Quantum algorithms (CO 781, Winter 2008)
Prof. Andrew Childs, University of Waterloo

LECTURE 15: Unstructured search and spatial search

Now we begin to discuss applications of quantum walks to search algorithms. We start with the most basic of all search problems, the unstructured search problem (that Grover's algorithm solves optimally). We discuss both discrete- and continuous-time quantum walk algorithms for this problem, and we prove that the running times of these algorithms are optimal. Finally, we discuss quantum walk algorithms for the search problem under locality constraints.

Unstructured search In the unstructured search problem, we are given a black box function $f : S \to \{0, 1\}$, where S is a finite set of size $|S| = N$. The inputs $x \in M$, where $M := \{x \in S : f(x) = 1\}$, are called *marked items*. In the decision version of the problem, our goal is to determine whether M is empty or not. We might also want to find a marked item when one exists.

It is quite easy to see that even the decision problem requires $\Omega(N)$ classical queries, and that N queries suffice, so the classical query complexity of unstructured search is $\Theta(N)$.

You should already be familiar with Grover's algorithm, which solves this problem using $O(\sqrt{n})$ quantum queries. Grover's algorithm works by starting from the state $|S\rangle := \sum_{x\in S} |x\rangle/\sqrt{N}$ and alternately applying the reflection about the set of marked items, $\sum_{x\in M} 2|x\rangle\langle x| - 1$, and the reflection about the state $|S\rangle$, $2|S\rangle\langle S| - 1$. The former can be implemented with two quantum queries to f, and the latter requires no queries to implement. It is straightforward to show that there is some $t = O(\sqrt{N/|M|})$ for which t steps of this procedure give a state with constant overlap on $|M\rangle$ (assuming M is non-empty), so that a measurement will reveal a marked item with constant probability.

Discrete-time quantum walk algorithm Consider the discrete-time random walk on the complete graph represented by the stochastic matrix

$$P = \frac{1}{N-1}\begin{pmatrix} 0 & 1 & 1 & \cdots & 1 \\ 1 & 0 & 1 & \cdots & 1 \\ 1 & 1 & 0 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 1 \\ 1 & 1 & \cdots & 1 & 0 \end{pmatrix} \tag{1}$$

$$= \frac{N}{N-1}|S\rangle\langle S| - \frac{1}{N-1}I. \tag{2}$$

It has eigenvalues 1 (which is non-degenerate) and $-1/(N-1)$ (with degeneracy $N-1$). Since the graph is highly connected, its spectral gap is very large: we have $\delta = 1 - \frac{1}{N-1} = \frac{N}{N-1}$.

This random walk gives rise to a very simple classical algorithm for unstructured search. In this algorithm, we start from a uniformly random item and repeatedly choose a new item uniformly at random from the other $N-1$ possibilities, stopping when we reach a marked item. The fraction of marked items is $\epsilon = |M|/N$, so the hitting time of this walk is

$$O\left(\frac{1}{\delta\epsilon}\right) = \frac{(N-1)N}{N|M|} = O(N/|M|) \tag{3}$$

(this is only an upper bound on the hitting time, but in this case we know it is optimal). Of course, if we have no a priori lower bound on $|M|$ in the event that M is non-empty, the best we can say is that $\epsilon \geq 1/N$, giving a running time $O(N)$.

The corresponding quantum walk search algorithm has a hitting time of

$$O\left(\frac{1}{\sqrt{\delta\epsilon}}\right) = O(\sqrt{N/|M|}), \tag{4}$$

corresponding to the running time of Grover's algorithm. To see that this actually gives an algorithm using $O(\sqrt{N/|M|})$ queries, we need to see that a step of the quantum walk can be performed using only $O(1)$ quantum queries. In the case where the first item is marked, the modified classical walk matrix is

$$P' = \frac{1}{N-1}\begin{pmatrix} N-1 & 1 & 1 & \cdots & 1 \\ 0 & 0 & 1 & \cdots & 1 \\ 0 & 1 & 0 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 1 \\ 0 & 1 & \cdots & 1 & 0 \end{pmatrix}, \tag{5}$$

so that the vectors $|\psi_j\rangle$ are $|\psi_1\rangle = |1,1\rangle$ and $|\psi_j\rangle = |j, S \setminus \{j\}\rangle = \sqrt{\frac{N}{N-1}}|j,S\rangle - \frac{1}{\sqrt{N-1}}|j,j\rangle$ for $j = 2, \ldots, N$. With a general marked set M, the projector onto the span of these states is

$$\Pi = \sum_{j \in M} |j,j\rangle\langle j,j| + \sum_{j \notin M} |j, S \setminus \{j\}\rangle\langle j, S \setminus \{j\}|, \tag{6}$$

so the operator $2\Pi - 1$ acts as Grover diffusion over the neighbors when when the vertex is unmarked, and as a phase flip when the vertex is marked. (Note that since we start from the state $|\psi\rangle = \sum_{j \notin M} |\psi_j\rangle$, we stay in the subspace of states $\operatorname{span}\{|j,k\rangle : (j,k) \in E\}$, and in particular have zero support on any state $|j,j\rangle$ for $j \in V$, so $2\Pi - 1$ acts as -1 when the first register holds a marked vertex.) Each such step can be implemented using two queries of the black box, one to compute whether we are at a marked vertex and one to uncompute it; and the subsequent swap operation requires no queries. Thus the query complexity is indeed $O(\sqrt{N/|M|})$.

This algorithm is not exactly the same as Grover's; for example, it works in the Hilbert space $\mathbb{C}^N \otimes \mathbb{C}^N$ instead of $\mathbb{C}^N$. Nevertheless, it is clearly closely related. In particular, notice that in Grover's algorithm, the unitary operation $2|S\rangle\langle S| - 1$ can be viewed as a kind of discrete-time quantum walk on the complete graph, where in this particular case no coin is necessary to define the walk.

The algorithm we have described so far only solves the decision version of unstructured search. To find marked item, we could use bisection, but this would introduce an overhead of $O(\log N)$. Alternatively, we could show that the final state of the quantum walk actually encodes a marked item when one exists. We will now show that this is in fact the case for the continuous-time version of this algorithm; the analysis of the discrete-time case is left as an exercise.

Continuous-time quantum walk algorithm Now let us fully analyze the behavior of a corresponding continuous-time quantum walk algorithm for unstructured search, assuming for simplicity that there is a unique marked item m, and our goal is to find it. Clearly this is sufficient to solve the decision problem with the promise that there are either 0 or 1 marked items, which is essentially the hardest case.

This algorithm is defined by a Hamiltonian given by the adjacency matrix of complete graph plus a marking term, namely

$$H' = \left(|S\rangle\langle S| - \frac{1}{N}I\right) + |m\rangle\langle m|. \tag{7}$$

Since the term proportional to I just generates a global phase, we can drop it to give

$$H = |S\rangle\langle S| + |m\rangle\langle m|. \tag{8}$$

Suppose we start from the state $|S\rangle$ (the only sensible starting state given the symmetry of the problem) and choose and evolution time so that we have a substantial probability of observing m if we measure in the vertex basis.

The calculation of the walk dynamics is particularly straightforward since the walk is confined to the two-dimensional subspace $\operatorname{span}\{|m\rangle, |S\rangle\}$. Let us write H in an orthonormal basis composed of $|m\rangle$ and the orthogonal state

$$|m^\perp\rangle = \frac{|S\rangle - \alpha|m\rangle}{\sqrt{1-\alpha^2}} \tag{9}$$

where $\alpha := \langle S|m\rangle = \frac{1}{\sqrt{N}}$. Then

$$|S\rangle = \alpha|m\rangle + \sqrt{1-\alpha^2}|m^\perp\rangle. \tag{10}$$

Thus in the basis $\{|m\rangle, |m^\perp\rangle\}$, we have

$$\begin{aligned} H &= \begin{pmatrix} \alpha^2 & \alpha\sqrt{1-\alpha^2} \\ \alpha\sqrt{1-\alpha^2} & 1-\alpha^2 \end{pmatrix} + \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} && (11) \\ &= I + \alpha(\sqrt{1-\alpha^2}\sigma_x + \alpha\sigma_z). && (12) \end{aligned}$$

Finally, we can calculate the evolution as

$$\begin{aligned} |\psi(t)\rangle &:= e^{-iHt}|S\rangle && (13) \\ &= \cos\alpha t|S\rangle - i\sin\alpha t(\sqrt{1-\alpha^2}\sigma_x + \alpha\sigma_z)|S\rangle && (14) \\ &= \cos\alpha t|S\rangle - i\sin\alpha t(\sqrt{1-\alpha^2}\sigma_x + \alpha\sigma_z)(\alpha|m\rangle + \sqrt{1-\alpha^2}|m^\perp\rangle) && (15) \\ &= \cos\alpha t|S\rangle - i\sin\alpha t(\alpha\sqrt{1-\alpha^2}|m^\perp\rangle + (1-\alpha^2)|m\rangle + \alpha^2|m\rangle - \alpha\sqrt{1-\alpha^2}|m^\perp\rangle) && (16) \\ &= \cos\alpha t|S\rangle - i\sin\alpha t|m\rangle. && (17) \end{aligned}$$

The probability of observing m if we stop the walk and measure in the vertex basis after time t is

$$|\langle m|\psi(t)\rangle|^2 = \alpha^2\cos^2\alpha t + \sin^2\alpha t. \tag{18}$$

In particular, when $t = \frac{\pi}{2\alpha} = \frac{\pi\sqrt{N}}{2}$, we observe the marked item m with probability 1.

Lower bound In fact, Grover's algorithm is optimal (up to a constant factor): any algorithm for unstructured search, even with the promise that there are either 0 or 1 marked items, requires $\Omega(\sqrt{N})$ queries. This is one of the most important results in quantum computing, so let's prove it.

Suppose we can employ an arbitrary time-dependent Hamiltonian $K(t)$ in the case where there is no marked item, and

$$H_m(t) = K(t) + |m\rangle\langle m| \tag{19}$$

in the case where there is a unique marked item m. (If the algorithm uses additional workspace, then the marking term simply acts as the identity on that part of the space.) Furthermore, we let the initial state $|\psi(0)\rangle$ be any fixed, m-independent state. We will show that the time to decide whether some vertex is marked is $\Omega(\sqrt{N})$.

Note that this will imply a lower bound of $\Omega(\sqrt{N})$ queries in the usual quantum query model, since by letting $K(t) = 0$ and evolving for time π, we can perform a phase flip; and by letting $K(t)$ be arbitrarily large for a short time period, we can perform a unitary gate without the oracle acting.

Now let us compare the evolution under $H_m(t)$, giving a state $|\psi_m(t)\rangle$, to the evolution under $K(t)$ alone, giving a state $|\phi(t)\rangle$. Define

$$\begin{aligned} d_m &:= \| |\psi_m(t)\rangle - |\phi(t)\rangle \|^2 && (20) \\ &= 2(1 - \mathrm{Re}\langle\psi_m(t)|\phi(t)\rangle). && (21) \end{aligned}$$

To be able to distinguish the two possibilities, we need $d_m \geq \epsilon$ for some $\epsilon > 0$. Since this must hold for each possible marked item m, we have

$$\sum_m d_m \geq N\epsilon. \tag{22}$$

Now

$$\begin{aligned} \left|\frac{\mathrm{d}d_m}{\mathrm{d}t}\right| &= 2\left|\frac{\mathrm{d}}{\mathrm{d}t}\mathrm{Re}\langle\psi_m(t)|\phi(t)\rangle\right| && (23) \\ &= 2|\,\mathrm{Re}(\langle\psi_m(t)|iK(t)|\phi(t)\rangle - \langle\psi_m(t)|iH_m(t)|\phi(t)\rangle)| && (24) \\ &= 2|\,\mathrm{Im}(\langle\psi_m(t)|m\rangle\langle m|\phi(t)\rangle)| && (25) \\ &\leq 2|\langle\psi_m(t)|m\rangle\langle m|\phi(t)\rangle| && (26) \\ &\leq 2|\langle m|\phi(t)\rangle|. && (27) \end{aligned}$$

(Here our notation assumes that there is no extra workspace; the reader is invited to check that the conclusions are not affected by this assumption.) Summing on m, we have

$$\begin{aligned} \left|\frac{\mathrm{d}}{\mathrm{d}t}\sum_m d_m\right| &\leq \sum_m \left|\frac{\mathrm{d}}{\mathrm{d}t}d_m\right| && (28) \\ &\leq 2\sum_m |\langle m|\phi(t)\rangle| && (29) \\ &\leq 2\sqrt{N} && (30) \end{aligned}$$

where in the last step we have used Cauchy-Schwarz. Finally, integrating $\sum_m d_m$ with the initial condition $d_m = 0$ at $t = 0$ (and using the triangle inequality), we find

$$\sum_m d_m \leq 2\sqrt{N}t. \tag{31}$$

Comparing to (22), we have $t \geq \frac{\epsilon}{2}\sqrt{N}$, which shows that $t = \Omega(\sqrt{N})$, as claimed.

Search on graphs Now let's consider a variant of unstructured search with additional locality constraints. We will view the items in S as the vertices of a graph $G = (S, E)$, and we require the algorithm to be local with respect to the graph. More concretely, the algorithms alternates between queries and unitary operations U constrained to satisfy $U|j, \psi\rangle = \sum_{k \in j \cup \partial(j)} \alpha_k |k, \phi_k\rangle$ for any $j \in S$ (where the second register represents possible ancillary space, and recall that $\partial(j)$ denotes the set of neighbors of j in G).

Since we have only added new restrictions that an algorithm must obey, the $\Omega(\sqrt{N})$ lower bound from the non-local version of the problem still applies. However, it is immediately clear that this bound cannot always be achieved. For example, if the graph is a cycle of N vertices, then simply propagating from one vertex of the cycle to an opposing vertex takes time $\Omega(N)$. So we would like to know, for example, how far from complete the graph can be such that we can still perform the search in $O(\sqrt{N})$ steps.

First, note that any expander graph (a graph with degree upper bounded by a constant and second largest eigenvalue bounded away from 1 by a constant) can be searched in time $O(\sqrt{N})$. Such graphs have $\delta = \Omega(1)$, and since $\epsilon \geq 1/N$ when there are marked items, the quantum hitting time is $O(1/\sqrt{\delta\epsilon}) = O(\sqrt{N})$ (whereas the classical hitting time is $O(1/\delta\epsilon) = O(N)$). Indeed, a randomly chosen d-regular graph for constant d is such an expander with high probability.

There are also many cases in which a quantum search can be performed in time $O(\sqrt{N})$ even though the eigenvalue gap of P is non-constant. For example, consider the n-dimensional hypercube (with $N = 2^n$ vertices). Recall that since the adjacency matrix acts independently as σ_x on each coordinate, the eigenvalues are equally spaced, and the gap of P is $2/n$. Thus the general bound in terms of the eigenvalues of P shows that the classical hitting time is $O(nN) = O(N \log N)$. In fact, this bound is loose; the hitting time is actually $O(N)$, which can be seen by directly computing $\|P_M\|$ with one marked vertex. So there is a local quantum algorithm that runs in the square root of this time, namely $O(\sqrt{N})$.

Perhaps the most interesting example is the d-dimensional square lattice with N sites (i.e., with linear size $N^{1/d}$). This case can be viewed as having N items distributed on a grid in d-dimensional space. For simplicity, suppose we have periodic boundary conditions; then the eigenstates of the adjacency matrix are given by

$$|\tilde{k}\rangle := \frac{1}{\sqrt{N}} \sum_x e^{2\pi i k \cdot x / N^{1/d}} |x\rangle \tag{32}$$

where k is a d-component vector of integers from 0 to $N^{1/d} - 1$. The corresponding eigenvalues are

$$2 \sum_{j=1}^{d} \cos \frac{2\pi k_j}{N^{1/d}}. \tag{33}$$

Normalizing to obtain a stochastic matrix, we simply divide these eigenvalues by $2d$. The 1 eigenvector has $k = (0, 0, \ldots, 0)$, and the second largest eigenvalue comes from (e.g.) $k = (1, 0, \ldots, 0)$, with an eigenvalue

$$\frac{1}{d}\left(d - 1 + \cos \frac{2\pi}{N^{1/d}}\right) \approx 1 - \frac{1}{2d}\left(\frac{2\pi}{N^{1/d}}\right)^2. \tag{34}$$

Thus the gap of the walk matrix P is about $\frac{2\pi^2}{2dN^{2/d}} = O(N^{-2/d})$. This is another case in which the bound on the classical hitting time in terms of eigenvalues of P is too loose (it gives only $O(N^{1+2/d})$), and instead we must directly estimate the gap of P_M. One can show that the classical hitting time is $O(N^2)$ in $d = 1$, $O(N \log N)$ in $d = 2$, and $O(N)$ for any $d \geq 3$. Thus there is a local quantum walk search algorithm that saturates the lower bound for any $d \geq 3$, and one that runs in time time $O(\sqrt{N \log N})$ for $d = 2$. We already argued that there could be no speedup for $d = 1$, and indeed we see that the quantum hitting time in this case is $O(N)$.

Note that similar results for spatial search can be obtained in the framework of continuous-time quantum walk.