Benchmarking a Neutral Atom Quantum Computer

Nathan Wagner

October 13, 2023

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Overview

- Quantum computing has been suggested to have wide-ranging applicability in various fields from drug discovery [1] to climate modelling [2] to machine learning [3].
- To realize a quantum computer capable of solving these problems, we must first demonstrate that quantum computers are scalable

Overview

- Several candidates have been proposed for a scalable quantum computer
- Our group is developing a quantum computer where qubits are encoded in a lattice of cesium atoms
- My work aims to develop and justify a noise model which quantifies sources of noise in our architecture, and project how our computer may perform as it is scaled
- We also show that a improvement in qubit topology realizable in our architecture will significantly improve circuit fidelities, in some cases by $10 - 15\%$

Neutral-Atom Quantum Computing

- In our computer, cesium atoms are loaded in a blue-detuned optical grid spaced $3\,\mu\text{m}$ from each other, and are positioned with optical tweezers [4]
- Qubits are encoded in the hyperfine states $|0\rangle = |6s_{1/2}, f = 3, m = 0\rangle, |1\rangle = |6s_{1/2}, f = 4, m = 0\rangle$ used in quantum clocks
- Atoms are cooled using laser cooling to temperatures under $5\,\mu\text{K}$
- Measurements are performed by blowing away qubits in the $|1\rangle$ state from the grid and observing the luminosity of the remaining atoms

Neutral-Atom Quantum Computing

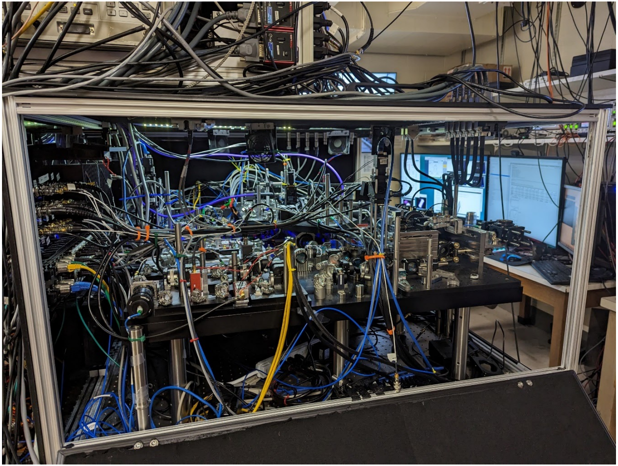

Figure: Image of our current quantum processor.

Neutral-Atom Quantum Computing

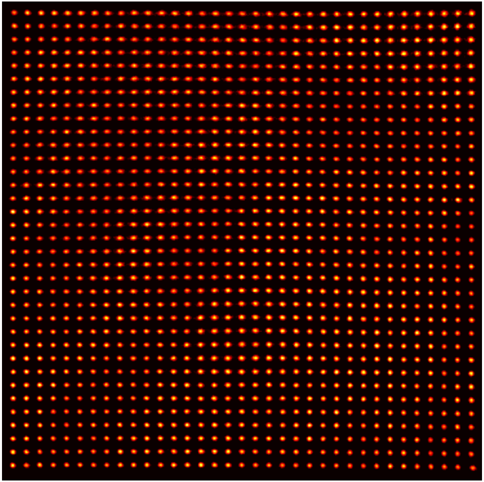

Figure: Cesium atoms loaded into our MOT in a square grid.

Neutral-Atom Quantum Computing

- Quantum gates are implemented by exciting the qubits with lasers
- Global $\mathsf{R_X}$ and $\mathsf{R_Y}$ qubits are implemented using a 9.2 GHz microwave laser tuned to the $|0\rangle \leftrightarrow |1\rangle$ transition
- Local $\mathsf{R_Z}$ gates are implemented by inducing a differential phase shift between the $|0\rangle$ and $|1\rangle$ states with a laser 0.76 GHz blue-detuned from the $|6s_{1/2}, f = 4, m = 0\rangle \leftrightarrow |7p_{1/2}\rangle$ transition
- By combining global $\mathsf{R_X}$ and $\mathsf{R_Y}$ gates with local $\mathsf{R_Z}$ gates we may implement local $\mathsf{R_X}$ and $\mathsf{R_Y}$ gates

Neutral-Atom Quantum Computing

- C_Z gates are realized using Rydberg blockage
- A Rydberg atom has a strong electric dipole, which shifts the energy levels of nearby atoms
- This effect may be used to entangle nearby atoms and realize a C_Z gate
- Atoms in the $|1\rangle$ state are excited to the $|7p_{1/2}\rangle$ state and then to the Rydberg state $|75s_{1/2}\rangle$ using a two-photon transition
- The R_X, R_Y, R_Z, and C_Z gates from a universal gate set

Modelling Noise

- Unique types of noise are caused by each gate
- In particular, an atom may end up in a state outside of the qubit basis due to a pulse area error or scattering in gates such as the C_Z gate
- This justifies the use of dark and bright decoherent states outside of the computational basis in our simulation, which do not partake in gate operations

Modelling Noise

- Our noise model was fit and compared to data collected in [4]
- We fit the model to our results from simulating MAXCUT circuits, and then compared our results to data from GHZ oscillations performed in the same paper
- As the error function is differentiable, the gradient was found using methods adapted from machine learning such as automatic differentiation, and optimization was performed using a gradient descent method

Modelling Noise

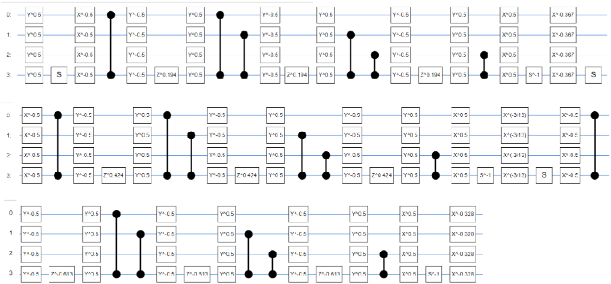

Figure: Sample MAXCUT circuit simulated in [4] which was used to fit our noise model

Modelling Noise

- The fidelity of simulated vs. observed MAXCUT data was 98.7% on average, with the lowest fidelity being 97.5% for the $p = 3$ circuit
- The noise model overpredicted the GHZ fidelity on average by 8.0%, with a standard deviation of 5.5%.

Modelling Noise

- Discrepancies are likely due to the fact that, on our physical computer, certain sites may have better or worse fidelity than others
- Since we aim to project how the computer will perform with circuits which have more qubits than we are able to implement now, we cannot allow for site-specific errors
- Also, more complicated sources of error, such as crosstalk between sites, are present in our hardware and were not modelled

Modelling Noise

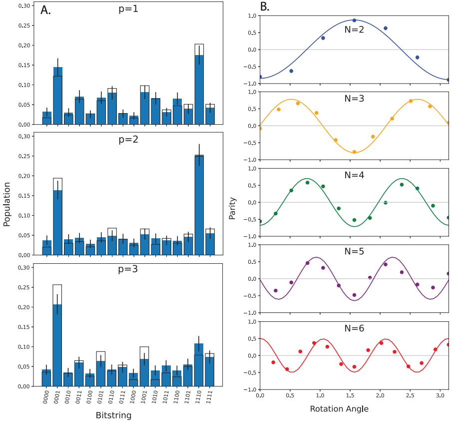

Figure: Comparison of simulated data vs. observed data for A. MAXCUT circuits and B. GHZ oscillations. In A, the blue filled bars represent simulated data and unfilled bars represent observed data.

Modelling Noise

- Global R_X and R_Y gates were found to be the least noisy, with an average gate fidelity over 1 π pulse to by 0.9922.
- Local R_Z gates were more noisy, with an average gate fidelity over 1 π pulse of 0.984
- The C_Z was found to be the most noisy, with an average gate fidelity of 0.969
- Note that the fidelities for each gate were obtained as an average over the full Bloch sphere, so for certain states the fidelity of the gate may be higher or lower
- In particular the fidelity of states populated heavily with $|1\rangle$ qubits were lower on average and the fidelity of states population heavily with $|0\rangle$ qubits were observed to be higher on average since most errors disproportionately affect the $|1\rangle$ state.

Modelling Noise

Gate/Process	Noise	Error	Avg. Gate Fidelity
Global R_X, R_Y	Depolarization (per π pulse)	1.8×10^{-6}	0.9922
Local R_Z	phase-flip (per π phase)	3.2×10^{-4}	0.984
	Loss to Dark State (per π pulse)	1.9×10^{-4}	
	Loss to Bright State (per π pulse)	2.7×10^{-4}	
	Polarization (per π pulse)	2.0×10^{-8}	
C_Z	Phaseflip	3.3×10^{-2}	0.969
	Loss to Dark State	1.8×10^{-2}	
	Loss to Bright State	2.9×10^{-2}	
	Polarization	2.1×10^{-5}	
	Phaseshift	-2.0×10^{-3} rads	
SPAM	Preparation	5.2×10^{-3}	NA
	Measurement	5.3×10^{-3}	
Decoherence	$T_1 = 10$ s	NA	NA
	$T_2^* = 3.5$ ms		
	$P_{\|0\rangle} = 0.42$ at $t = \infty$		

Table: List of the types of error used in the noise model, along with the probabilities of each being applied. The indicated gate fidelities are SPAM-corrected.

Benchmark Results

- We applied this model to predict how the computer may perform against a benchmark published by the Quantum Economic Development Consortium [5]
- We simulated 10 types of circuits using the noise model developed

Benchmark Results

- Each circuit type which was simulated belonged to one of 3 broad categories: circuits based on an oracle, circuits commonly used as subroutines, and application-based circuits
- For each category, 3-4 types of circuits were chosen to be simulated
- For each circuit type, up to 3 circuits were chosen for each possible circuit width

Benchmark Results

- Our qubit topology is currently based on the square grid used to contain the qubits, in that only neighboring qubits may be entangled using C_Z gates
- However, the Rydberg interaction used in our architecture to entangle qubits has a long range, meaning that it may be used to entangle far-away qubits
- We compiled and simulated each circuit with both topologies; one being a square grid of qubits connected horizontally and vertically (nearest-neighbor connectivity), and one where any two qubits may be entangled (all-to-all connectivity)
- Our noise model predicts a significant improvement in fidelity in circuits with all-to-all connectivity compared to those with nearest-neighbor connectivity

Benchmark Results

Figure: Implementations of a 4-qubit phase estimation circuit on our architecture with both topologies.

Benchmark Results

- We modelled the full density matrix evolution of the quantum system, treating each qubit as a 4-level quantum object with states $|0\rangle, |1\rangle, |l\rangle_0, |l\rangle_1$
- Error channels were modelled as Kraus operators which acted on the density matrix after each gate
- Most circuits were simulated in Cirq, with some high-width circuits being manually implemented in Numpy and Scipy for efficiency

Benchmark Results

- Circuits including an oracle which we simulated include the Deutsch-Jozsa, Bernstein-Vazirani, and hidden shift algorithms
- The simulator runs the Bernstein-Vazirani and Deutsch-Jozsa circuits with up to 5 qubits with $> 70\%$ fidelity, and the Hidden Shift circuit with up to 4 qubits with $> 60\%$ fidelity
- Accuracy is retained at high qubit counts, with 11-qubit implementations of Bernstein-Vazirani and Deustch-Jozsa circuits giving $> 50\%$ fidelity on average

Benchmark Results

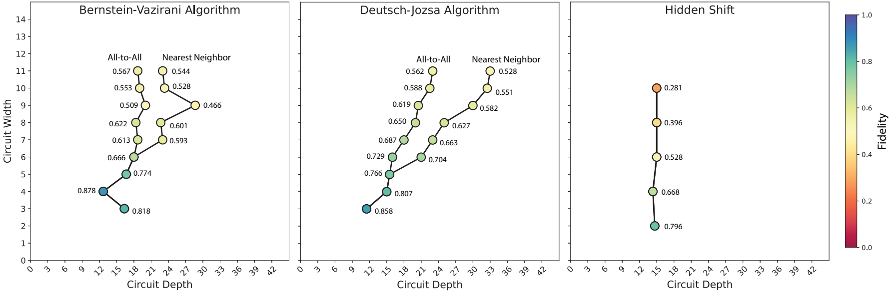

Figure: Simulated results for the Deutsch-Jozsa, Bernstein-Vazirani, and Hidden Shift benchmarks for all-to-all and nearest-neighbor connectivity.

Benchmark Results

- Subroutine circuits simulated included the inverse quantum Fourier Transform (Method 2), a QFT chained with an inverse QFT (Method 1), and phase and amplitude estimation circuits
- All-to-all connectivity was found to significantly improve the performance of the inverse quantum Fourier Transform and phase estimation circuits
- The fidelity of the inverse quantum Fourier transform increased by 13% in the 3-5 qubit range on average, peaking at an increase of 17% at the 5-qubit implementation
- The average increase in fidelity of phase estimation circuits on the same range was 22%, and peaks at a 26% increase at the 4 qubit implementation

Benchmark Results

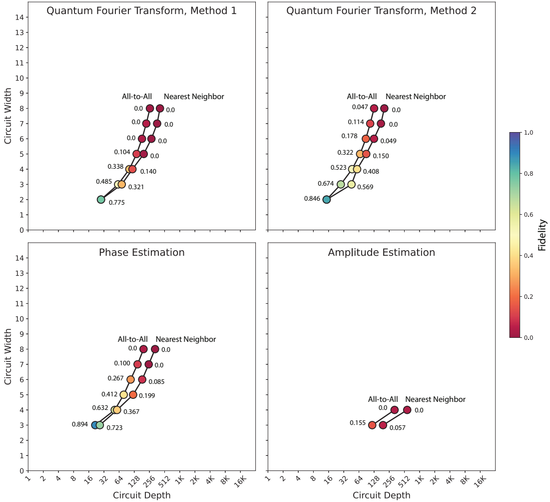

Figure: Simulation results for the phase estimation, amplitude estimation, and both Quantum Fourier Transform benchmarks.

Benchmark Results

- Application-oriented circuit which we simulated included Grover's search algorithm, a simulation of a Hamiltonian, and a quantum Monte Carlo algorithm
- Both the Grover's search and Hamiltonian estimation had high-fidelity results for 2 qubits
- All Hamiltonian circuits with 5 or more qubits, along with all other circuits with 3 or more qubits, returned an average fidelity of $< 15\%$ for both nearest-neighbor and all-to-all connectivity simulations

Benchmark Results

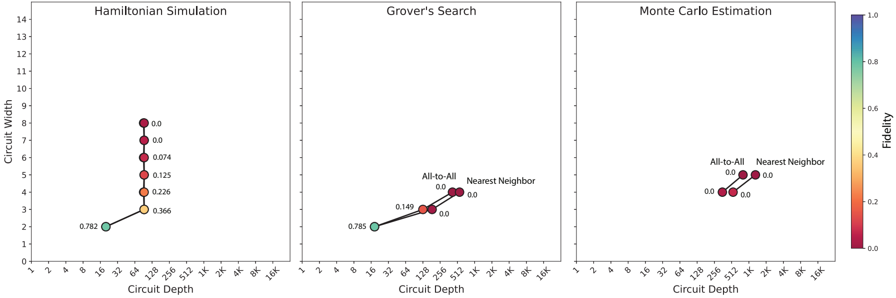

Figure: Simulation results for Hamiltonian simulation, Grover's search, and Monte Carlo estimation.

Discussion

- In this work, we propose a noise model quantifying sources of noise in our computer, and predict that utilizing the long range of Rydberg blockage will significantly improve our computer's results
- Rydberg blockage additionally has the potential to make neutral-atom quantum computer easier to scale than other platforms

Discussion

- We will investigate how neutral-atom quantum computers may be scaled taking Rydberg blockage into account, and will focus on increasing the range and fidelity of our C_Z gates
- We believe that this approach is promising for near-future development of quantum simulators and computers
- We also plan on continuing to use application-oriented benchmarks as metrics to quantify our computer's performance

Discussion

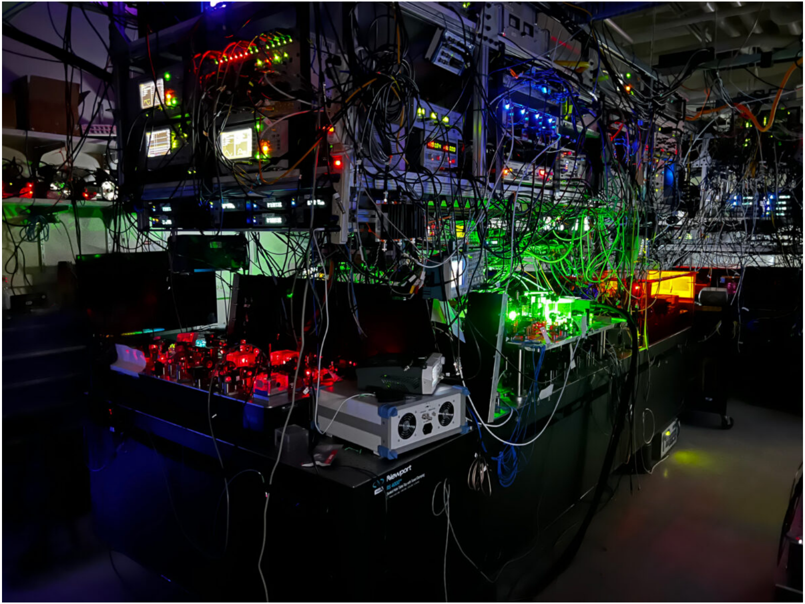

References

[1] Y. Cao, J. Romero, and A. Aspuru-Guzik. "Potential of quantum computing for Drug Discovery". In: *IBM Journal of Research and Development* 62.6 (2018). DOI: 10.1147/jrd.2018.2888987.

[2] Vojtěch Havlíček et al. "Supervised learning with quantum-enhanced feature spaces". In: *Nature* 567.7747 (2019), pp. 209–212. DOI: 10.1038/s41586-019-0980-2.

[3] Michael Lubasch et al. "Variational quantum algorithms for nonlinear problems". In: *Physical Review A* 101.1 (2020). DOI: 10.1103/physreva.101.010301.

[4] T. M. Graham et al. "Multi-qubit entanglement and algorithms on a neutral-atom quantum computer". In: *Nature* 604.7906 (2022), pp. 457–462. DOI: 10.1038/s41586-022-04603-6.

[5] Thomas Lubinski et al. "Application-Oriented Performance Benchmarks for Quantum Computing". In: (2021). DOI: 10.48550/ARXIV.2110.03137. URL: https://arxiv.org/abs/2110.03137.