<u>***Benchmarking Emerging Quantum Machine Learning Devices***</u>

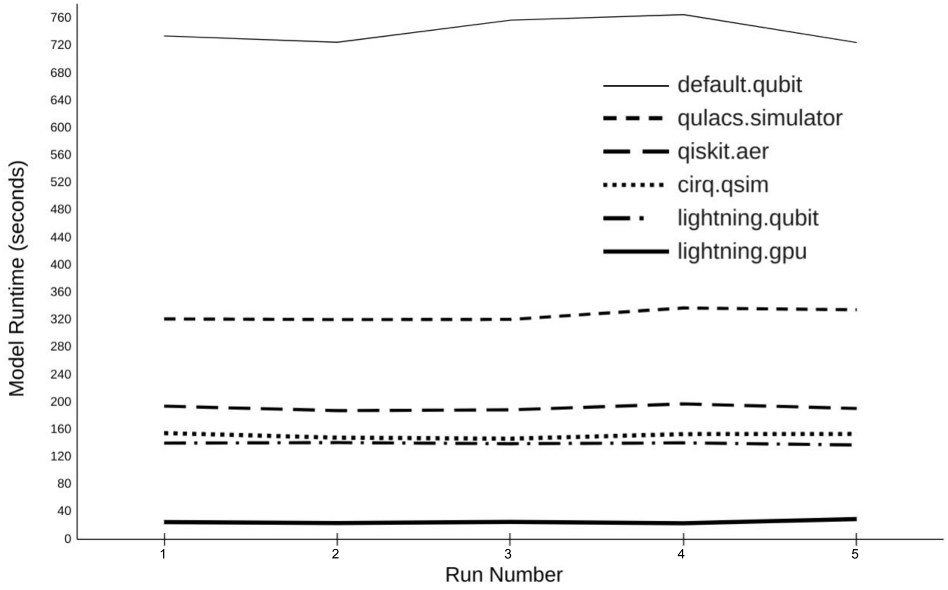

Run	default.qubit	qulacs.simulator	qskit.aer	cirq.qsim	lightning.qubit	lightning.gpu
1	733.88	320.99	193.68	154.45	139.78	24.43
2	724.66	319.98	187.19	147.73	140.56	23.13
3	756.91	320.25	188.26	146.23	138.89	24.70
4	765.06	336.96	197.05	152.91	140.15	23.01
5	724.20	334.37	190.31	153.22	137.01	29.08
Avg.	**740.94**	**326.51**	**191.30**	**150.90**	**139.28**	**24.87**

Table 1: Single threaded CPU simulators vs. Lightning.gpu GPU simulator
Note: Qulacs.simulator and qiskit.aer also have available GPU options
Cirq.qsim 90 core Xeon processor was likely only utilized single-threaded

Benchmarking Emerging Quantum Machine Learning Devices

"By far, the greatest danger of Artificial Intelligence is that people conclude too early that they understand it."

- Eliezer Yudkowsky, Machine Intelligence Research Institute

Introduction: Benchmarking a readily available graph quantum machine learning model using CPU and GPU quantum devices is practical for creating faster models. Here, permutation equivariant graph embeddings with trainable qubit rotations over the course of the run were successfully optimized. [1]

The quantum circuit size was increased to 25 qubits from 5 qubits to better distinguish single threaded CPU vs. single GPU speeds.
Pure states were utilized without quantum noise for convenient representations based on the surface of the "bloch sphere". [2]
Adjoint differentiation was employed for time and memory efficient state vector circuits, likely through the diff_method='best' setting. [3]
Exact expectation values were obtained each time the circuit was assessed, which is suitable for Machine Learning workflows. [4]

Results: PennyLane Lightning.gpu had the fastest average runtimes by 5x over the closest CPU method. Lightning.qubit and Cirq.qsim experienced the best CPU quantum simulator runtimes. Qiskit.aer simulator was slightly slower than the simulators mentioned, which was then followed by Qulacs.simulator. Default.qubit was over twice as slow as any other cpu quantum simulator, but is currently PennyLane's most compatible device.

Discussion: 5 runs with each quantum simulator clearly displayed time separations for the PennyLane "An equivariant graph embedding" demo without using multi-threading or multi-gpus on a 25 qubit quantum algorithm. Additional CPU benefits may be experienced with cirq.qsim's advanced CPU in theory, however several models only run the slowest default.qubit simulator unless modifications to code are made or a new notebook is written. Faster and more compatible quantum simulators are probable, based on the frequency of device updates in platform release notes. Lightning.gpu was accessed through the Colab paid V100 GPU tier, according to installation instructions. [5] CPU devices were installed using PennyLane-"Platform" in Colab with methods from the Devices and ecosystem documentation. [6]

Additional 2023 device benchmarkings were performed by Terra Quantum [7], and a comprehensive QED-C study was also published. [8]

Benchmarking Emerging Quantum Machine Learning Devices

References

1 Schuld, M. (2023b). An equivariant graph embedding. PennyLane Demos. https://pennylane.ai/qml/demos/tutorial_equivariant_graph_embedding
2 States | Quantiki. (n.d.-c). https://www.quantiki.org/wiki/states
3 Lee, C. (2021b). Adjoint differentiation. PennyLane Demos. https://pennylane.ai/qml/demos/tutorial_adjoint_diff
4 Measurements — PennyLane. (n.d.-b). https://docs.pennylane.ai/en/stable/introduction/measurements.html
5 Difficulties with Lightning-GPU installation. (2023, September 18). Xanadu Discussion Forum. https://discuss.pennylane.ai/t/difficulties-with-lightning-gpu-installation/3439/4
6 Plugins and ecosystem | PennyLane. (n.d.). https://pennylane.ai/plugins/
7 Benchmarking simulated and physical quantum processing units using quantum and hybrid algorithms. https://arxiv.org/abs/2211.15631
8 Optimization Applications as Quantum Performance Benchmarks. https://arxiv.org/abs/2302.02278

GPU Quantum Simulator Documentation

9 Qiskit — NVIDIA cuQuantum 23.06.1 documentation. (n.d.). https://docs.nvidia.com/cuda/cuquantum/latest/appliance/qiskit.html
10 Lightning-GPU device — PennyLane. (n.d.). Lightning-GPU 0.33.0-dev Documentation. https://docs.pennylane.ai/projects/lightning-gpu/en/latest/devices.html
11 Lightning-fast simulations with PennyLane and the NVIDIA cuQuantum SDK | PennyLane Blog. (n.d.). https://pennylane.ai/blog/2022/07/lightning-fast-simulations-with-pennylane-and-the-nvidia-cuquantum-sdk/
12 https://docs.nvidia.com/cuda/cuquantum/latest/appliance/cirq.html

Quantum Inspired Industry

13 Stanwyck, S. (2023, September 12). Quantum Boost: CuQuantum, PennyLane Let simulations ride supercomputers | NVIDIA Blogs. NVIDIA Blog. https://blogs.nvidia.com/blog/2023/09/12/quantum-supercomputers-pennylane/
14 Achieving Supercomputing-Scale Quantum Circuit Simulation with the NVIDIA cuQuantum Appliance | NVIDIA Technical Blog. (2023, July 11). NVIDIA Technical Blog. https://developer.nvidia.com/blog/achieving-supercomputing-scale-quantum-circuit-simulation-with-the-cuquantum-appliance/
15 Pires, F. (2023, May 3). Single-GPU systems will beat quantum computers for a while: research. Tom's Hardware. https://www.tomshardware.com/news/research-single-gpu-systems-will-continue-to-beat-quantum-computers-for-a-while