Benchmarking Classical Quantum Simulators for Quantum Fourier Transform on SCAYLE

Santiago Merino Bajo[1]*
Lidia Sánchez-González[2]
Cristina Esteban Blanco[1]
Álvaro Fanego Lobo[1]
Hilde Pérez[1]

1 SCAYLE Supercomputing Castilla y León, León, Spain
2 Robotics Group, I4 Institute, Universidad de León, Spain

*santiago.merino@scayle.es

This study presents a comprehensive comparative benchmark of classical quantum simulators executing the Quantum Fourier Transform (QFT)—a fundamental algorithm in quantum computing and precursor to applications such as Shor's algorithm and quantum phase estimation—on the high-performance infrastructure of the SCAYLE supercomputer. Given that classical simulation constitutes an essential stage for the validation and debugging of quantum algorithms prior to their implementation on actual hardware, this work establishes a systematic benchmark analyzing the scalability and efficiency of the main quantum development platforms currently available.
The experimental methodology was implemented on three distinct compute nodes with different architectures: Sapphire Rapids processors (2 CPUs, 2 TB RAM), IceLake (2 CPUs, 1 TB RAM), and NVIDIA Grace (1 CPU, 470 GB RAM). Four quantum simulation frameworks—Qiskit, Qulacs, PennyLane, and Qibo—were evaluated by generating QFT circuits ranging from 21 to 35 qubits, a limit determined by the computational capabilities of each environment. The implementations utilized exclusively CPU-based backends. The primary evaluation metric was total execution time, allowing quantification of the incremental computational cost associated with the exponential growth of the Hilbert space.

Results demonstrate an exponential increase in simulation time with respect to the number of qubits for all evaluated frameworks. In the range of 21 to 33 qubits, Qiskit emerged as the most efficient platform, achieving the best performance in large-scale simulations (33 qubits in 4346 seconds on Grace) and reaching the maximum of 35 qubits in 30205 seconds using the Sapphire Rapids nodes. Qulacs positioned itself as the second most competitive option, maintaining consistent times across different scales, while PennyLane exhibited significant degradation in high-dimensional circuits and Qibo recorded the highest execution times in the study. Regarding hardware infrastructure, the NVIDIA Grace processor consistently outperformed x86 architectures, followed by Sapphire Rapids.
These findings underscore the critical importance of selecting both the framework and hardware architecture for high-scale quantum simulations, evidencing that modern CPUs (particularly Grace) enable substantial improvements in the feasibility of simulating 30+ qubit systems in classical environments. Source code is available in the repository: https://github.com/smerib00/QuantumComputing.

References

[1] H. Abraham, A. Akhalwaya, et al., Qiskit: An Open-source Framework for Quantum Computing, 2021. https://qiskit.org
[2] IBM Quantum, "Quantum Fourier Transform (QFT) — IBM Quantum Learn", IBM Quantum (2024). https://quantum.cloud.ibm.com/learning/en/modules/computer-science/qft
[3] Y. Suzuki, et al., Qulacs: a fast and versatile quantum circuit simulator for research purpose, Quantum 5 (2021) 253
[4] [6] S. Efthymiou, S. Ramos-Calderer, C. Bravo-Prieto, A. Pérez-Salinas, D. García-Martín, A. Garcia-Saez, J. I. Latorre, S. Carrazza, Qibo: a framework for quantum simulation with hardware acceleration, arXiv:2009.01845v2 (2021). https://doi.org/10.48550/arXiv.2009.01845

Figures

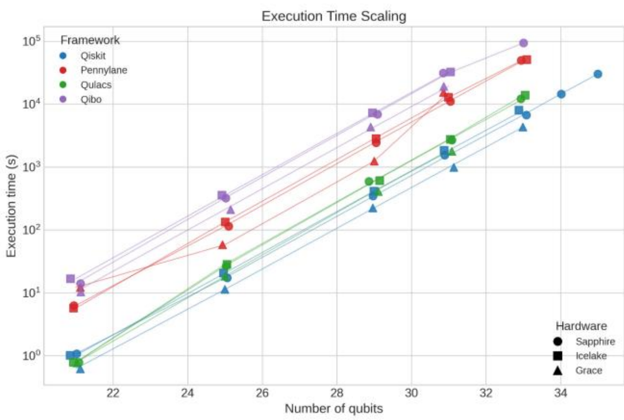

Figure 1: Execution Time vs Qubit Count for QFT Across Frameworks and Nodes