Implementation of Partial Quantum Search

James Chen, Debayan Bandyopadhyay, Amy Searle, Victor Li, Sahil Patel, Henry Sun,
Glenn LeBlanc, Ashwin Rastogi, Ryan Jia

Summary

Grover's Algorithm, one of the best-known algorithms of quantum computing, provides a square root speedup for unstructured database search. Eight years after publishing his paper on the algorithm, Grover worked with Radhakrishnan on a variation of the problem: partial search, where only the first few bits of the target element were desired. Grover and Radhakrishnan showed that it was possible to perform a partial search in $O(\sqrt{N}(1-\frac{1}{\sqrt{K}}))$ queries to the database, where N is the total size of the database and K is the number of blocks. Our group implemented this partial search algorithm using Rigetti's pyquil, a Python library built on top of QUIL (quantum instruction language).

Quantum Computing - Introduction

Qubits

- A **qubit** is a two-level quantum system with basis states $|0\rangle$ and $|1\rangle$.
- All possible states of a qubit are unit-norm complex linear combinations of the two basis states. For example, $\frac{1}{\sqrt{2}}|0\rangle + \frac{i}{\sqrt{2}}|1\rangle$ is a possible qubit state. The coefficients of $|0\rangle$ and $|1\rangle$ are their **amplitudes**.
- The joint state of multiple qubits is the tensor product of all individual qubit states. For example, if our first qubit is in the $|1\rangle$ state and our second qubit is in the $\frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle$ state, the joint state would be $|1\rangle \otimes (\frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle) = \frac{1}{\sqrt{2}}|10\rangle - \frac{1}{\sqrt{2}}|11\rangle$.
- To extract information from qubits, we have to measure. While qubits can be in a superposition of multiple basis states (as shown above), measurement only gives one basis state. The probability of measuring the basis state b is $|\alpha_b|^2$, where α_b is the amplitude of the state along b. For example, if we have a two-qubit sytem $\frac{1}{\sqrt{2}}|00\rangle - \frac{1}{\sqrt{2}}|11\rangle$, we have a $\left|\frac{1}{\sqrt{2}}\right|^2 = \frac{1}{2}$ probability of measuring $|00\rangle$ and a $\left|-\frac{1}{\sqrt{2}}\right|^2 = \frac{1}{2}$ probability of measuring $|11\rangle$.

Quantum Gates

- A **quantum gate** G is a unitary linear transformation on the vector of amplitudes, i.e. $G^\dagger G = I$.
- One of the most commonly used single-qubit gates is the Hadamard gate, H. Written as a matrix: $H = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}$.
- We can tensor product quantum gates to create larger quantum gates. For example, $H^{\otimes 2} = H \otimes H$ is a two-qubit quantum gate which performs a Hadamard gate on both qubits.

Grover's Algorithm – Introduction

Problem:

Consider a function $f : \{0,1\}^n \to \{1,-1\}$, where f maps an unknown target bitstring ω to -1 and all other bitstrings to 1. How many queries do we need to f to determine ω?

Classically:

Since we have no knowledge on the state, our best strategy is to guess randomly. Then we expect to find ω after $\frac{N}{2}$ queries, where $N = 2^n$. On average, we require $O(N)$ queries.

Quantum Search

- We start with a uniform superposition over all n-qubit states: $|s\rangle = \sum_{b\in\{0,1\}^n} \frac{1}{\sqrt{N}}|b\rangle$.
- We pass this state into f, which negates the amplitude of the target $|\omega\rangle$.
- Next, we perform a gate which reflects our state about the initial uniform superposition $|s\rangle$ (see below). This reflects all amplitudes about the mean.
- Performing the two steps above constitutes a **global amplification** step. Repeating the global amplification step rotates our quantum state towards the target state $|\omega\rangle$. The probability of measuring $|\omega\rangle$ is highest when we perform global amplification $\frac{\pi}{4}\sqrt{N}$ times, for $O(\sqrt{N})$ total queries.

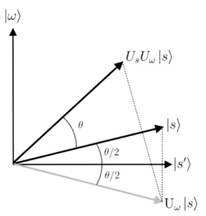

Grover's Algorithm - Implementation

- First, we initialize an n-qubit quantum virtual machine (QVM).
- To generate the initial superposition, we perform Hadamard gates on all n qubits (known as a Hadamard Transform, or $H^{\otimes n}$), which start in the $|00\cdots0\rangle$ state.
- To create the global mean inversion gate, we prepare a special diagonal matrix, and then prepend and append Hadamard Transfroms (see below).

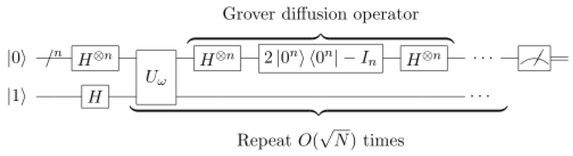

Grover's Algorithm - Results

Below are graphs for the probabilities of the target state and non-target states for a search over 5 qubits. The optimal number of iterations in this case was 4.

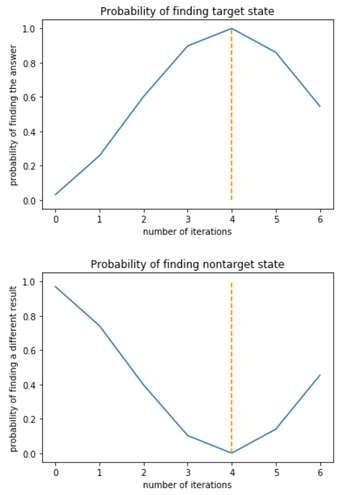

Below is the graph for the optimal number of global amplification steps for Grovers original algorithm as a function of the total number of qubits.

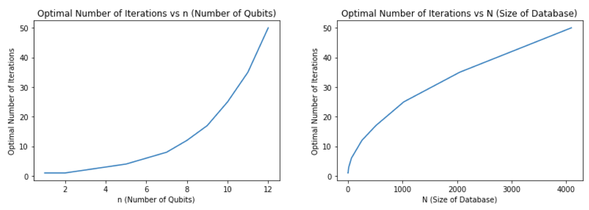

Grover's Algorithm - Analysis

- In the first two graphs, the sinusoidal nature of the target and non-target probabilities demonstrates our state rotating towards the target state.
- After performing the optimal number of global amplification steps (4), the state begins to overrotate past the target state. This overrotation, which is bad for global search, plays a key role in the partial search algorithm (in the step where negative amplitude is pushed into the target block).
- In the bottom right graph above, we see the $O(\sqrt{N})$ behavior of the optimal number of global amplification steps. While this is still exponential in the number of qubits, we can do no better for unstructured search.

Partial Quantum Search - Introduction

Problem:

Consider a function $f : \{0,1\}^n \to \{1,-1\}$, where f maps an unknown target bitstring ω to -1 and all other bitstrings to 1. How many queries do we need to f to determine the first k bits of ω?

Classically:

Since we have no knowledge on the state, our best strategy is to search all bitstrings with all but one of the possible k bit prefixes (say, $00\cdots0$). If the target is not found, we know it must have the prefix that was not searched. Then we expect to find ω after $\frac{N}{2}(1-\frac{1}{K})$ queries, where $N = 2^n$ and $K = 2^k$. On average, we require $O(N(1-\frac{1}{K}))$ queries.

Partial Quantum Search

- Again, we start with a uniform superposition over all n-qubit states: $|s\rangle = \sum_{b\in\{0,1\}^n} \frac{1}{\sqrt{N}}|b\rangle$.
- We perform the aforementioned global amplification step $l_1(\varepsilon) = \frac{\pi}{4}(1-\varepsilon)\sqrt{N}$ times, where ε is a tunable parameter. This pushes amplitude into the target state so we will have enough negative amplitude in the target block for the next step.
- A **block amplification** step constitutes an oracle query of the current state and a gate which inverts about the mean in each block.
- We perform the perform the block amplification step $l_2 = (\frac{\theta_1+\theta_2}{2})\sqrt{\frac{N}{K}}$ times. This process makes the amplitudes in the target block increasingly negative.
- The above step gradually decreases the global amplitude mean. Once the global amplitude mean reaches exactly half of the amplitudes in the non-target blocks, we perform one final global inversion, eliminating all the amplitude in the non-target blocks.

Partial Quantum Search - Implementation

- Again, we first initialize an n-qubit quantum virtual machine (QVM) and generate a uniform superposition using a Hadamard Transform.
- We create the global mean inversion gate as before.
- To create the block mean inversion gate, we prepare a slightly different diagonal matrix, and then prepend and append Hadamard Transforms.
- We used the optimal number of iterations given above, taking $\varepsilon = \frac{1}{\sqrt{K}}$.

Partial Quantum Search - Results

Below are graphs for the probabilities of the target state, non-target block states, all block states, and non-block states, in order. The partial search performed below was on 9 total qubits, with 2 target qubits (hence 4 total blocks, each of size 128). We performed 9 global iterations and 5 block iterations.

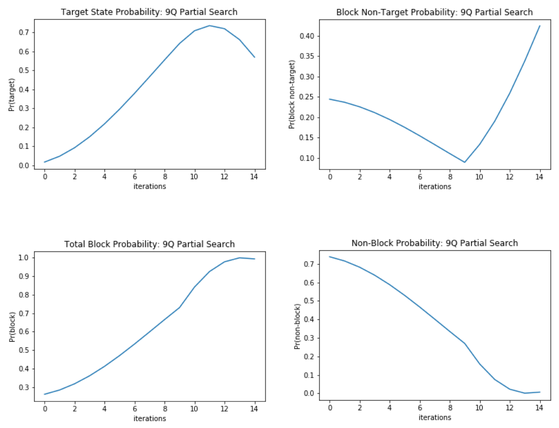

Partial Quantum Search - Results

Below, on the left, we have a graph of the optimal number of global and block iterations as a function of the number of target qubits, with the total number of qubits fixed at 9. On the right, we have a graph of the optimal number of global and block iterations as a function of the number of qubits, with the total number of blocks fixed at 4.

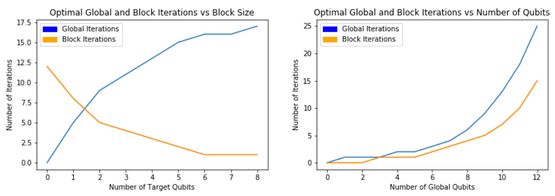

Below, we have graphs of the probability in the target block as a function of the number of oracle queries for the original and partial search algorithms. The number of target qubits is 1 and 3 for the left and right graphs, respectively.

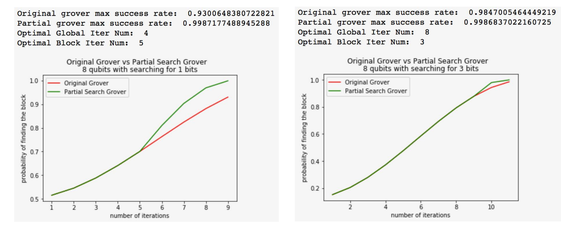

Partial Quantum Search - Analysis

- As shown in the last two plots above, partial quantum search outperforms Grover's original algorithm, reaching a block probability of 1 in fewer iterations.
- The advantage is most noticable when $K \ll N$ (confirmed result in paper). In the case of 8 total qubits, the speedup is significant for 1 or 2 target qubits, but drops off around 4 target qubits, at which point the global and partial search algorithms are identical (i.e. 0 block iterations is optimal).
- The closed form for the optimal number of iterations is unstable at 0 target qubits. In this case, no queries are necessary at all. However, the optimal number of block iterations is 12.

Conclusion

Our group successfully implemented Grover and Radhakrishnan's partial quantum search algorithm in pyquil. We are looking forward to possible future extensions, including:

- Exact Rotations: It was shown that Grover's original algorithm could be modified to give a result with certainty, by underrotating slightly on the last iteration to line up exactly with the target state. Grover and Radhakrishnan noted that this was possible to do in the partial search problem as well. This would be interesting to implement.
- Binary Search by Repeated Partial Search: It is possible to use the partial quantum search as a submodule in a binary search algorithm (i.e. identify the first bit of the target, then identify the second bit, etc.). Korepin and Xu showed that this process was no faster than performing a global search. However, it might be interesting to experimentally verify this.
- Searching an Oracle with Multiple Target States: In the original search problem, an oracle with multiple target states is handled by simply decreasing the number of iterations appropriately. However, the result is not extended to the partial search case. Examining whether a partial search analog exists and deriving it seems like a challenging extension.

References

[1] Lov K. Grover. A fast quantum mechanical algorithm for database search, 1996; arXiv:quant-ph/9605043.

[2] Lov K. Grover and Jaikumar Radhakrishnan. Is partial quantum search of a database any easier?, 2004; arXiv:quant-ph/0407122.