For the hardest instances of NP-hard problems, no better method is known than exhaustive search for a solution. For example, suppose that for some efficiently computable Boolean function

$$f: \{0,1\}^n \longrightarrow \{0,1\}$$

we wish to determine whether there is an x such that $f(x) = 1$. We could search for a solution by trying all of the $N = 2^n$ possible values of the input x, but that might require time $O(N \,\mathrm{polylog}\, N)$ — assuming we can evaluate f in time $O(\mathrm{polylog}\, N)$. That's very slow, but if f has no structure that we know how to exploit, we might not know how to do better.

We can model this situation in the blackbox setting. Suppose we are promised that the function evaluated by the box has the form

$$f_\omega(x) = \begin{cases} 0 & x \neq \omega \\ 1 & x = \omega \end{cases} \qquad \text{where } x \in \{0,1,2,\ldots,N-1\}$$

for some unknown ω. Our task is to find ω, the "marked string." Classically, we'll need to query the box more than $N/2$ times to find ω with success probability above 1/2. This is a black-box version of an NP-hard problem, where there is a unique witness accepted by a circuit, but the problem has no structure, so there is no better option than exhaustive searching.

Now we ask, can exhaustive search be done faster on a quantum computer? The answer is yes, using "Grover's algorithm." With quantum queries, we can find the marked string using $O(\sqrt{N})$ queries. Thus, we can solve NP-hard problems by exhaustive search in time $O(\sqrt{N}\,\mathrm{polylog}\, N)$

We say that Grover's algorithm achieves a "quadratic speedup" relative to exhaustive search on a classical computer. Though the speedup is only quadratic rather than exponential, Grover's algorithm is interesting because of its broad applicability. And it is rather remarkable: in effect we can interrogate N potential witnesses by asking $O(\sqrt{N})$ questions.

In the quantum setting, the black box applies the unitary

$$U_\omega : |x\rangle \otimes |y\rangle \longmapsto |x\rangle \otimes |y \oplus f_\omega(x)\rangle .$$

where $x \in \{0,1\}^n$ and $y \in \{0,1\}$. By the standard trick, U_ω becomes a "phase oracle":

$$U_\omega : |x\rangle \otimes |-\rangle \longmapsto (-1)^{f_\omega(x)} |x\rangle \otimes |-\rangle .$$

where $$(-1)^{f_\omega(x)} = \begin{cases} 1 & \omega \neq x \\ -1 & \omega = x \end{cases}$$

Ignoring the output register (which is unaffected by U_ω), we can express U_ω acting on input as

$$U_\omega = I - 2|\omega\rangle\langle\omega|$$

We can express a general n-qubit state $|\psi\rangle$ as

$$|\psi\rangle = a|\omega\rangle + b|\psi^\perp\rangle \xrightarrow[U_\omega]{} -a|\omega\rangle + |\psi^\perp\rangle$$

where $\langle\omega|\psi^\perp\rangle = 0$. That is, we resolve $|\psi\rangle$ into a component along $|\omega\rangle$ and a component in the hyperplane orthogonal to $|\omega\rangle$.

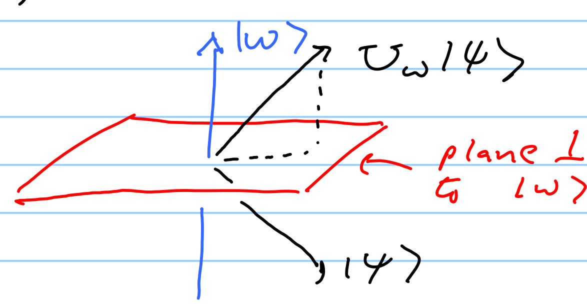

U_ω induces a reflection of the vector $|\psi\rangle$ about this hyperplane.

The first step in Grover's algorithm is to prepare the uniform superposition of all values of x:

$$|s\rangle = H^{\otimes n}|0\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle$$

with the marked string $|\omega\rangle$, this state $|s\rangle$ has overlap

$$\langle \omega | s\rangle = \frac{1}{\sqrt{N}}$$

The next step is to apply the "Grover iteration" many times in succession, where each iteration enhances the overlap of the quantum sup. with the marked state $|\omega\rangle$, while suppressing the amplitude for each $|x\rangle$ with $x \neq \omega$. This iteration is

$$U_{Grover} = U_s U_\omega$$

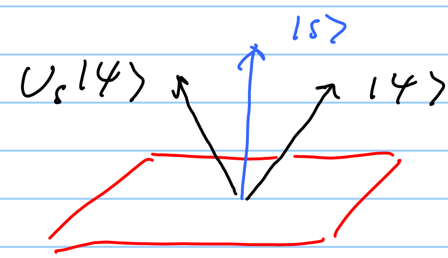

where U_ω is the query and

$$U_s = 2|s\rangle\langle s| - I,$$

which reflects a vector about the axis determined by $|s\rangle$. Note that U_s is easy to construct as a quantum circuit. It can be expressed as

$$U_s = H^{\otimes n}\left(2|0\rangle\langle 0| - I\right) H^{\otimes n}$$

since $H^{\otimes n}: |s\rangle \longmapsto |0\rangle$, where $H = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}$ is the single-qubit Hadamard Gate. Furthermore

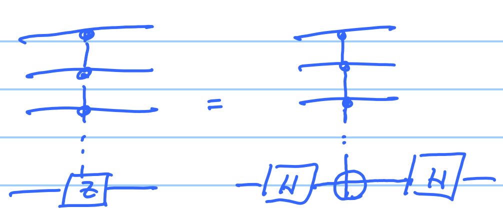

and $\Lambda^{n-1}(X)$ can be constructed from $O(n)$ Toffoli gates. Finally, we can conjugate by $X^{\otimes n}$ so phase is triggered by $|00\cdots0\rangle$ rather than $|11\cdots1\rangle$.

So U_s is realized by a circuit of size $O(\log N)$

What does U_{Grover} do? It preserves the plane spanned by $|s\rangle$ and $|w\rangle$, so we may confine our attention to that plane.

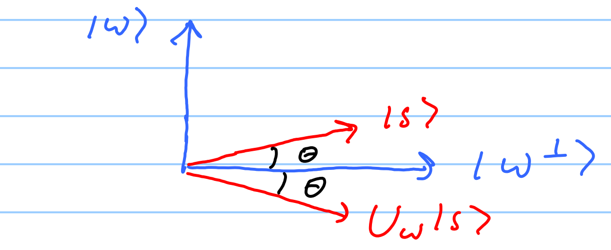

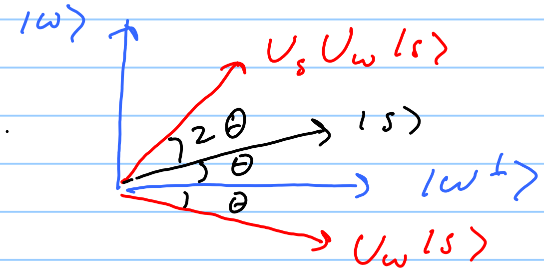

U_w reflects $|s\rangle$ about the axis $|w^\perp\rangle$ (the vector $\perp$ to $|w\rangle$ in the span of $|w\rangle$ and $|s\rangle$)

Then U_s reflects $U_w|s\rangle$ about the axis $|s\rangle$. The net effect of U_{Grover}, then, is to rotate CCW by 2θ, where θ is initial angle between $|s\rangle$ and $|w^\perp\rangle$

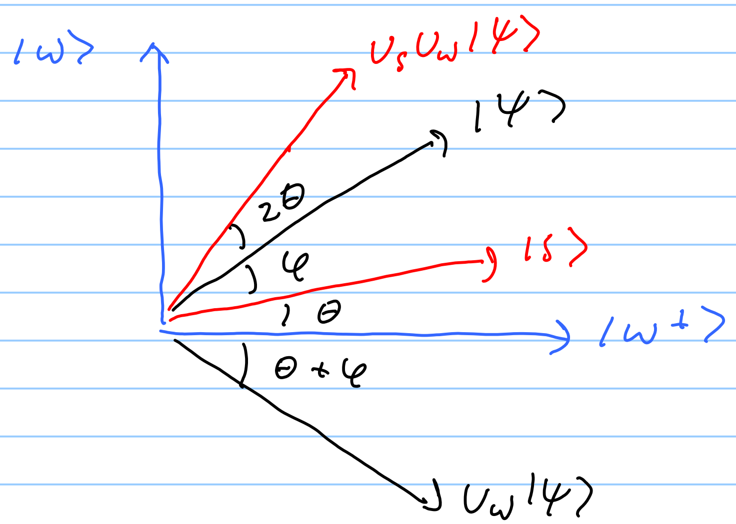

Each time we repeat the Grover iteration, the vector rotates further CCW by 2θ.

The initial angle θ between $|s\rangle$ and $|w^\perp\rangle$ is given by

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

For $N \gg 1$, then

$$\theta = \frac{1}{\sqrt{N}} + O\left(\frac{1}{N^{3/2}}\right).$$

If we repeat the Grover iteration T times, then the vector is rotated away from the $|w^\perp\rangle$ axis by $(2T+1)\theta$

we may choose T so that $(2T+1)\theta = \frac{\pi}{2} + \delta$ where $|\delta| \le \frac{\theta}{2} \simeq \frac{1}{2\sqrt{N}}$. Then if we measure in the computational basis, we find the outcome $|\omega\rangle$ with probability $\mathrm{Prob}(\omega) = \cos^2\delta \geq 1 - \delta^2 \geq 1 - \frac{1}{4N}$

Thus we find $|\omega\rangle$ with success probability close to 1 using $T \approx \frac{\pi}{4\theta} \approx \frac{\pi}{4}\sqrt{N}$ Grover iterations

and therefore also T quantum queries to the black box. This is Grover's quadratic speedup.

Suppose there are r marked states, where r is known. Classically, with each query the prob. of finding a solution (ω_i such that $f(\omega_i) = 1$) is r/N, so we need $O(N/r)$ queries to find a solution with constant success probability. Quantumly, the uniform superposition of the marked states

$$|\text{Marked}\rangle = \frac{1}{\sqrt{r}} \sum_{i=1}^{r} |\omega_i\rangle$$

has overlap with $\quad |s\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle$

$$\langle \text{Marked} | s \rangle = \sqrt{\frac{r}{N}} = \sin\theta$$

and the Grover iteration again rotates by 2θ in the plane spanned by $|s\rangle$ and $|\text{Marked}\rangle$ (because query reflects about the axis $\perp$ to $|\text{Marked}\rangle$).

As above, then, for $N/r \gg 1$, we achieve success prob $\mathrm{Prob} = 1 - O(r/N)$ in $T \approx \pi/4 \sqrt{N/r}$ queries

Again, the speedup is quadratic. The number of quantum queries needed to find a solution is

$$\#\text{ quantum queries} = O\left(\sqrt{\#\text{ classical queries}}\right).$$

What if r is not known a priori? As a function of the number of queries the success probability oscillates, where the period of the oscillation is

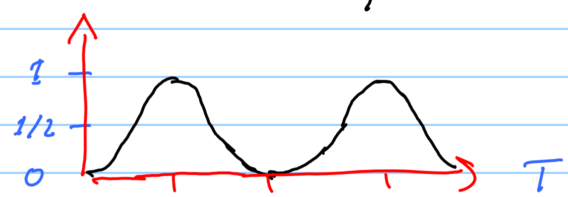

$$T \approx \frac{\pi}{2}\sqrt{\frac{N}{r}}$$

If we choose T uniformly at random in the interval $T \in \{0, 1, 2, \ldots, \approx \frac{\pi}{4}\sqrt{N}\}$, then if there is a solution $(r>0)$, then a soln. will be found with prob $\geq \frac{1}{2} - O(\frac{1}{N})$

If we repeat m times, we will find a soln with failure prob $\lesssim 2^{-m}$. Therefore, we can use Grover's algorithm to solve a decision problem in NP with high success probability, in time

$$\text{time} = O\left(\sqrt{N}\ \text{poly}\log N\right)$$

since we can compute the circuit that evaluates $f(x)$ in (classical or quantum) time $O(\text{poly}\log N)$.

Generalized Search

In some cases, the problem may have structure that can be exploited to search faster for a solution.

In that case, some strings are better candidates than others to be solutions, and so we ought to be able to search more efficiently by spending more time testing more likely solutions than less likely ones.

For the case of Grover's exhaustive search of a function without any apparent structure, we started the algorithm by preparing

$$|s\rangle = H^{\otimes n}|0\rangle = \frac{1}{\sqrt{N}}\sum_{x=0}^{N-1}|x\rangle$$

which has overlap $\sin\theta = 1/\sqrt{N}$ with the solution $|\omega\rangle$, and hence would yield the uniform distribution on x if we measured in the computational basis. For a function with structure, we may be able to construct an efficiently computable unitary U such that $U|0\rangle = \sin\theta\,|\omega\rangle + \cos\theta\,|\psi^{\perp}\rangle$

where $\sin^2\theta > 1/N$. In that case, we can conduct Grover's algorithm with $H^{\otimes n}$ replaced by U, and U_s replaced by

$$\tilde{U}_s = U\left(2|0\rangle\langle 0| - I\right)U^{\dagger}$$

Thus $\tilde{U}_s$ reflects in the axis $U|0\rangle$ rather than $|s\rangle$.

The analysis of the algorithm is the same as before, and in the case where there exists a unique solution, we can find it with high probability in $T = \frac{\pi}{4\theta} < \frac{\pi}{4}\sqrt{N}$ queries

Specifically, because of the structure of the function, we might be able to exclude all except $M < N$ inputs as potential solutions. Then classically we could find the solution in $O(M)$ queries, while quantumly only $O(\sqrt{M})$ queries suffice, if we can construct U such that

$$U|0\rangle = \frac{1}{\sqrt{M}} \sum_{i=1}^{M} |x_i\rangle,$$ the uniform sup. of the candidate solutions.

For example, suppose that classical search for a solution can be accelerated by a "classical heuristic" — that is, a function g that takes a randomly generated "seed" r in a set R to a trial solution:

$$g : r \longmapsto g(r) \quad \text{where } r \in R$$

The heuristic is useful if trial solutions generated by the heuristic are more likely to be accepted than trial solutions chosen uniformly at random

$$\left\langle \frac{\#\text{ of soln. in } g(R)}{|R|} \right\rangle > \left\langle \frac{\text{total \# soln.}}{N} \right\rangle$$

where the bracket $\langle\ \rangle$ indicates the expectation value evaluated for a probability distribution on black-box functions. Then the number of classical queries to find a soln, using the heuristic, with constant success probability is

$$T_{class} = O\left(\left\langle \frac{|R|}{\#\text{ soln in } g(R)} \right\rangle\right)$$

To exploit the heuristic in quantum searching, we apply Grover's algorithm to searching in the space of seeds instead of the full search space. The heuristic is realized as an efficiently computable unitary:

$$|r\rangle \otimes |0\rangle \longmapsto |r\rangle \otimes |g(r)\rangle$$

We can query the box with $|g(r)\rangle$ and then run the evaluation of g backwards to erase garbage:

$$|r\rangle \otimes |0\rangle \otimes |y\rangle \longmapsto |r\rangle \otimes |g(r)\rangle \otimes |y\rangle$$
$$\longmapsto |r\rangle \otimes |g(r)\rangle |y \oplus f(g(r))\rangle \longmapsto |r\rangle \otimes |0\rangle \otimes |y \oplus f(g(r))\rangle.$$

This composite oracle can be consulted to search R for a state marked by the function $f \circ g$ (i.e., for a state marked by f in $g(R)$, the range of g).

The number of quantum queries used is

$$T_{quantum} = O\left\langle \sqrt{\frac{|R|}{\#\text{ of soln in } g(R)}} \right\rangle$$

(for each black box there is a quadratic speed up)
Furthermore, the square root function is concave:

$$\langle \sqrt{F} \rangle = \sum_i p_a \sqrt{F_a} \le \sqrt{\sum_i p_a F_a} = \sqrt{\langle F \rangle}$$

where $\{p_a\}$ is a prob. distribution, and here F_a represents the classical query complexity for a black box function labeled by a

Thus $$T_{quantum} \le O\left(\sqrt{\left\langle \frac{|R|}{\#\text{ soln in } g(R)} \right\rangle}\right) = O\sqrt{T_{classical}}.$$

The speedup is quadratic, as for unstructured search.