Quantum Algorithms 2/21/2023

Lecture 9: Quantum Search

Instructor: Dieter van Melkebeek

In this lecture we cover quantum search. We begin with the problem statement for search problems and then present Grover's algorithm. Finally, we discuss a generalization of the approach to the search problem and explore the possibility of error elimination.

1 Solution to Exercise #7

Part (a). For each $i \in [n]$ such that $a_i = 1$, include a CNOT with $|x_i\rangle$ as control and $|y\rangle$ as the target. This has the accumulated effect of mapping $|x\rangle\,|y\rangle$ to $|x\rangle\,|y \oplus \sum_{i=1}^{n} a_i x_i\rangle = |x\rangle\,|y \oplus f(x)\rangle$. See the following figure for an example with $n = 5$ and $a = 11001$.

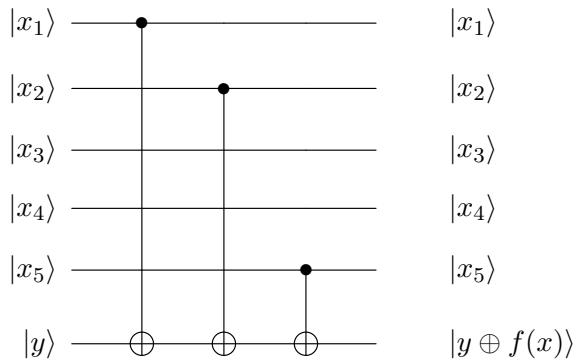

Part (b). We want to argue the following equivalence:

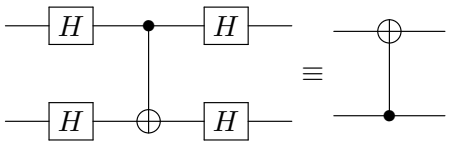

We will present four solutions for this part. They all boil down to the same but offer different perspectives.

Traces. Given that both circuits are unitary, it suffices to show that they behave the same on each of the standard basis states $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. There are several ways to organize this computation and express the intermediate states. The following representation of the traces of the circuit on the left is inspired by the third solution:

$$|00\rangle \mapsto |+\rangle\,|+\rangle \mapsto |+\rangle\,|+\rangle \mapsto |00\rangle \qquad |10\rangle \mapsto |-\rangle\,|+\rangle \mapsto |-\rangle\,|+\rangle \mapsto |10\rangle$$
$$|01\rangle \mapsto |+\rangle\,|-\rangle \mapsto |-\rangle\,|-\rangle \mapsto |11\rangle \qquad |11\rangle \mapsto |-\rangle\,|-\rangle \mapsto |+\rangle\,|-\rangle \mapsto |01\rangle$$

In all four cases the net effect is that same as the CNOT gate on the right.

One can also use another basis than the standard one. See the third solution for a judicious way of choosing the basis.

Transition matrices. Explicitly compute the transition matrices for both unitary circuits and show that they are the same. The one for the circuit on the left is

$$\begin{aligned} H^{\otimes 2} \cdot \mathrm{CNOT} \cdot H^{\otimes 2} &= \frac{1}{\sqrt{2}} \begin{bmatrix} H & H \\ H & -H \end{bmatrix} \cdot \begin{bmatrix} I & 0 \\ 0 & X \end{bmatrix} \cdot \frac{1}{\sqrt{2}} \begin{bmatrix} H & H \\ H & -H \end{bmatrix} \\ &= \frac{1}{2} \begin{bmatrix} H^2 + HXH & H^2 - HXH \\ H^2 - HXH & H^2 + HXH \end{bmatrix} = \frac{1}{2} \begin{bmatrix} I + Z & I - Z \\ I - Z & I + Z \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \end{aligned}$$

We used the fact that NOT is the same as X, that the Hadamard transform H is its own inverse ($H^2 = I$) and that the underlying basis change transforms an X gate (bit flip) into a Z gate (phase flip): $HXH = Z$. The final transition matrix swaps $|01\rangle$ and $|11\rangle$, which is what the circuit on the right does.

(Reverse) phase kickback. For a controlled operation like CNOT, it is generally interesting to consider an eigenbasis of the operator for the controlled qubits. On an eigenvector $|\phi\rangle$ with eigenvalue λ, the effect is a noop when the control qubit is in $|0\rangle$, and a phase change given by λ when the control qubit is in $|1\rangle$.

$$\begin{aligned} |0\rangle\,|\phi\rangle &\mapsto |0\rangle\,|\phi\rangle \\ |1\rangle\,|\phi\rangle &\mapsto |1\rangle\,\lambda\,|\phi\rangle \end{aligned}$$

By kicking back the phase change to the control qubit, the operation can be interpreted as leaving the controlled qubits untouched and performing an operation (phase change) on the controlling qubit, where the operation is determined by the state of the controlled qubits. This is exactly the change in perspective we seek. Note that the kickback goes in the reverse direction to the one we used last lecture.

In this case the operator is X, which has $|+\rangle$ as an eigenvector with eigenvalue $\lambda_+ = 1$, and $|-\rangle$ with eigenvalue $\lambda_i = -1$. The effect for $|+\rangle$ on the control qubit is nothing. The effect for $|-\rangle$ on the control qubit is that of a phase flip:

$$\alpha_0\,|0\rangle + \alpha_1\,|1\rangle \mapsto \alpha_0\,|0\rangle\ -\alpha_1\,|1\rangle\,.$$

Since H maps the standard basis $(|0\rangle\,, |1\rangle)$ to $(|+\rangle\,, |-\rangle)$, this shows that:

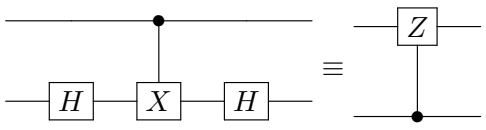

As the noop in the case of $|+\rangle$ can be written as H^2, and the Z gate in case of $|-\rangle$ as HXH, we have:

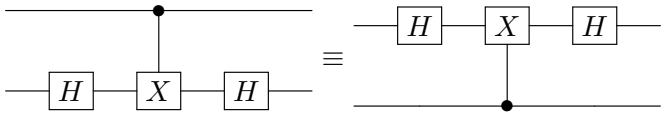

Applying an Hadamard gate on the top qubit before and after yields the desired equivalence:

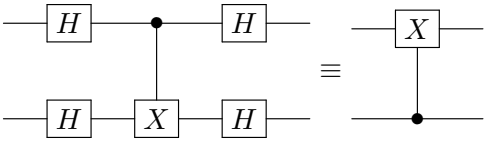

Symmetry of controlled Z. An ad-hoc solution makes use of the fact that the effect of a controlled Z-gate is symmetric in its arguments as it leaves each of the tree basis states other than $|11\rangle$ unaffected, and maps $|11\rangle$ to $-|11\rangle$. In combination with $H^2 = I$ and $HXH = Z$, this leads to the following circuit manipulations:

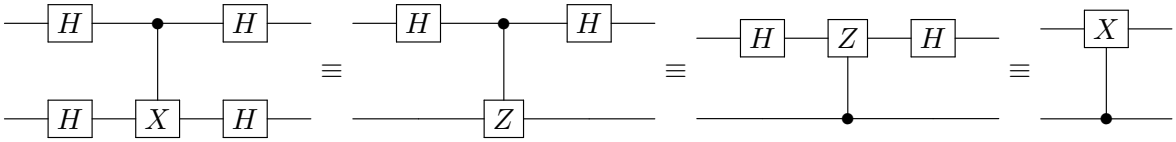

Part (c). Plugging in the circuit for U_f from part (a) into our circuit for learning linear functions (without the final measurement) yields the circuit on the left in Figure **??**. Inserting the noops H^2 in the bottom qubit line after each CNOT, and applying the equivalence from part (b) to each CNOT yields the equivalent circuit on the right.

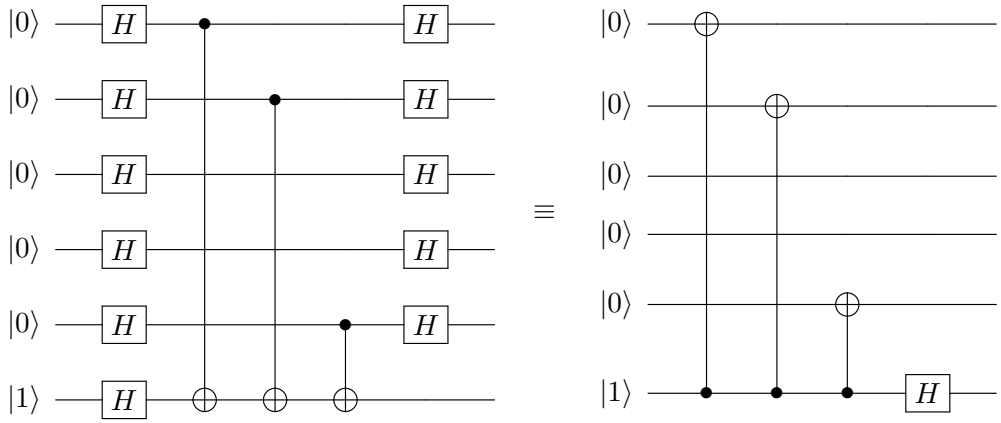

Figure 6: Alternate explanation of learning linear functions.

Note that the control bit for each of the CNOTs in the above circuit is 1, so each of the controlled qubits gets flipped from 0 to 1. As such, the top n qubits are in state exactly $|a\rangle$ in the end (and the last qubit in state $|-\rangle$). This gives an alternate explanation why a final measurement of the top n qubits yields a with certainty.

2 Search

2.1 Problem Statement

We are given access to a black-box function $f : \{0,1\}^n \to \{0,1\}$ and want to find an input $x \in \{0,1\}^n$ such that $f(x) = 1$. We call such items "'good" and the other ones "bad."

In this lecture, we assume that there is at least one good item and that the number of good items $t \doteq |f^{-1}(1)|$ is known. Next lecture we will address the case where t is unknown and can be zero. The overall problem can be seen as a search for good items in an unsorted list of items. The hardest case when $t > 0$ is $t = 1$ since there is only a single good item for the algorithm to find.

2.2 Algorithms

Consider a search for a good item in a list with $N = 2^n$ items, t good items, and $b = N - t$ bad items. We analyze the query complexity in the three usual settings.

Deterministic setting. To give a deterministic answer, in the worst-case scenario, the algorithm will have to go through all $N-t$ bad items before knowing a good item, namely any of the remaining items. Therefore, the query complexity is $N - t$.

Probabilistic setting. The chance of getting a good item from queries of items at random is $p = t/N$. This can be viewed as a Bernoulli experiment with success probability p. The expected of trials until the first success in a Bernoulli experiment is $1/p$, and the number of trials to guarantee success with any constant level of confidence less than 1 is $\Theta(N/t)$.

Quantum setting. For quantum algorithms, we will see that $\Theta(\sqrt{N/t})$ queries suffices to guarantee success with 100%. The improvement over probabilistic algorithms is quadratic for any fixed level of confidence less than 100%.

2.3 Applications

Although the quantum search algorithm seems useful for unsorted database search problems, it does not offer any clear advantage over classical algorithms. To utilize quantum search algorithms, quantum black-box functions must be constructed, and the entire database must be uploaded to the black-box for interference and superposition to occur. Furthermore, the database is normally sorted or can be sorted, which allows for other more efficient algorithms than a linear search, such as binary search, to complete the task.

One useful application is the satisfiability problem, where the algorithm checks each item if it satisfies a condition and assigns a Boolean output. In this case, the running time of classical algorithms will be 2^n. There are other more efficient classical algorithms known for this problem, but they all still run in time $O(2^n)$ where n is the number of variables. Using quantum search, we can improve the running time to $\tilde{O}(2^{n/2})$.

3 Quantum Approach

The approach can be broken into 3 main procedures.

1. Create a uniform superposition of all possible x

$$|\psi_0\rangle = \sum_x \alpha_x |x\rangle = \frac{1}{\sqrt{N}} \sum_x |x\rangle \tag{1}$$

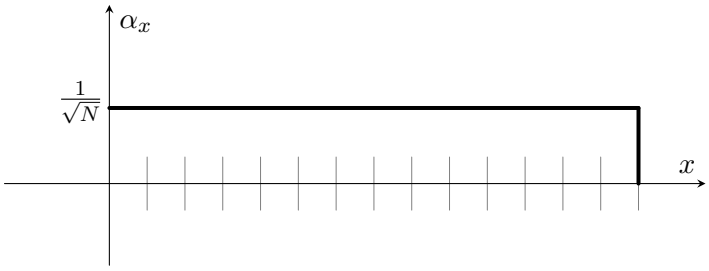

Figure 7: The initial state of the system. Each state is equally likely to be observed if a measurement is taken. We refer to this state as the uniform superposition.

2. Apply unitary operations to boost the weight of the good x's.
3. Measure $|\psi_{final}\rangle$ and output the observed x.

The first step of the algorithm can be accomplished by applying $H^{\otimes n}$ to $|0^n\rangle$. The second step is more complex and can be broken down into two unitary procedures. The first flips good components about the x-axis and the second flips all components about the average.

Flip good components about the x-axis: R_{bad}. The goal of the R_{bad} operator is to flip the sign of good items x and leave bad items invariant such that

$$R_{bad} |x\rangle = (-1)^{f(x)} |x\rangle \tag{2}$$

Thus, the operator R_{bad} flips the phase iff x_i is a good item, that is $f(x_i) = 1$. The resulting amplitudes are shown in figure 8 assuming $t = 3$.

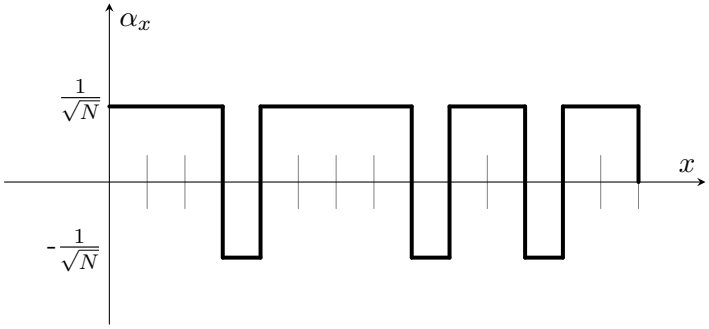

Figure 8: The state of the system after a phase kickback on all states where $f(x) = 1$. In this example, there were three states affected, which were reflected across the x-axis.

Recall from our discussion on phase kickback in the previous lecturre that this is exactly the unitary operator U_f with an ancilla qubit in the $|-\rangle$ state.

Flip about the average: $R_{average}$. This operator $R_{average}$ flips all the amplitudes about the average. Applying $R_{average}$ to figure 9 and assuming that t is small compared to N, the mean is slightly below $\frac{1}{\sqrt{N}}$. For a good state, the operator results in an ampligude of approximately $\frac{3}{\sqrt{N}}$. For a bad state, it results in an amplitude slightly below the mean as shown in figure 9.

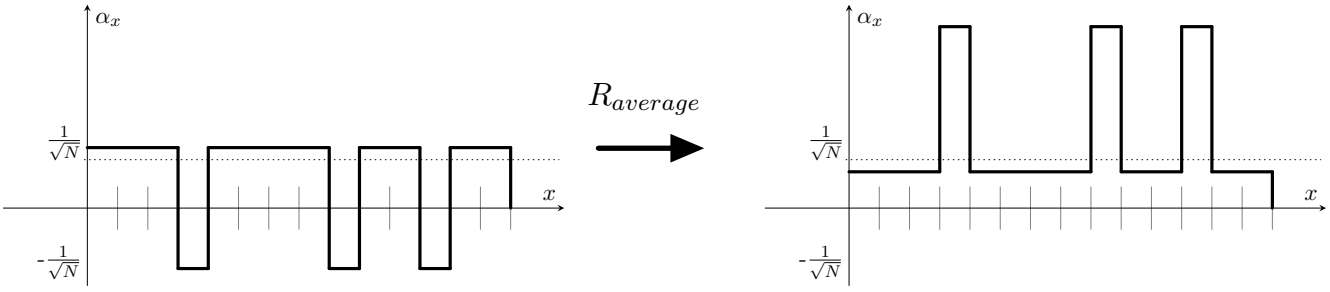

Figure 9: The state of the system after being reflected across the average, which is indicated by a dotted line. Note that the states where $f(x) = 1$ are now more probable if a measurement is taken.

Proposition 1. *$R_{average}$ is unitary.*

Proof. To prove $R_{average}$ is unitary, we must show that it is both linear and 2-norm preserving. To show $R_{average}$ is linear, we can consider how it might be implemented. Note that reflecting about the average is equivalent to subtracting the average, reflecting about the x-axis, and adding the average back. Formally,

$$R_{average}\,|\psi\rangle = -\left(|\psi\rangle - AVG\,(\alpha_x)\sum_x |x\rangle\right) + AVG(\alpha_x)\sum_x |x\rangle \tag{3}$$

where

$$AVG(\alpha_x) = \frac{1}{N}\sum_x \alpha_x. \tag{4}$$

This is clearly linear in a_x, as $AVG(\alpha_x)$ is simply a linear combination of the a_x's and all of the operators are linear.

To show $R_{average}$ is 2-norm preserving, we show that all the eigenvalues of $R_{average}$ have a magnitude of 1. First, consider what happens when we apply $R_{average}$ to the initial state shown in figure 7. Nothing happens as the reflection across the average transforms this state to itself. Thus, the uniform distribution is an eigenvector with an eigenvalue of 1.

Now consider the case shown in figure 10. On the left, we have a system where the average is zero, and after applying $R_{average}$, we have the system mirrored across the x-axis. Thus, this state is another eigenvector and the eigenvalue is -1. All the eigenvectors orthogonal to the uniform distribution have an average amplitude of zero, and thus are states that $R_{average}$ reflects about the x-axis. Therefore, all eigenvalues are either 1 or -1. In fact, $R_{average}$ is a reflection about the unique axis given by the eigenvectors corresponding to the eigenvalue 1, namely $\sum_x |x\rangle$. □

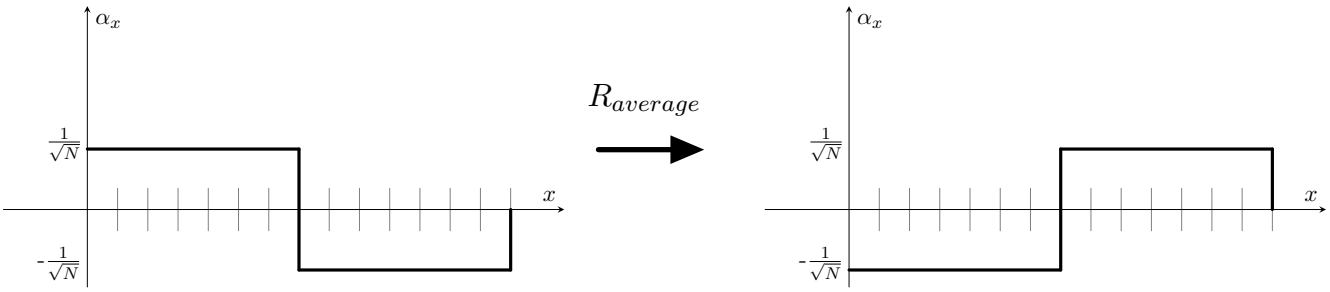

Figure 10: The state of the system before applying $R_{average}$ is on the left, and the system afterwards is on the right. As the average is zero, the system is merely reflected across the x-axis.

The physical implementation of $R_{average}$ can be done in three steps as follows

1. Bring the average axis (dash-line) to the x-axis by applying $H^{\otimes n}$.
2. Flip the phase of all components about the x-axis: $R_{|0^n\rangle}$
3. Bring the average axis back from the x-axis by applying the inverse of step 1: $(H^{\otimes n})^{-1} = H^{\otimes n}$

In summary, the flip about the average is implemented as: $R_{average} = H^{\otimes n} R_{|0^n\rangle} H^{\otimes n}$

4 Quantum Algorithm

By repeatably applying R_{bad} and $R_{average}$, the amplitudes of good states can be amplified after each iteration. However, the mean of the states will slightly decrease each iteration. As soon as the mean reaches a negative value, applying more iterations of R_{bad} and $R_{average}$ will in fact decrease the probability of finding a good item. Therefore, the number of iterations k is critical and must be carefully determined. Note that both R_{bad} and $R_{average}$ are unitary operators which can be represented as rotations. In general, repeatably applying unitary operators will result in an approximately cyclical behavior.

The full algorithm is thus as follows:

1. Create a uniform superposition of all possible x.

$$|\psi_0\rangle = H^{\otimes} |0^n\rangle = \frac{1}{\sqrt{N}} \sum_x |x\rangle \tag{5}$$

2. Apply $R_{average} R_{bad}$ k times.
3. Measure $|\psi_k\rangle$ and output the observed x.

Quantum circuit. We are now ready to describe the quantum circuit that implements Grover's algorithm. We can repeat the combination of R_{bad} and $R_{average}$ as many times as desired. The full

circuit is shown below.

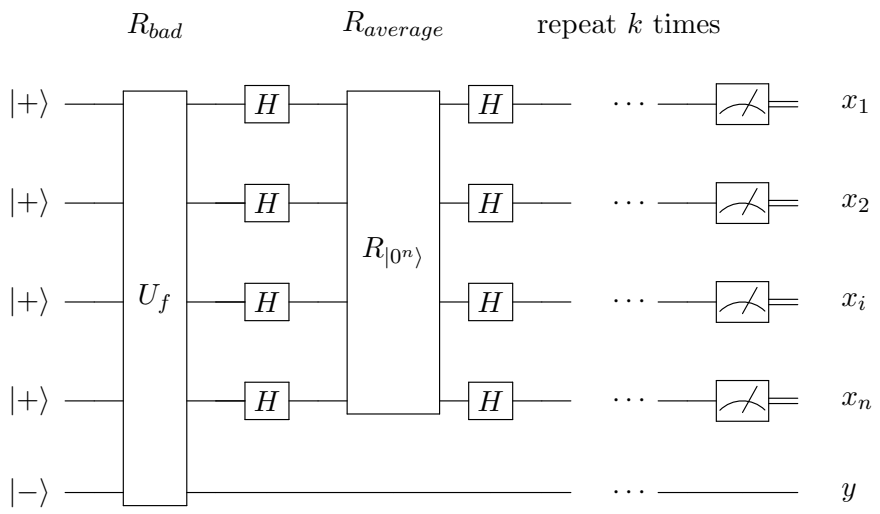

Two-dimensional state representation. We now seek to determine the optimal value of k, where k is the number of iterations of $R_{average}R_{bad}$. The amplitude α_x of $|x\rangle$ at any point in time depends only whether $f(x) = 0$ or $f(x) = 1$. That is, at any point in time, the amplitude of all the bad x's are equal and the amplitude of all the good x's are equal. Moreover, the amplitudes remain real throughout the computation. Thus, we can write the state after the i-th iteration as

$$|\psi_i\rangle = \beta_i |B\rangle + \gamma_i |G\rangle , \tag{6}$$

where $|B\rangle \doteq \frac{1}{\sqrt{N-t}} \sum_{x:f(x)=0} |x\rangle$ denotes the uniform superposition over all bad items and $|G\rangle \doteq \frac{1}{\sqrt{t}} \sum_{x:f(x)=1} |x\rangle$ the uniform superposition over all the good items. Moreover, $\beta_i, \gamma_i \in \mathbb{R}$ and are constrained by

$$\beta_i^2 + \gamma_i^2 = 1. \tag{7}$$

This constraint allows us to write β_i and γ_i in terms of trigonometric functions: $\beta_i = \cos(\theta_i)$ and $\gamma_i = \sin(\theta_i)$ for some $\theta_i \in \mathbb{R}$, i.e.,

$$|\psi_0\rangle = \cos(\theta_0) |B\rangle + \sin(\theta_0) |G\rangle . \tag{8}$$

We can thus describe the system as a two-dimensional system with parameters β and γ, where (β, γ) lie on the unit circle, as shown in figure 11. Here we plot the initial state with β_0 on the B axis and γ_0 on the G axis. We have that $\sin(\theta_0) = \gamma_0 = \sqrt{\frac{t}{N}}$. Note that $\theta_0 \geq \sin(\theta_0)$ always holds, and that $\theta_0 \approx \sin(\theta_0)$ when t is small relative to N.

Analysis. Given some point (β, γ) on this unit circle, what will the effect of the R_{bad} and $R_{average}$ operators be on this point? Since R_{bad} is a phase flip for the good components, it transforms (β, γ) by

$$(\beta, \gamma) \xrightarrow{R_{bad}} (\beta, -\gamma) \tag{9}$$

which is simply a reflection across the B-axis. $R_{average}$ reflects the point across the line defined by the origin and the point (β_0, γ_0). Taken together, these two reflections form a rotation of $2\theta_0$

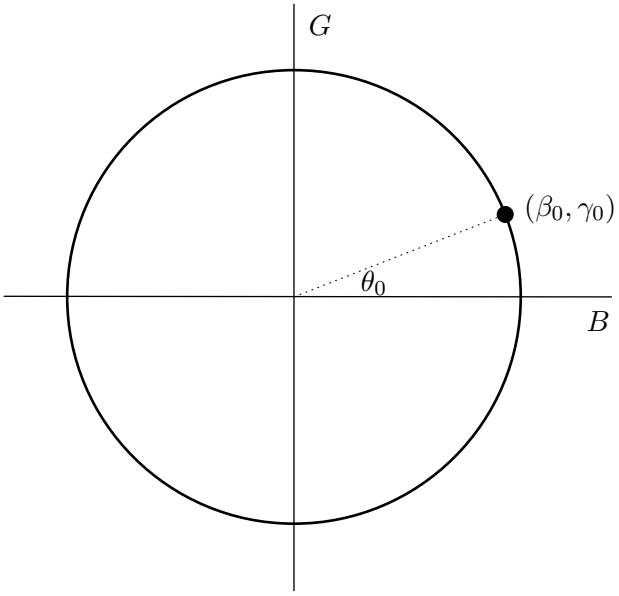

Figure 11: β and γ can be mapped on a unit circle, with β on the B axis and γ on the G axis.

counterclockwise. It follows that after i iterations,

$$\begin{aligned} \theta_i &= (2i+1)\theta_0, \\ \beta_i &= \cos(\theta_i), \\ \gamma_i &= \sin(\theta_i). \end{aligned}$$

From looking at the unit circle, the best time to make a measurement is when (β, γ) is on or very close to the G-axis, as that is when the amplitudes of the valid states are highest. It follows that the ideal value of k would satisfy

$$(2k+1)\theta_0 = \frac{\pi}{2} \tag{10}$$

which leads to

$$k = k^* \doteq \frac{1}{2}\left(\frac{\pi}{2\theta_0} - 1\right). \tag{11}$$

Note that this may not be possible because k^* may not be an integer. We choose for k the integer value closest to k^*:

$$k = \lceil k^* \rfloor \doteq \left\lceil \frac{1}{2}\left(\frac{\pi}{2\theta_0} - 1\right) \right\rfloor \tag{12}$$

We claim that this choice of k guarantees that

$$\Pr\left[\text{measure } x \in f^{-1}(1)\right] \geq \frac{1}{2}. \tag{13}$$

We know this as k must bring us to a point within the top quarter of the unit circle, as illustrated in Figure 12. The figure shows an example where we have applied $R_{average}R_{bad}$ twice, which brings us into the shaded part the of unit circle. Each application of $R_{average}R_{bad}$ rotates us by $2\theta_0$ counterclockwise, and there is no value of $\theta_0 < \pi/2$ that will allow us to completely jump over the shaded area when applying $R_{average}R_{bad}$. The advantage of being in the shaded area is that, in terms of absolute value, the amplitudes of the good states exceed the amplitudes of the bad states in absolute value, thus giving us a probability of at least $1/2$ of finding a good x when taking a measurement.

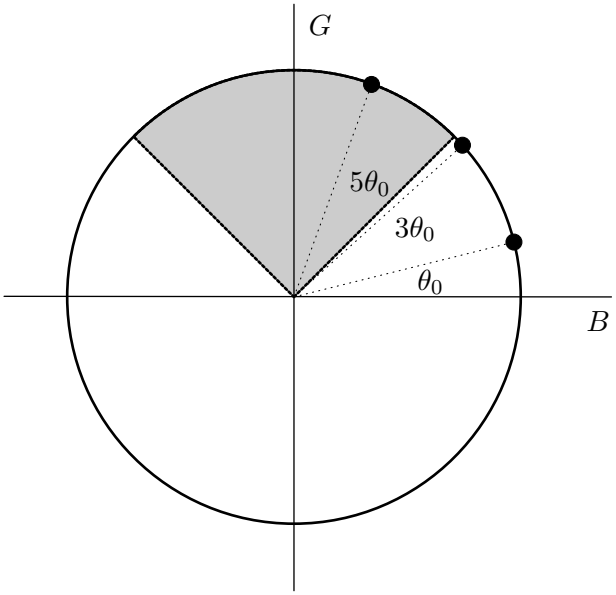

Figure 12: Optimal choice of k.

Given that $\theta_0 \geq \sin(\theta_0) = \sqrt{\frac{t}{N}}$, it follows that some

$$k = \Theta\left(\sqrt{\frac{N}{t}}\right) \tag{14}$$

iterations guarantee a probability of success of at least $1/2$, and we know the value of k provided we know t.

5 Generalization

In the previous section, we considered a search problem starting from a uniform superposition of states and discussed Grover's algorithm. In this section, we generalize the search problem where we start from a superposition state $|\psi\rangle$ which is not necessarily uniform but can be generated by applying a unitary A to the basis state $|0^n\rangle$.

Problem statement. We are again given access to a black-box function $f : \{0,1\}^n \to \{0,1\}$. We are now provided with a unitary A on n qubits such that $|\psi_0\rangle = A|0^n\rangle$ has amplitude $\alpha_x > 0$ for some good x. Finally, we assume for now that the weight p of the good x's is given where $p = \sum_{x:f(x)=1} |\alpha_x|^2$.

Our goal is the same as before and we would like to find some $x \in \{0,1\}^n$ such that $f(x) = 1$. Note that the previous problem statement was a special case of this new statement where $A = H^{\otimes n}$

Algorithm. For the algorithm construction, we start with the initial state $|\psi_0\rangle = A|0^n\rangle$ and apply k iterations of $R_{average}R_{bad}$.

The generalized $R_{average}$, however, must be adapted to our generalized case. We can generalize the $R_{average}$ from the special case where $A = H^{\otimes n}$ to get

$$R_{average} = AR_{|0^n\rangle}A^{-1} \tag{15}$$

The quantum circuit for this general case is given by

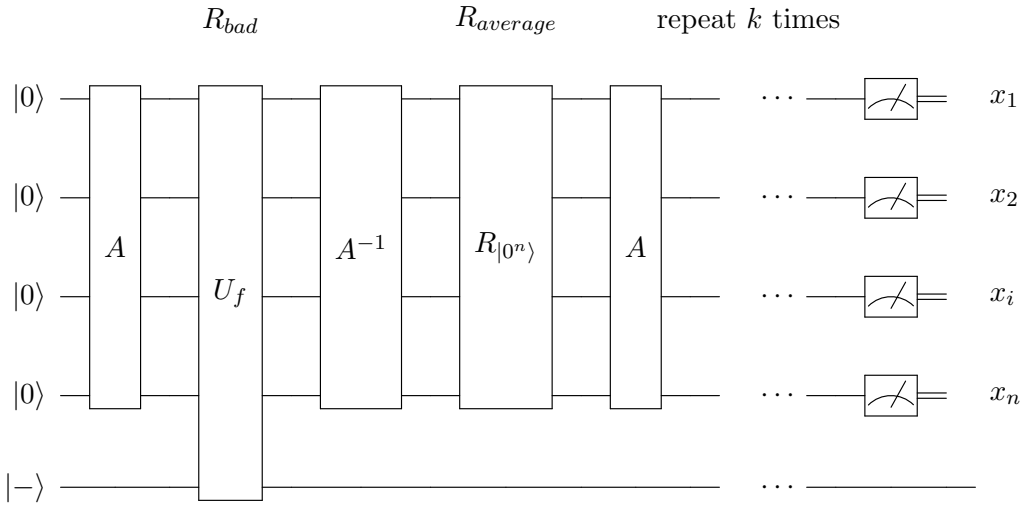

If $A = H^{\otimes n}$, the circuit is identical to the circuit in section 4.

Analysis. To determine the optimum k iterations, we must consider the initial angle θ_0. We know that the probablity of measuring a good item in $|\psi_0\rangle$ is p. Therefore,

$$\sin(\theta_0) = \sqrt{p} \tag{16}$$

Following the derivation in the previous section, the optimum number of iterations is $k = \lceil k^* \rfloor \doteq \lceil \frac{1}{2}(\frac{\pi}{2\theta_0} - 1) \rfloor = O(\frac{1}{\sqrt{p}})$. After exactly that many iterations, we get the same guarantee shown in (13), and we can compute the number provided we know p.

6 Error Elimination

To eliminate all errors after applying $k = \lceil k^* \rfloor$ iterations of $R_{average}R_{bad}$, the final state must be positioned exactly at $\frac{\pi}{2}$. This happens if and only if $k^* \in \mathbb{Z}$. The idea to arrive at this situation in general is to slightly manipulate the angle θ_0 to $\tilde{\theta}_0$ such that after about k^* iterations, the final state lands exactly at $\frac{\pi}{2}$. This can be achieved by using an additional ancilla qubit to reduce the probability of measuring a good item slightly from p to $\tilde{p}$, where $\tilde{p}$ yields an integral value for $\tilde{k}^*$ as shown in figure 13. The new good states are the ones where the ancilla qubit is in state $|1\rangle$.

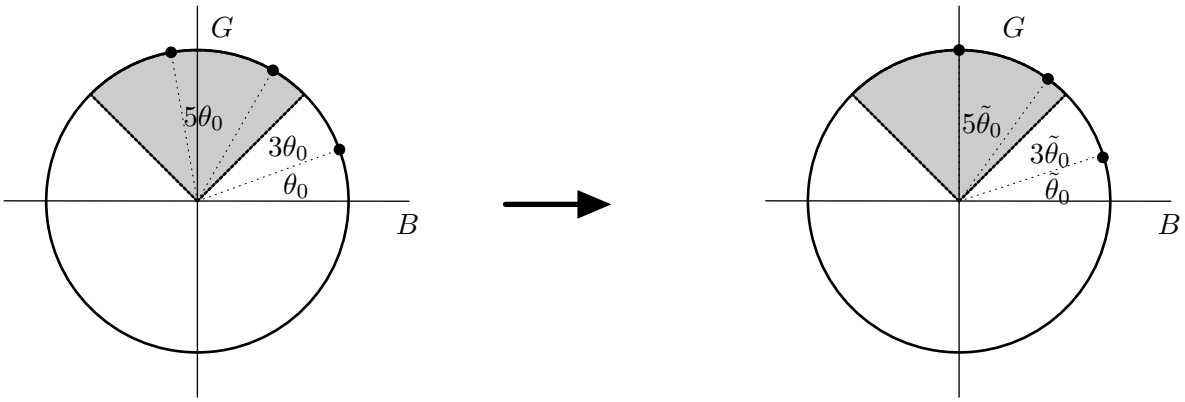

Figure 13: The left figure shows the starting state, $\theta_0 = \sqrt{p}$, and the state after first iteration and second iterations. Note that the final state is not exactly at the ideal stopping point of $\frac{\pi}{2}$. In the right figure, by slightly decreasing θ_0 to $\tilde{\theta}_0$, the final state is exactly at $\frac{\pi}{2}$.

Exercise #8. Work out a way to eliminate the error in quantum search using the above ideas, without increasing the number of queries by more than one. Hint: Use the generalization shown in section 5 with $\tilde{A} = U \otimes A$ for some U.