones, A. Brown, I. Bush, and S.C. Benjamin. Quest and high performance simulation of quantum computers. *Scientific Reports*, 9(1):10736, 2019.

[8] N. Mahmud, E. El-Araby, and D. Caliga. Scaling reconfigurable emulation of quantum algorithms at high precision and high throughput. *Quantum Engineering*, 1(2):e19, 2019.

[9] N. Mahmud, A. MacGillivray, M. Chaudhary, and E. El-Araby. Decoherence-optimized circuits for multi-dimensional and multi-level decomposable quantum wavelet transform. *IEEE Internet Computing*, 26(1):47–55, 2022.

[10] Mingyoung Jeng, SM Ishraq Ul Islam, David Levy, Andrew Riachi, Manu Chaudhary, Md Alvir Islam Nobel, Dylan Kneidel, Vinayak Jha, Jack Bauer, Anshul Maurya, et al. Improving quantum-to-classical data decoding using optimized quantum wavelet transform: M. jeng et al. *The Journal of Supercomputing*, 79(18):20532–20561, 2023.