1 Readings

Literature:

Grover's algorithm and amplitude amplification: quant-ph/9605043

Diffusion transform and other motivations from physics: Grover, quant-ph/0109116

Quantum bomb detection: Elitzur and Vaidman, Vistas in Astronomy 37, 253 (1993); Found. of Physics 23, 987 (1993); Vaidman, Found. of Physics 33, 491 (2003); Kwiat et al. PRL 74, 4763 (1995); Rudolph and Grover, quant-ph/0206066

2 Amplitude amplification in Grover search

The Diffusion operator D has two properties:

1. It is unitary and can be efficiently realized.
2. It can be seen as an "inversion about the mean."

We discussed the first property in the previous lecture. We now analyze the second property.

For $N = 2^n$, we have

$$D = -I + 2\left|\psi_0\right\rangle\left\langle\psi_0\right|$$

where $\left|\psi_0\right\rangle$ is the zero vector in the Hadamard basis. We saw last time that D can be decomposed as:

$$
\begin{aligned}
D &= H_N \begin{pmatrix} +1 & 0 & \cdots & 0 \\ 0 & -1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & -1 \end{pmatrix} H_N \\
&= H_N \left(\begin{pmatrix} +2 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} + I \right) H_N \\
&= H_N \begin{pmatrix} +2 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} H_N - I \\
&= \begin{pmatrix} -2/N & -2/N & \cdots & -2/N \\ -2/N & -2/N & \cdots & -2/N \\ \vdots & \vdots & \ddots & \vdots \\ -2/N & -2/N & \cdots & -2/N \end{pmatrix} - I \\
&= \begin{pmatrix} +2/N-1 & +2/N & \cdots & +2/N \\ +2/N & +2/N-1 & \cdots & +2/N \\ \vdots & \vdots & \ddots & \vdots \\ +2/N & +2/N & \cdots & +2/N-1 \end{pmatrix}
\end{aligned}
$$

The indexing here is such that the first state (top left hand corner of the matrices) is the target state $|a\rangle$. Note that the central matrix in D is a conditional phase shift matrix, i.e., it puts a phase shift in front of all states except the target.

Consider the action of D on a vector $|\alpha\rangle$ to generate another vector $|\beta\rangle$:

$$
D \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_i \\ \vdots \\ \alpha_N \end{pmatrix} = \begin{pmatrix} \beta_1 \\ \vdots \\ \beta_i \\ \vdots \\ \beta_N \end{pmatrix}
$$

Define $\mu = \sum_i \alpha_i / N$ as the mean amplitude Then

$$
\begin{aligned}
\beta_i &= \frac{2}{N} \sum_j \alpha_j - \alpha_i \\
&= 2(\mu - \alpha_i) \\
&= \mu + (\mu - \alpha_i)
\end{aligned}
$$

which corresponds to a reflection of α_i about the mean value μ. This is illustrated in Figure 1 below (note that if we start from the uniform superposition, the non-target states will have equal amplitude, we have just

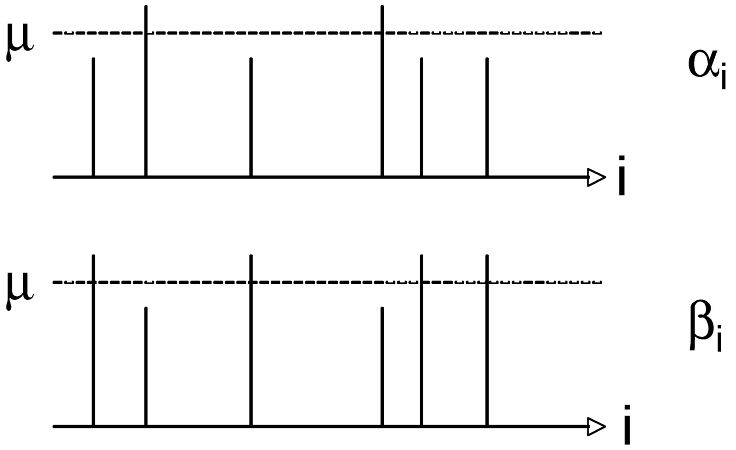

Figure 1: Inversion of amplitudes α_i about their mean value μ.

illustrated the principle for a general state here). Thus, the amplitude of $\beta_i = -\frac{2}{N}\sum_j \alpha_j + \alpha_i = -2\mu + \alpha_i$ can be considered an "inversion about the mean" with respect to α_i. Now if we first change the sign of the amplitude of the target state, by applying the oracle as in the previous lecture, the target state is now significantly further away from the mean. The inversionabout the mean further amplifies this, as shown in Figure 2 below.

This shows how quantum search algorithm iteratively improves the probability of measuring a solution by increasing the component of the target state at each iteration. The overall procedure is summarized as follows

1. Start state is $|\psi_0\rangle = \sum_x \frac{1}{\sqrt{N}}|x\rangle$
2. Invert the phase of $|a\rangle$ using f
3. Then invert about the mean using D
4. Repeat steps 2 and 3 $O(\sqrt{N})$ times, so in each iteration α_a increases by $\frac{2}{\sqrt{N}}$

Suppose we just want to find a with probability $\frac{1}{2}$. Until this point, the rest of the basis vectors will have amplitude at least $\frac{1}{\sqrt{2N}}$. In each iteration of the algorithm, α_a increases by at least $\frac{2}{\sqrt{2N}} = \sqrt{\frac{2}{N}}$. Eventually, $\alpha_a = \frac{1}{\sqrt{2}}$. The number of iterations to get to this α_a is $\leq \sqrt{N}$.

2.1 Diffusion transform

See the discussion of motivation and form of matrix D by Grover in his article quant-ph/0109116.

2.2 Applications of quantum search

Grover's algorithm is often called a "database" search algorithm, where you can query in superposition. Other things you can do with a similar approach:

1. Find the minimum.
2. Approximately count elements, or generate random ones.

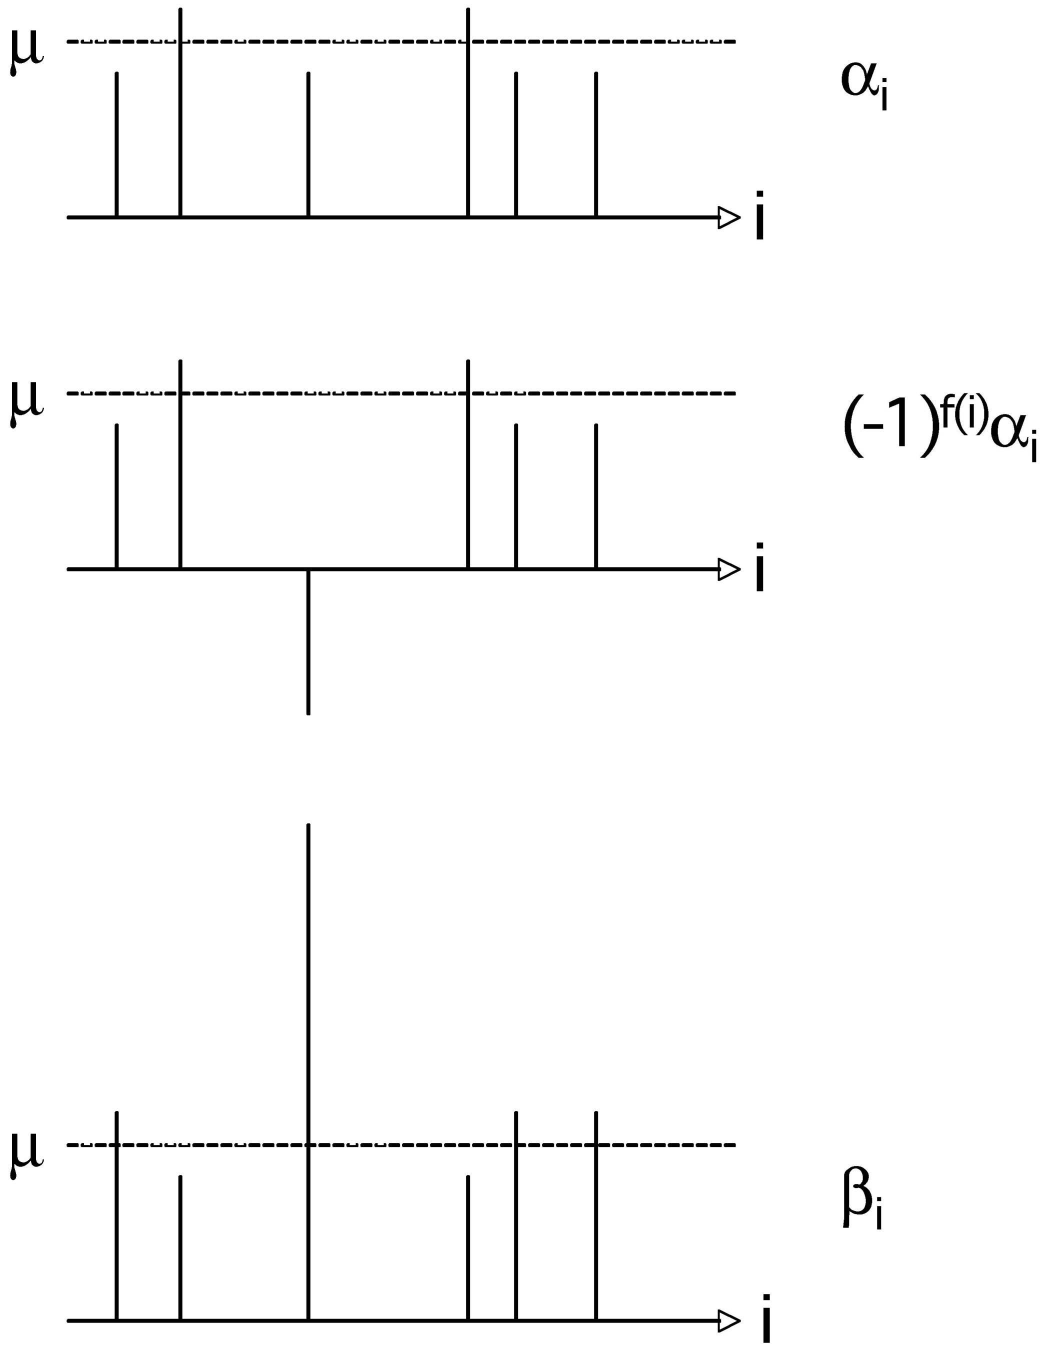

Figure 2: Application of oracle to invert sign of target state followed by Inversion of amplitudes α_i about their mean value μ gives rise to amplification of the target component.

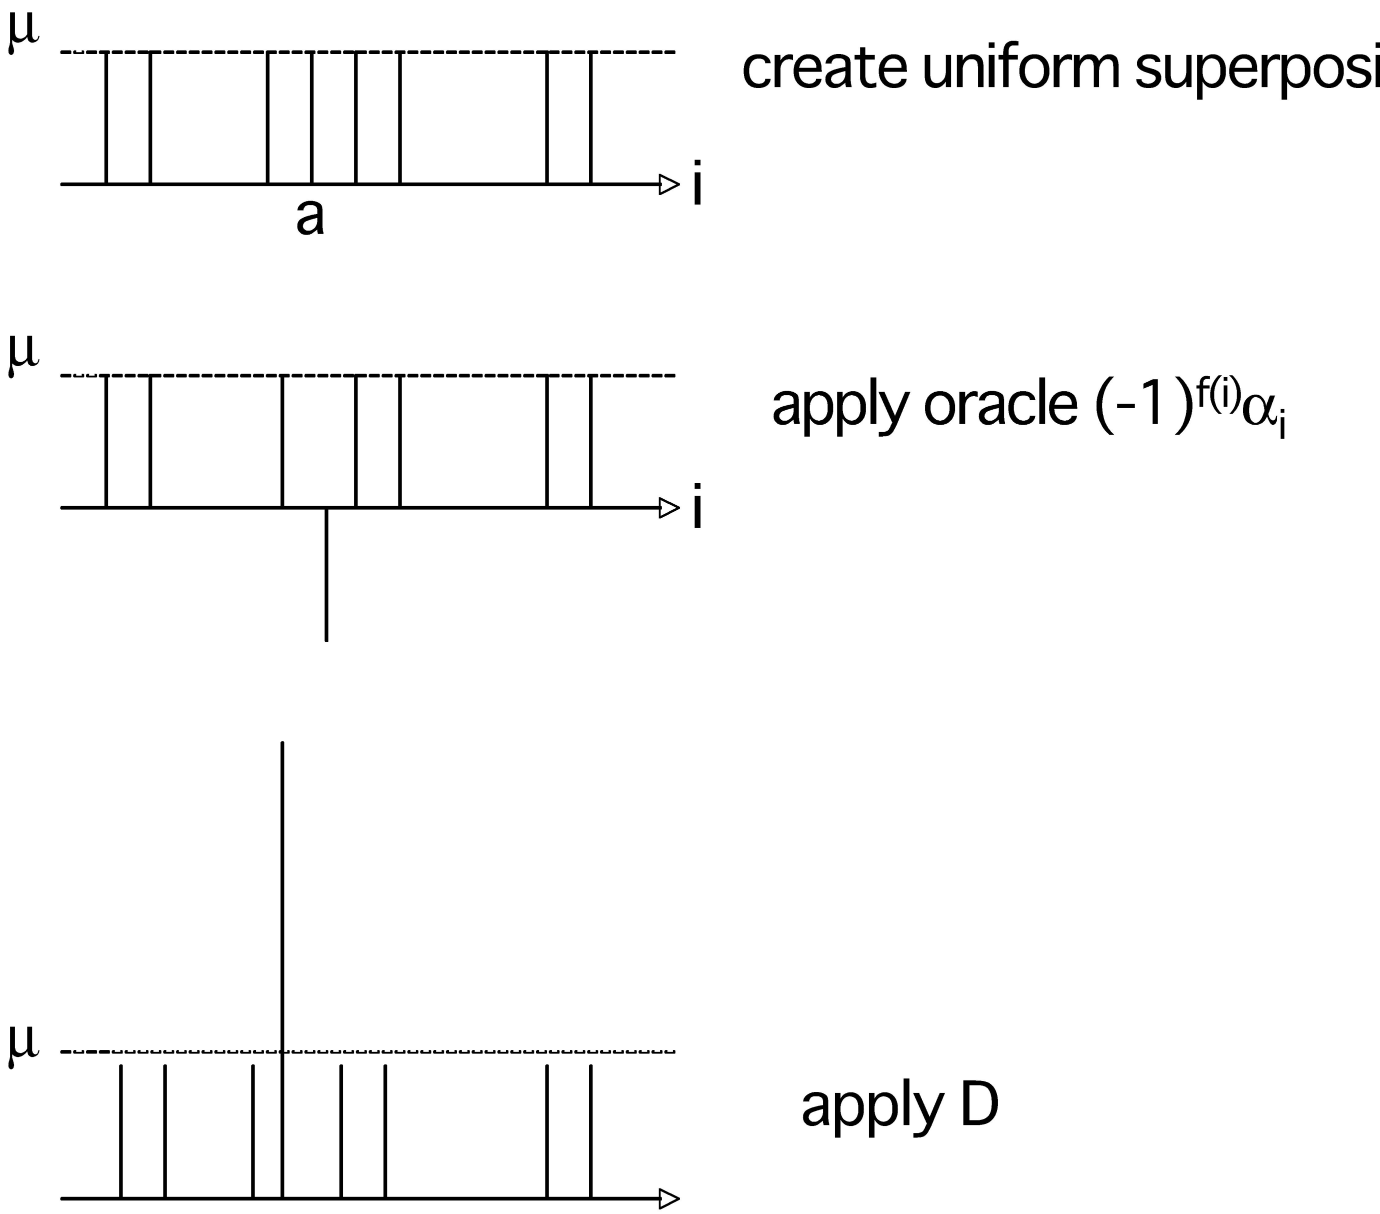

Figure 3: The first three steps of Grover's algorithm. We start with a uniform superposition of all basis vectors in the top panel. In the middle panel we have used the function f to invert the phase of α_k. After running the diffusion operator D in the bottom panel, we have amplified α_k while decreasing all other amplitudes.

3. Speed up the collision problem.

4. Speed up the test for matrix multiplication. In this problem we are given three matrices, A, B, and C, and are told that the product of the first two equals the third. We wish to verify that this is indeed true. An efficient (randomized) way of doing this is picking a random array r, and checking to see whether $Cr = ABr = A(Br)$. Classically, we can do the check in $O(n^2)$ time, but using a similar approach to Grover's algorithm we can speed it up to $O(n^{1.75})$ time.

5. Speedup exhaustive search in NP-complete problems, although this alone is not enough to provide efficient solution. See Ambainis, quant-ph/0504012 for a review of applications to NP-complete problems.

3 Quantum bomb detection

To illustrate some of the concepts behind Grover's algorithm, we can consider a problem known as Vaidman's bomb. In this problem, we have a package that may or may not contain a bomb. However, the bomb is so sensitive that simply looking to see if the bomb exists will cause it to explode. So, can we determine whether the package contains a bomb without setting it off? Paradoxically, quantum mechanics says that we can. This is achieved by combining a technique referred to in the physics literature as 'interaction free measurement', which relies on interferometry, with amplitude amplfication. By making an iterative amplitude amplification, we can arrive at a sequence of N cycles of single qubit operations such that if the package contains a bomb, we will look (i.e., interact with it and blow up) with probability only $1/N$, while the rest of the time we have two distinct outcomes for the qubit state which tell us whether or not the bomb is present.

3.1 The Quantum Zeno Effect

We will make use of a phenomenon known as the Quantum Zeno Effect (also referred to as the "watched pot" or the "watchdog" effect, or the "hare and tortoise syndrome"). Consider a quantum state consisting of a single qubit. This qubit starts at $|0\rangle$, and at every step we will rotate it toward $|1\rangle$ by $\theta = \pi/2N$. After one rotation, we have $|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$, where where $\beta = \sin\theta \approx 1/N$. After 2 steps, we have $|\phi\rangle = \cos(2\theta)|0\rangle + \sin(2\theta)|1\rangle$, ...etc. so that after N steps we have $|\phi\rangle = \cos(N\theta)|0\rangle + \sin(N\theta)|1\rangle$. Now since $\beta = \sin\theta \approx 1/N$, then this final state after N steps will be $|1\rangle$, or very close to this, so that any measurement will then return $|1\rangle$ with high probability.

Now what if we decide to measure the state after each rotation? After the first rotation, we will measure $|0\rangle$ with high probability, *but this measurement collapses the state back to* $|0\rangle$. Thus, each measurement has a high probability of yielding $|0\rangle$; the probability of getting $|1\rangle$ by the end is approximately $N\frac{1}{N^2} = \frac{1}{N}$, as opposed to the extremely high probability in the previous case.

Essentially, the Quantum Zeno Effect says that if we have a quantum state that is in transition toward a different state, making frequent measurements can delay that transition by repeatedly collapsing the qubit back to its original state.

This is another kind of amplitude amplification - in Grover's algorithm our amplification was unitary, but this is not.

3.2 Looking for the Bomb

To determine whether Vaidman's bomb exists without actually looking at it, we want to take advantage of the Quantum Zeno Effect. Figure 4 shows a scheme for the interaction free detection with photonic qubits,

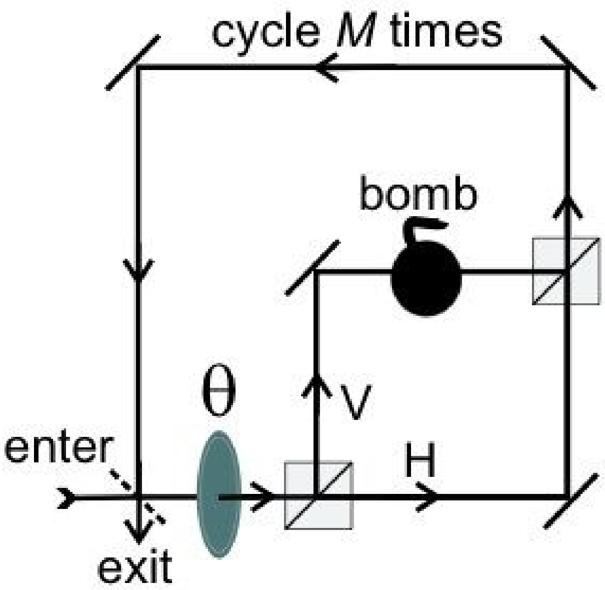

Figure 4: Interaction free detection of a single bomb.

labelled by their polarization states $|0\rangle \equiv |H\rangle$ and $|1\rangle \equiv |V\rangle$ (taken from Rudolph and Grover, quant-ph/00206066). Let us follow the action of this circuit. The initial input is $|0\rangle$, which is then rotated by the lens (denoted θ) to the state $\cos\theta|0\rangle + \sin\theta|1\rangle$, with $\theta = \pi/2N$ and N large so that $\sin\theta \sim \pi/2N$. This state is then input into a beam splitter (for spin qubits, a Stern-Gerlach magnet would be used) which sends $|0\rangle$ along the lower path and $|1\rangle$ along the upper path. These two paths are completely independent and can go, for instance, through different rooms in a building. Our protocol will be set up to establish whether there is or is not a bomb in the upper path - so one will always use the apparatus in a configuration that the upper path traverses the suspect region and the lower path goes through a known safe (no bomb) region.

First, suppose that there is no bomb present. Then the second beamsplitter combines the amplitudes from the two paths coherently and the final state is $\cos\theta|0\rangle + \sin\theta|1\rangle$. If we send this back to the lens and rotate by θ again, then input into the beam splitter, ... etc. we get the final state $\cos 2\theta|0\rangle + \sin 2\theta|1\rangle$. So repeating the cycle N times will give the final state $|1\rangle$ as described above. So if there is no bomb and we cycle N times before measuring the qubit, we will find $|1\rangle$ to certainty or a very high probability.

Now, suppose that there is a bomb present on the upper path. Then in the first cycle there are two possible outcomes. With probability $\sin^2\theta$ the qubit passes the bomb and causes an explosion. With probability $\cos^2\theta$ the qubit goes through the lower path, does not see the bomb, and emerges as $|0\rangle$. Thus the presence of the bomb is like a quantum tortoise, which forces a quantum Zeno effect - but note that in this interferometric situation the bomb does not need to actually be measured... hence the name 'interaction free measurement'. Now the probability that the qubit emerges unscathed as $|0\rangle$ from the second cycle is then $\cos^2\theta\cos^2\theta$, while the probability that the bomb explodes on the second cycle if $\cos^2\theta sin^2\theta$. You can thus see how to continue this to get the distribution of probabilities for all possible results after N cycles: the probability for no explosion in any cycle is $\cos^{2N}\theta$ and the qubit emerges as $|0\rangle$, and the probability for having the bomb explode in any one cycle between 1 and N is $1-\cos^{2N}\theta$. In the latter case there is no qubit....

So we have three possible outcomes after making N cycles with our single qubit starting as $|0\rangle$:

1. no bomb: final state is $|1\rangle$ with (near) certainty

2. bomb present: no explosion, final state is $|0\rangle$ with probability $\cos^{2N}\theta$
3. bomb present: explosion happened, no final state..., probability $1-\cos^{2N}\theta$

What are these probabilities? For $\theta = \pi/2N$ and N large, we have $\cos^{2N}\theta \sim 1-\pi^2/4N \sim 1$ and $1-\cos^{2N}\theta \sim \pi^2/4N$. So we have very effectively reduced the probability of exploding the bomb to $\sim 1/N$ by this combination of quantum Zeno and cycling with a qubit interferometer.

The approach can also be modified to the particularly unfortunate case where we have N packages, $N-1$ of which contain bombs. We want to find the one package that does not contain a bomb, though we don't mind setting off a few of the bombs in the process. This can be done by a modification of the above interferometric scheme, as described in quant-ph/0206066. Note that here the amplitude amplification occurs by the implicit (deferred) measurement in the interferometer. In contrast, in Grover's algorithm the amplitude of one target basis vector is amplified while all others are constantly diminished or reset in a unitary manner.

One important thing that the N bomb example illustrates about quantum search is that it's highly counter-intuitive to be able to search in $\sqrt{N}$ steps. By querying in superposition, we manage to search using fewer steps than there are locations to search!