Ph/CS 219A
Quantum Computation

Lecture 16. Quantum Searching

We've seen that in the black box setting, we can speed up Period Finding exponentially by allowing quantum queries, and that, using this black box algorithm, we can also vastly speed up finding the periods of efficiently computable functions (like the modular exponential function).

Other exponential quantum speedups are known, and we will consider another important example soon: quantum simulation.

But today's lecture is about a quantum speedup in the black box setting that is ``merely'' polynomial (quadratic in fact). Though far less spectacular than exponential speedups, polynomial quantum speedups are also interesting and might possibly be of practical interest as well.

An updated version of the Chapter 6 Lecture Notes has been posted on the course website.
Problems Set 4 is due on December 4.

Exhaustive Search

For an NP-hard combinatorial problem, there is no obvious structure that can be exploited to greatly speedup the search for a solution. We may not be able to do much better than exhaustive search for a solution.

To formalize this notion in the black box setting, suppose that the function accepts only one possible n-bit input *w* (the "marked string").

$$f_w : \{0,1,2,3,\ldots,N-1\}^n \to \{0,1\}; \quad f_w(x) = \begin{cases} 0 & x \neq w \\ 1 & x = w \end{cases}$$

Classically, we need to query more than $N/2$ times to find *w* with success probability > ½.
Quantumly, we can find *w* using only $O(\sqrt{N})$ queries, a "quadratic speedup" using *Grover's algorithm*.

We can apply this method to any problem in NP, where *f* is the efficiently computable verifier.
Remarkably, by querying in superposition, we can interrogate *N* potential witnesses in time $O(\sqrt{N})$.

Turn the black box into a "phase oracle":

$$U_w : |x\rangle \otimes |y\rangle \mapsto |x\rangle \otimes |y \oplus f_w(x)\rangle \quad \Rightarrow \quad U_w : |x\rangle \otimes |-\rangle \mapsto (-1)^{f_w(x)} |x\rangle \otimes |-\rangle, \quad (-1)^{f_w(x)} = \begin{cases} 1 & x \neq w \\ -1 & x = w \end{cases}$$

$$U_w = I - 2|w\rangle\langle w| \quad \Rightarrow \quad U_w : |\psi\rangle = a|w\rangle + b|\psi^\perp\rangle \mapsto -a|w\rangle + b|\psi^\perp\rangle.$$

The phase query reflects a state vector in the hyperplane orthogonal to $|w\rangle$.

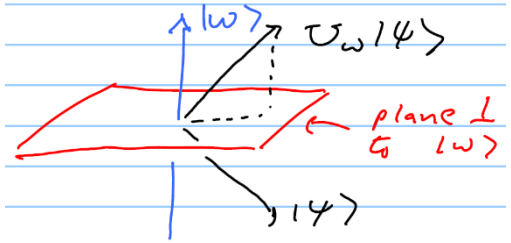

Grover's algorithm

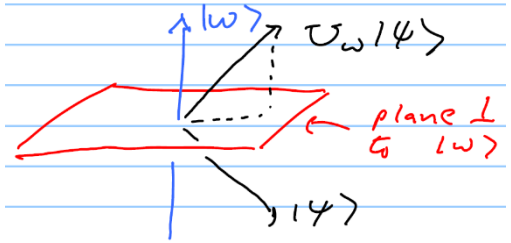

$$U_w = I - 2|w\rangle\langle w|$$

The phase query reflects a state vector in the hyperplane orthogonal to $|w\rangle$.

(1) Prepare $|s\rangle = \frac{1}{\sqrt{N}}\sum_{x=1}^{N-1}|x\rangle \quad \Rightarrow \quad \langle w|s\rangle = \frac{1}{\sqrt{N}} = \sin\theta, \quad \theta \approx \frac{1}{\sqrt{N}}.$

(2) Apply many times the "Grover iteration" $U_{\text{Grover}} = U_s U_w, \quad U_s = 2|s\rangle\langle s| - I.$

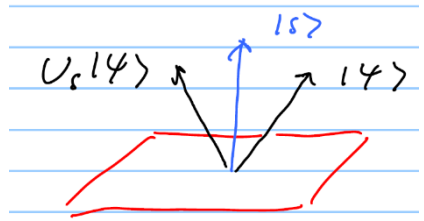

U_w is the query, U_s we apply ourselves. It reflects a state vector in the axis defined by $|s\rangle$.

(3) Measure. But how many times to apply U_{Grover} ?

Action of U_{Grover} in the plane spanned by $|s\rangle$ and $|w\rangle$:

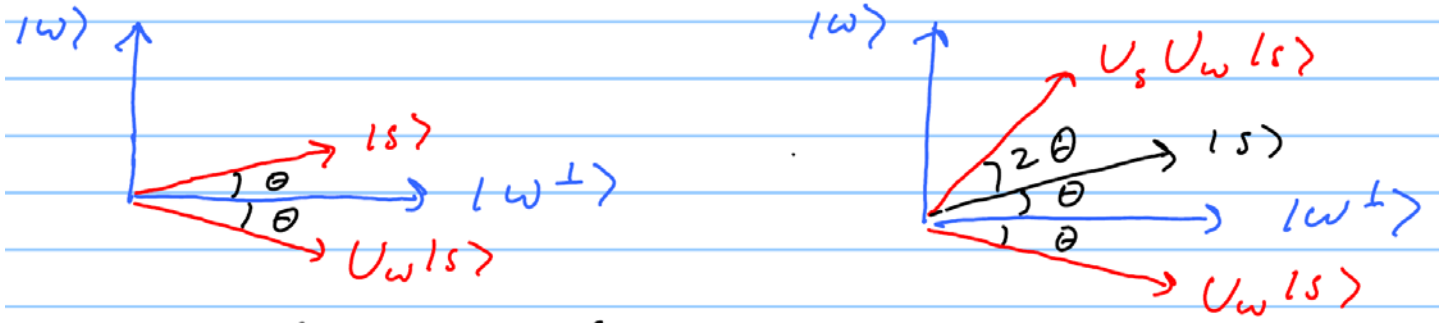

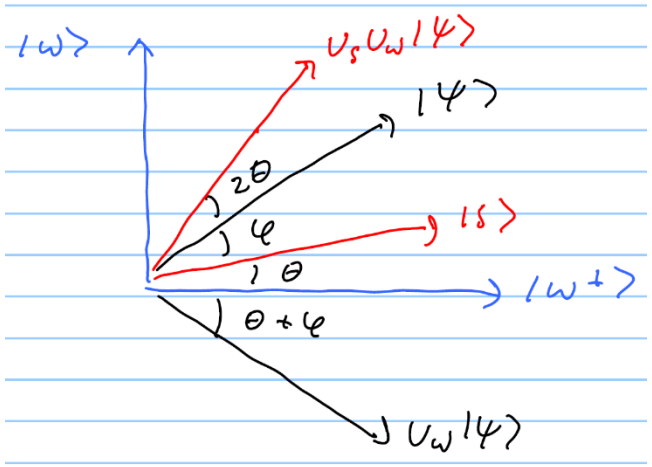

Each Grover iteration rotates the state vector counterclockwise by angle 2θ.

Grover's algorithm

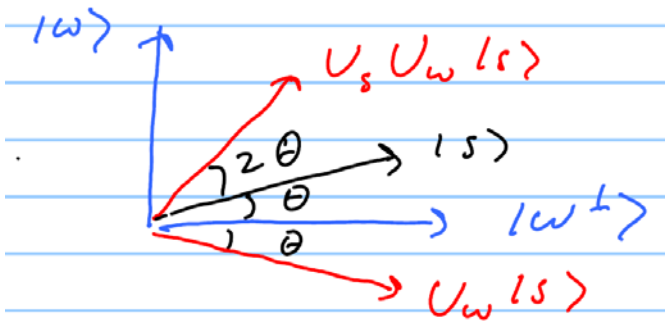

$$\langle w|s\rangle = \frac{1}{\sqrt{N}} = \sin\theta, \quad \theta \approx \frac{1}{\sqrt{N}}.$$

Each Grover iteration rotates the state vector counterclockwise by angle 2θ.

After T such iterations, where $(2T+1)\ \theta \approx \pi/2$, the state vector lines us with w, apart from a misalignment no larger than θ. Measuring in the standard basis, we find the string w with probability ≈ 1.

Number of iterations, and hence number of queries to the black box, is: $T \approx \frac{\pi}{4\theta} \approx \frac{\pi}{4}\sqrt{N}$.

It works because probabilities are *squares* of amplitudes. This means that, while t classical queries increase the success probability linearly in t, t quantum queries increase the success probability *quadratically* in t.

What if the number of marked strings is $r > 1$ instead of just 1? Finding a marked string is still hard, if $r \ll N$. Classically, we need $O(N/r)$ queries to find any marked string with probability $O(1)$. Quantumly:

$$|\text{marked}\rangle = \frac{1}{\sqrt{r}}\sum_{i=1}^{r}|w_i\rangle \quad \Rightarrow \quad \langle \text{marked}|s\rangle = \sqrt{\frac{r}{N}} = \sin\theta$$

A state near $|\text{marked}\rangle$ is obtained after T iterations, $(2T+1)\ \theta \approx \pi/2$; then measurement samples uniformly from marked states.

$$\Rightarrow \quad T \approx \frac{\pi}{4\theta} \approx \frac{\pi}{4}\sqrt{\frac{N}{r}} \quad \Rightarrow \quad \text{again } \#\text{quantum queries} = O(\sqrt{\#\text{classical queries}}).$$

Grover's algorithm

What if we don't have a priori knowledge of *r*, the number of "satisfying assignments"?

The success probability oscillates as a function of the number of queries, with period $\approx \frac{\pi}{2}\sqrt{\frac{N}{r}}$.

Choose the number of queries *t* by sampling from {1, 2, 3, …, *T*}, where $T \approx \frac{\pi}{4}\sqrt{N}$.

If solutions exists, one is found with probability ≈ ½ in each trial.

For some problems, there are better classical methods than unstructured exhaustive search. We can exploit the structure of the problem to restrict the search for a solution to a smaller set of candidate solutions, e.g. $\{x_1, x_2, x_3, \ldots, x_M\}$. Suppose we can construct an efficient unitary *U* such that

$U|0\rangle = \frac{1}{\sqrt{M}}\sum_{i=1}^{M}|x_i\rangle$. For the Grover iteration, use $U\left(2|0\rangle\langle 0| - I\right)U^\dagger$, which is also efficient.

There is a quadratic speedup, analyzed just as for Grover. O(*M*) queries classically, O($\sqrt{M}$) queries quantumly

The Grover algorithm is optimal

We've seen that Grover's algorithm achieves a quadratic speedup for exhaustive search relative to classical methods. Can a quantum computer do better than that? It turns out the answer is no.

Consider the case where there is a single marked string w, so the phase oracle performs the unitary $U_\omega = \mathrm{I} - 2\,|\,\omega\rangle\langle\,\omega\,|$. In between queries, we can perform any unitary transformation we please. After *T* queries, with arbitrary unitaries between queries, the state becomes.

$$|\psi_\omega(T)\rangle = U(\omega,T)\,|\psi(0)\rangle = U_\omega U_T U_\omega U_{T-1}\ldots U_2 U_\omega U_1\,|\psi(0)\rangle.$$

This is one of *N* possible states, one for each possible marked string. Our transformations U_1, U_2, U_3, … do not cause these *N* possible vectors to splay outward; rather they just rotate the *N* vectors while maintaining their relative orientation. Only the queries can cause the vectors to separate. If we want to be able to distinguish these *N* vectors almost perfectly, they should be very close to mutually orthogonal.

Keep track of the splaying vectors relative to a *reference state*, the case of the "empty oracle":

$$|\varphi(t)\rangle = U_t U_{t-1}\ldots U_2 U_1\,|\psi(0)\rangle,$$

$$|\psi_\omega(t)\rangle = |\varphi(t)\rangle + |E_\omega(t)\rangle.$$

How does this deviation from the reference state grow as *t* increases?

The Grover algorithm is optimal

$$|\psi_\omega(t)\rangle = U_\omega U_t U_\omega U_{t-1} \ldots U_2 U_\omega U_1 |\psi(0)\rangle = |\varphi(t)\rangle + |E_\omega(t)\rangle,$$

$$|\varphi(t)\rangle = U_t U_{t-1} \ldots U_2 U_1 |\psi(0)\rangle \quad \Rightarrow \quad |\psi_\omega(t+1)\rangle = U_\omega U_{t+1} |\varphi(t)\rangle + U_\omega U_{t+1} |E_\omega(t)\rangle.$$

How the deviation from the reference state changes in one query:

$$U_\omega = I + (U_\omega - I) \quad \Rightarrow \quad |\psi_\omega(t+1)\rangle = |\varphi(t+1)\rangle + (U_\omega - I)|\varphi(t+1)\rangle + U_\omega U_{t+1} |E_\omega(t)\rangle$$

$$= |\varphi(t+1)\rangle + |E_\omega(t+1)\rangle, \quad \text{where} \quad |E_\omega(t+1)\rangle = (U_\omega - I)|\varphi(t+1)\rangle + U_\omega U_{t+1} |E_\omega(t)\rangle.$$

How the size of the deviation grows in each query:

$$\| |E_\omega(t+1)\rangle \| \le \| |E_\omega(t)\rangle \| \quad \| (U_\omega - I)|\varphi(t+1)\rangle \| \qquad \text{(triangle inequality).}$$

$$U_\omega = I - 2|\omega\rangle\langle\omega| \quad \Rightarrow \quad \| |E_\omega(t+1)\rangle \| - \| |E_\omega(t)\rangle \| \le 2\left|\langle\omega|\varphi(t+1)\rangle\right|$$

After T such steps: $\| |E_\omega(T)\rangle \| \le 2\sum_{t=1}^{T} \left|\langle\omega|\varphi(t)\rangle\right|$

Cauchy-Schwarz: $\sum_{t=1}^{T} |c_t| \le \sqrt{T}\left(\sum_{t=1}^{T} |c_t|^2\right)^{1/2}$

$$\Rightarrow \quad \| |E_\omega(T)\rangle \|^2 = \| |\psi_\omega(T)\rangle - |\varphi(T)\rangle \|^2 \le 4T \sum_{t=1}^{T} \left|\langle\omega|\varphi(t)\rangle\right|^2$$

Upper bound on deviation.

The Grover algorithm is optimal

$$\| |E_\omega(T)\rangle . \|^2 = \| |\psi_\omega(T)\rangle - |\varphi(T)\rangle \|^2 \le 4T \sum_{t=1}^{T} \left| \langle \omega | \varphi(t) \rangle \right|^2$$

Upper bound on deviation.

Now, sum up the deviations from the reference state over all possible strings marked by the oracle:

$$\sum_{\omega=0}^{N-1} \| |\psi_\omega(T)\rangle - |\varphi(T)\rangle \|^2 \le 4T \sum_{t=1}^{T} \sum_{\omega=0}^{N-1} \left| \langle \omega | \varphi(t) \rangle \right|^2 = 4T^2 .$$

If we can identify the marked string with success probability close to 1, then these *N* states are almost mutually orthogonal. For any orthonormal basis $\{|i\rangle\}$ and any fixed vector:

$$\sum_{i=0}^{N-1} \| |i\rangle - |\varphi\rangle \|^2 = \sum_{i=0}^{N-1} \left(2 - 2\,\mathrm{Re} \langle i | \varphi \rangle \right) \ge 2N - 2 \sum_{i=0}^{N-1} |\langle i | \varphi \rangle| \ge 2N - 2\sqrt{N} = 2N\left(1 - o(N)\right).$$

$$\Rightarrow \quad 4T^2 \ge 2N\left(1 - o(N)\right) \quad \Rightarrow \quad T \ge \sqrt{\frac{N}{2}} \approx .707\sqrt{N}, \quad \text{vs.} \quad T_{\text{Grover}} \approx \frac{\pi}{4}\sqrt{N} \approx .785\sqrt{N}.$$

In fact, the lower bound can be tightened to match the upper bound on query complexity.

The Grover algorithm is optimal

What if we are willing to accept a success probability of at least (1-ε)? Then for each marked string ω we want

$$|\psi_\omega(T)\rangle = A_\omega\,|\omega\rangle + B_\omega\,|\omega^\perp\rangle, \quad |B_\omega| \le \sqrt{\epsilon} \quad \Rightarrow \quad \left|\langle\psi_\omega(T)\,|\,\varphi\rangle\right| \le \left|\langle\omega\,|\,\varphi\rangle\right| + \sqrt{\epsilon}$$

$$\Rightarrow \quad 4T^2 \ge \sum_{\omega=0}^{N-1} \||\psi_\omega(T)\rangle - |\varphi(T)\rangle\|^2 \ge 2N - 2\sum_{\omega=0}^{N-1}\left|\langle\psi_\omega(T)\,|\,\varphi\rangle\right| \ge 2N - 2\sqrt{N} - 2N\sqrt{\epsilon}$$

$$\Rightarrow \quad T \ge \sqrt{\frac{N}{2}}\left(1-\sqrt{\epsilon}-N^{-1/2}\right)^{1/2} = \Omega(\sqrt{N}).$$

The quadratic speedup for exhaustive search is optimal in the black box model.

We might regard this as indirect evidence that NP is not contained in BQP. At least we can say that the sheer magic of quantum superposition does not suffice to achieve a better-than-quadratic speedup. To do better, we would have to somehow exploiting the structure of the problem.

If there is any structure shared by all the NP-complete problems, it seems to be well hidden --- for hard instances we don't know how to improve very much on exhaustive search for a solution. If there is such structure, it would be a (delightful) surprise if it turns out to be well matched with the advantages of the quantum circuit model.