Quantum algorithms (CO 781/CS 867/QIC 823, Winter 2013)
Andrew Childs, University of Waterloo

LECTURE 10: Continuous-time quantum walk

We now turn to our second major topic in quantum algorithms, the concept of quantum walk. In this lecture we will introduce continuous-time quantum walk as a natural analog of continuous-time classical random walk, and we'll see some examples of how the two kinds of processes differ.

Continuous-time quantum walk

Random walks come in two flavors: discrete- and continuous-time. It is easiest to define a quantum analog of a continuous-time random walk, so we consider this case first. Given a graph $G = (V, E)$, we define the continuous-time random walk on G as follows. Let A be the *adjacency matrix* of G, the $|V| \times |V|$ matrix with

$$A_{j,k} = \begin{cases} 1 & (j,k) \in E \\ 0 & (j,k) \notin E \end{cases} \tag{1}$$

for every pair $j, k \in V$. In particular, if we disallow self loops, then the diagonal of A is zero. There is another matrix associated with G that is nearly as important: the *Laplacian* of G, which has

$$L_{j,k} = \begin{cases} -\deg(j) & j = k \\ 1 & (j,k) \in E \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

where $\deg(j)$ denotes the degree of vertex j. (The Laplacian is sometimes defined differently than this—e.g., sometimes with the opposite sign. We use this definition because it makes L a discrete approximation of the Laplacian operator ∇^2 in the continuum.)

The continuous-time random walk on G is defined as the solution of the differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t} p_j(t) = \sum_{k \in V} L_{jk}\, p_k(t). \tag{3}$$

Here $p_j(t)$ denotes the probability associated with vertex j at time t. This can be viewed as a discrete analog of the diffusion equation. Note that

$$\frac{\mathrm{d}}{\mathrm{d}t} \sum_{j \in V} p_j(t) = \sum_{j,k \in V} L_{jk}\, p_k(t) = 0 \tag{4}$$

(since the columns of L sum to 0), which shows that an initially normalized distribution remains normalized: the evolution of the continuous-time random walk for any time t is a *stochastic process*. The solution of the differential equation can be given in closed form as

$$p(t) = e^{Lt} p(0). \tag{5}$$

Now notice that the equation (3) is very similar to the Schrödinger equation

$$\mathrm{i} \frac{\mathrm{d}}{\mathrm{d}t} |\psi\rangle = H |\psi\rangle \tag{6}$$

except that it lacks the factor of i. If we simply insert this factor, and rename the probabilities $p_j(t)$ as quantum amplitudes $q_j(t) = \langle j|\psi(t)\rangle$ (where $\{|j\rangle : j \in V\}$ is an orthonormal basis for the Hilbert space), then we obtain the equation

$$\mathrm{i}\frac{\mathrm{d}}{\mathrm{d}t} q_j(t) = \sum_{k \in V} L_{jk}\, q_k(t), \tag{7}$$

which is simply the Schrödinger equation with the Hamiltonian given by the Laplacian of the graph. Since the Laplacian is a Hermitian operator, these dynamics preserve normalization in the sense that $\frac{\mathrm{d}}{\mathrm{d}t}\sum_{j\in V} |q_j(t)|^2 = 0$. Again the solution of the differential equation can be given in closed form, but here it is $|\psi(t)\rangle = e^{-iLt}|\psi(0)\rangle$.

We could also define a continuous-time quantum walk using any Hermitian Hamiltonian that respects the structure of G. For example, we could use the adjacency matrix A of G, even though this matrix cannot be used as the generator of a continuous-time classical random walk.

Random and quantum walks on the hypercube

Let's begin by investigating a simple, dramatic example of a difference between the behavior of random and quantum walks. Consider the Boolean hypercube, the graph with vertex set $V = \{0,1\}^n$ and edge set $E = \{(x,y) \in V^2 : \Delta(x,y) = 1\}$, where $\Delta(x,y)$ denotes the Hamming distance between the strings x and y. When $n = 1$, the hypercube is simply an edge, with adjacency matrix

$$\sigma_x := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \tag{8}$$

For general n, the graph is the direct product of this graph with itself n times, and the adjacency matrix is

$$A = \sum_{j=1}^{n} \sigma_x^{(j)} \tag{9}$$

where $\sigma_x^{(j)}$ denotes the operator acting as σ_x on the jth bit, and as the identity on every other bit.

For simplicity, let's consider the quantum walk with the Hamiltonian given by the adjacency matrix. (In fact, since the graph is regular, the walk generated by the Laplacian would only differ by an overall phase.) Since the terms in the above expression for the adjacency matrix commute, the unitary operator describing the evolution of this walk is simply

$$\begin{aligned} e^{-\mathrm{i}At} &= \prod_{j=1}^{n} e^{-\mathrm{i}\sigma_x^{(j)} t} && (10) \\ &= \bigotimes_{j=1}^{n} \begin{pmatrix} \cos t & -\mathrm{i}\sin t \\ -\mathrm{i}\sin t & \cos t \end{pmatrix}. && (11) \end{aligned}$$

After time $t = \pi/2$, this operator flips every bit of the state (up to an overall phase), mapping any input state $|x\rangle$ to the state $|\bar{x}\rangle$ corresponding to the opposite vertex of the hypercube.

In contrast, consider the continuous- or discrete-time random walk starting from the vertex x. It is not hard to show that the probability of reaching the opposite vertex $\bar{x}$ is exponentially small at any time, since the walk rapidly reaches the uniform distribution over all 2^n vertices of the hypercube. So this simple example shows that random and quantum walks can exhibit radically different behavior.

Random and quantum walks in one dimension

Perhaps the best-known example of a random walk is the case of an infinite path, with $V = \mathbb{Z}$ and $(j,k) \in E$ iff $|j-k| = 1$. It is well known that the random walk on this graph starting from the origin (in either continuous or discrete time) typically moves a distance proportional to $\sqrt{t}$ in time t. Now let's consider the corresponding quantum walk.

To calculate the behavior of the walk, it is helpful to diagonalize the Hamiltonian. The eigenstates of the Laplacian of the graph are the momentum states $|\hat{p}\rangle$ with components

$$\langle j|\hat{p}\rangle = e^{\mathrm{i}pj} \tag{12}$$

where $-\pi \le p \le \pi$. We have

$$\begin{aligned}
\langle j|L|\hat{p}\rangle &= \langle j+1|\hat{p}\rangle + \langle j-1|\hat{p}\rangle - 2\langle j|\hat{p}\rangle &(13)\\
&= (e^{\mathrm{i}p(j+1)} + e^{\mathrm{i}p(j-1)} - 2e^{\mathrm{i}pj}) &(14)\\
&= e^{\mathrm{i}pj}(e^{\mathrm{i}p} + e^{-\mathrm{i}p} - 2) &(15)\\
&= 2(\cos p - 1)\langle j|\hat{p}\rangle, &(16)
\end{aligned}$$

so the corresponding eigenvalue is $2(\cos p - 1)$. Thus the amplitude for the walk to move from j to k in time t is

$$\begin{aligned}
\langle k|e^{-\mathrm{i}Lt}|j\rangle &= \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{-2\mathrm{i}t(\cos p - 1)}\langle k|\hat{p}\rangle\langle\hat{p}|j\rangle \,\mathrm{d}p &(17)\\
&= \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{ip(k-j)-2it(\cos p-1)} \,\mathrm{d}p &(18)\\
&= e^{2\mathrm{i}t}(-\mathrm{i})^{k-j} J_{k-j}(2t) &(19)
\end{aligned}$$

where J_ν is the Bessel function of order ν. This expression can be understood using basic asymptotic properties of the Bessel function. For large values of ν, the function $J_\nu(t)$ is exponentially small in ν for $\nu \gg t$, of order $\nu^{-1/3}$ for $\nu \approx t$, and of order $\nu^{-1/2}$ for $\nu \ll t$. Thus (19) describes a wave propagating with speed 2.

We can use a similar calculation to exactly describe the corresponding continuous-time classical random walk, which is simply the analytic continuation of the quantum case with $t \to \mathrm{i}t$. Here the probability of moving from j to k in time t is

$$[e^{Lt}]_{kj} = e^{-2t} I_{k-j}(2t), \tag{20}$$

where I_ν is the modified Bessel function of order ν. For large t, this expression is approximately $\frac{1}{\sqrt{4\pi t}}\exp(-(k-j)^2/4t)$, a Gaussian of width $\sqrt{2t}$, in agreement with our expectations for a classical random walk in one dimension.

Black-box traversal of the glued trees graph

We have seen that the behavior of a quantum walk can be dramatically different from that of its classical counterpart. Next we will see an even stronger example of the power of quantum walk: a black-box problem that can be solved exponentially faster by a quantum walk than by *any* classical algorithm.

Consider a graph obtained by starting from two balanced binary trees of height n, and joining them by a random cycle of length $2 \cdot 2^n$ that alternates between the leaves of the two trees. For example, such a graph for $n = 4$ could look like the following:

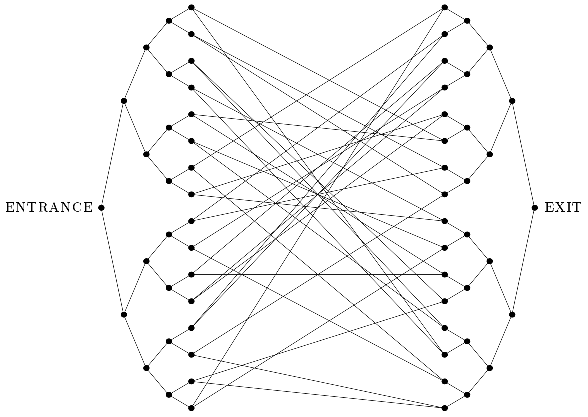

Suppose we take a random walk on the graph starting from the root of the left tree. It is not hard to see that such a walk rapidly gets lost in the middle of the graph and never has a substantial probability of reaching the opposite root. In fact, by specifying the graph in such a way that it can only be explored locally, we can ensure that no classical procedure starting from the left root can efficiently reach the right root. However, a quantum walk starting from the left root produces a state with a large (lower bounded by $1/\operatorname{poly}(n)$) overlap on the right root in a short (upper bounded by $\operatorname{poly}(n)$) amount of time.

To establish a provable separation between classical and quantum strategies, we will formulate the graph traversal problem in terms of query complexity.

Let $G = (V, E)$ be a graph with N vertices. To represent G by a black box, let m be such that $2^m \geq N$, and let k be at least as large as the maximum degree of G. For each vertex $a \in V$, assign a distinct m-bit string (called the *name* of a), not assigning $11\ldots1$ as the name of any vertex. For each $b \in V$ with $(a, b) \in E$, assign a unique label from $\{1, 2, \ldots, k\}$ to the ordered pair (a, b). For $a \in \{0, 1\}^m$ (identifying the vertex with its name) and $c \in \{1, 2, \ldots, k\}$, define $v_c(a)$ as the name of the vertex reached by following the outgoing edge of a labeled by c, if such an edge exists. If there is no vertex of G named a or no outgoing edge from a labeled c, then let $v_c(a) = 11\ldots1$. The black box for G takes $a \in \{0, 1\}^m$ and $c \in \{1, 2, \ldots, k\}$ as input and returns $v_c(a)$.

The black box graph traversal problem is as follows. Let G be a graph and let ENTRANCE and EXIT be two vertices of G. Given a black box for G as described above, with the additional promise that the name of the ENTRANCE is $00\ldots0$, the goal is to output the name of the EXIT. We say an algorithm for this problem is efficient if its running time is polynomial in m.

Of course, a random walk is not necessarily the best classical strategy for this problem. For example, there is an efficient classical algorithm for traversing the n-dimensional hypercube (exercise: what is it?) even though a random walk does not work. However, no classical algorithm can efficiently traverse the glued trees, whereas a quantum walk can.

Quantum walk algorithm to traverse the glued trees graph

Given a black box for a graph G as specified above, we can efficiently compute a list of neighbors of any desired vertex, provided $k = \operatorname{poly}(m)$ (i.e., provided the maximum degree of the graph is not too large). Thus it is straightforward to simulate the dynamics of the continuous-time quantum walk on any such G, and in particular, on the glued trees graph (which has maximum degree 3). Our strategy for solving the traversal problem is simply to run the quantum walk and show that the resulting state has a substantial overlap on the EXIT for some $t = \operatorname{poly}(n)$.

Let G be the glued trees graph. The dynamics of the quantum walk on this graph are dramatically simplified because of symmetry. Consider the basis of states $|\mathrm{col}\, j\rangle$ that are uniform superpositions over the vertices at distance j from the ENTRANCE, i.e.,

$$|\mathrm{col}\, j\rangle := \frac{1}{\sqrt{N_j}} \sum_{\delta(a,\textsc{entrance})=j} |a\rangle \tag{21}$$

where

$$N_j := \begin{cases} 2^j & 0 \le j \le n \\ 2^{2n+1-j} & n+1 \le j \le 2n+1 \end{cases} \tag{22}$$

is the number of vertices at distance j from the ENTRANCE, and where $\delta(a,b)$ denotes the length of the shortest path in G from a to b. It is straightforward to see that the subspace $\mathrm{span}\{|\mathrm{col}\, j\rangle : 0 \le j \le 2n+1\}$ is invariant under the action of the adjacency matrix A of G. At the ENTRANCE and EXIT, we have

$$A|\mathrm{col}\, 0\rangle = \sqrt{2}|\mathrm{col}\, 1\rangle \tag{23}$$

$$A|\mathrm{col}\, 2n+1\rangle = \sqrt{2}|\mathrm{col}\, 2n\rangle. \tag{24}$$

For any $0 < j < n$, we have

$$\begin{aligned} A|\mathrm{col}\, j\rangle &= \frac{1}{\sqrt{N_j}} \sum_{\delta(a,\textsc{entrance})=j} A|a\rangle && (25)\\ &= \frac{1}{\sqrt{N_j}} \left(2 \sum_{\delta(a,\textsc{entrance})=j-1} |a\rangle + \sum_{\delta(a,\textsc{entrance})=j+1} |a\rangle \right) && (26)\\ &= \frac{1}{\sqrt{N_j}} (2\sqrt{N_{j-1}}|\mathrm{col}\, j-1\rangle + \sqrt{N_{j+1}}|\mathrm{col}\, j+1\rangle) && (27)\\ &= \sqrt{2}(|\mathrm{col}\, j-1\rangle + |\mathrm{col}\, j+1\rangle). && (28) \end{aligned}$$

Similarly, for any $n+1 < j < 2n+1$, we have

$$\begin{aligned} A|\mathrm{col}\, j\rangle &= \frac{1}{\sqrt{N_j}} (\sqrt{N_{j-1}}|\mathrm{col}\, j-1\rangle + 2\sqrt{N_{j+1}}|\mathrm{col}\, j+1\rangle) && (29)\\ &= \sqrt{2}(|\mathrm{col}\, j-1\rangle + |\mathrm{col}\, j+1\rangle). && (30) \end{aligned}$$

The only difference occurs at the middle of the graph, where we have

$$\begin{aligned} A|\mathrm{col}\, n\rangle &= \frac{1}{\sqrt{N_n}} (2\sqrt{N_{n-1}}|\mathrm{col}\, n-1\rangle + 2\sqrt{N_{n+1}}|\mathrm{col}\, n+1\rangle) && (31)\\ &= \sqrt{2}|\mathrm{col}\, n-1\rangle + 2|\mathrm{col}\, n+1\rangle && (32) \end{aligned}$$

and similarly

$$\begin{aligned} A|\mathrm{col}\, n+1\rangle &= \frac{1}{\sqrt{N_{n+1}}} (2\sqrt{N_n}|\mathrm{col}\, n\rangle + 2\sqrt{N_{n+2}}|\mathrm{col}\, n+2\rangle) && (33)\\ &= 2|\mathrm{col}\, n\rangle + \sqrt{2}|\mathrm{col}\, n+2\rangle. && (34) \end{aligned}$$

In summary, the matrix elements of A between basis states for this invariant subspace can be depicted as follows:

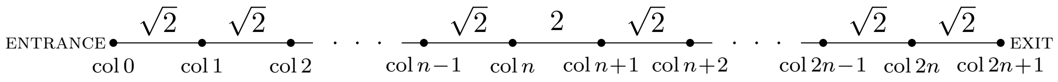

By identifying the subspace of states $|\mathrm{col}\, j\rangle$, we have found that the quantum walk on the glued trees graph starting from the ENTRANCE is effectively the same as a quantum walk on a weighted path of $2n+2$ vertices, with all edge weights the same except for the middle one. Given our example of the quantum walk on the infinite path, we can expect this walk to reach the EXIT with amplitude $1/\operatorname{poly}(n)$ in time linear in n. To prove that the walk indeed reaches the EXIT in polynomial time, we will use the notion of the *mixing time* of a quantum walk.

Classical and quantum mixing

Informally, the mixing time of a random walk is the amount of time it takes to come close to a stationary distribution. Recall that the continuous-time random walk on a graph $G = (V, E)$ with Laplacian L is defined as the solution of the differential equation $\frac{\mathrm{d}p(t)}{\mathrm{d}t} = Lp(t)$, where $p(t) \in \mathbb{R}^{|V|}$ denotes a vector of probabilities for the walk to be at each vertex at time t. The uniform distribution over the vertices, $u := (1, 1, \ldots, 1)/|V|$, is an eigenvector of L with eigenvalue 0. Indeed, if G is connected, then this is the unique eigenvector with this eigenvalue. Letting v_λ denote a normalized eigenvector of L with eigenvalue λ (so that $L = \sum_{\lambda\neq 0} \lambda v_\lambda v_\lambda^T$), we have

$$\begin{aligned}
p(t) &= e^{Lt} p(0) && (35)\\
&= \left(|V| u u^T + \sum_{\lambda \neq 0} e^{\lambda t} v_\lambda v_\lambda^T \right) p(0) && (36)\\
&= \langle |V| u, p(0)\rangle u + \sum_{\lambda\neq 0} e^{\lambda t} \langle v_\lambda, p(0)\rangle v_\lambda && (37)\\
&= u + \sum_{\lambda\neq 0} e^{\lambda t} \langle v_\lambda, p(0)\rangle v_\lambda && (38)
\end{aligned}$$

(where in exponentiating L we have used the fact that $\sqrt{|V|}u$ is a normalized eigenvector of L, so that $|V|uu^T$ is the projector onto the corresponding subspace). The Laplacian is a negative semidefinite operator, so the contributions $e^{\lambda t}$ for $\lambda \neq 0$ decrease exponentially in time; thus the walk asymptotically approaches the uniform distribution. The deviation from uniform is small when t is large compared to the inverse of the largest (i.e., least negative) nonzero eigenvalue of L.

Since a quantum walk is a unitary process, we should not expect it to approach a limiting quantum state, no matter how long we wait. Nevertheless, it is possible to define a notion of the limiting distribution of a quantum walk as follows. Suppose we pick a time t uniformly at random between 0 and T, run the quantum walk starting at $a \in V$ for a total time t, and then measure in the vertex basis. The resulting distribution is

$$\begin{aligned}
p_{a\to b}(T) &= \frac{1}{T}\int_0^T |\langle b|e^{-iHt}|a\rangle|^2 \mathrm{d}t && (39)\\
&= \sum_{\lambda,\lambda'} \langle b|\lambda\rangle\langle\lambda|a\rangle\langle a|\lambda'\rangle\langle\lambda'|v\rangle \frac{1}{T}\int_0^T e^{-i(\lambda-\lambda')t}\mathrm{d}t && (40)\\
&= \sum_{\lambda} |\langle a|\lambda\rangle\langle b|\lambda\rangle|^2 + \sum_{\lambda\neq\lambda'} \langle b|\lambda\rangle\langle\lambda|a\rangle\langle a|\lambda'\rangle\langle\lambda'|b\rangle \frac{1-e^{-i(\lambda-\lambda')T}}{i(\lambda-\lambda')T} && (41)
\end{aligned}$$

where we have considered a quantum walk generated by an unspecified Hamiltonian H (it could be the Laplacian or the adjacency matrix, or some other operator as desired), and where we have assumed for simplicity that the spectrum of $H = \sum_\lambda \lambda|\lambda\rangle\langle\lambda|$ is nondegenerate. We see that the distribution $p_{a\to b}(T)$ tends toward a limiting distribution

$$p_{a\to b}(\infty) := \sum_\lambda |\langle a|\lambda\rangle\langle b|\lambda\rangle|^2. \tag{42}$$

The timescale for approaching this distribution is again governed by the spectrum of H, but now we see that T must be large compared to the inverse of the smallest gap between any pair of distinct eigenvalues, not just the smallest gap between a particular pair of eigenvalues as in the classical case.

Let's apply this notion of quantum mixing to the quantum walk on the glued trees. It will be simplest to consider the walk generated by the adjacency matrix A. Since the subspace of states $|\mathrm{col}\, j\rangle$ has dimension only $2n+1$, it should not be surprising that the limiting probability of traversing from ENTRANCE to EXIT is bigger than $1/\mathrm{poly}(n)$. To see this, notice that A commutes with the reflection operator R defined as $R|\mathrm{col}\, j\rangle = |\mathrm{col}\, 2n+1-j\rangle$, so these two operators can be simultaneously diagonalized. Now $R^2 = 1$, so it has eigenvalues ± 1, which shows that we can choose the eigenstates $|\lambda\rangle$ of A to satisfy $\langle\text{ENTRANCE}|\lambda\rangle = \pm\langle\text{EXIT}|\lambda\rangle$. Therefore,

$$\begin{aligned} p_{\text{ENTRANCE}\to\text{EXIT}}(\infty) &= \sum_\lambda |\langle\text{ENTRANCE}|\lambda\rangle\langle\text{EXIT}|\lambda\rangle|^2 && (43)\\ &= \sum_\lambda |\langle\text{ENTRANCE}|\lambda\rangle|^4 && (44)\\ &\ge \frac{1}{2n+2}\left(\sum_\lambda |\langle\text{ENTRANCE}|\lambda\rangle|^2\right)^2 && (45)\\ &= \frac{1}{2n+2} && (46)\end{aligned}$$

where the lower bound follows by the Cauchy-Schwarz inequality. Thus it suffices to show that the mixing time of the quantum walk is $\mathrm{poly}(n)$.

To see how long we must wait before the probability of reaching the EXIT is close to its limiting value, we can calculate

$$\begin{aligned} &|p_{\text{ENTRANCE}\to\text{EXIT}}(\infty) - p_{\text{ENTRANCE}\to\text{EXIT}}(T)|\\ &\quad= \left|\sum_{\lambda\ne\lambda'} \langle\text{EXIT}|\lambda\rangle\langle\lambda|\text{ENTRANCE}\rangle\langle\text{ENTRANCE}|\lambda'\rangle\langle\lambda'|\text{EXIT}\rangle\frac{1-e^{-i(\lambda-\lambda')T}}{i(\lambda-\lambda')T}\right| && (47)\\ &\quad\le \frac{2}{\Delta T}\sum_{\lambda,\lambda'} |\langle\text{EXIT}|\lambda\rangle\langle\lambda|\text{ENTRANCE}\rangle\langle\text{ENTRANCE}|\lambda'\rangle\langle\lambda'|\text{EXIT}\rangle| && (48)\\ &\quad= \frac{2}{\Delta T}\sum_{\lambda,\lambda'} |\langle\text{ENTRANCE}|\lambda\rangle|^2|\langle\text{ENTRANCE}|\lambda'\rangle|^2 && (49)\\ &\quad= \frac{2}{\Delta T}, && (50)\end{aligned}$$

where Δ denotes the smallest gap between any pair of distinct eigenvalues of A. All that remains is to lower bound Δ.

To understand the spectrum of A, recall that an infinite path has eigenstates of the form e^{ipj}. For any value of p, the state $|\lambda\rangle$ with amplitudes $\langle \mathrm{col}\, j|\lambda\rangle = e^{ipj}$ satisfies $\langle \mathrm{col}\, j|A|\lambda\rangle = \lambda\langle \mathrm{col}\, j|\lambda\rangle$, where the eigenvalue is $\lambda = 2\sqrt{2}\cos p$, for all values of j except $0, n, n+1, 2n+1$. We can satisfy the eigenvalue condition for $j = 0, 2n+1$ by taking linear combinations of $e^{\pm ipj}$ that vanish for $j = -1$ and $j = 2n+2$, namely

$$\langle \mathrm{col}\, j|\lambda\rangle = \begin{cases} \sin(p(j+1)) & 0 \le j \le n \\ \pm\sin(p(2n+2-j)) & n+1 \le j \le 2n+1. \end{cases} \tag{51}$$

Finally, we can enforce the eigenvalue condition at $j = n$ (which automatically enforces it at $j = n+1$ by symmetry), which restricts the values of p to a finite set. We have

$$\sqrt{2}\sin(pn) \pm 2\sin(p(n+1)) = 2\sqrt{2}\cos(p)\sin(p(n+1)), \tag{52}$$

which can be simplified to

$$\frac{\sin(p(n+2))}{\sin(p(n+1))} = \pm\sqrt{2}. \tag{53}$$

The left hand side of this equation decreases monotonically, with poles at integer multiples of $\pi/(n+1)$. For example, with $n = 4$, we have the following:

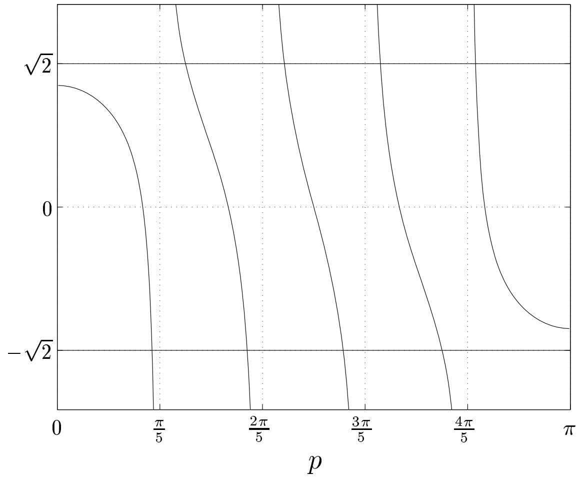

With a bit of analysis (see quant-ph/0209131 for details), one can show that the solutions of this equation give $2n$ values of p, each of which is separated from the integer multiples of $\pi/(n+1)$ by $\Omega(1/n^2)$. The spacings between the corresponding eigenvalues of A, $\lambda = 2\sqrt{2}\cos p$, are $\Omega(1/n^3)$. The remaining two eigenvalues of A can be obtained by considering solutions with p imaginary, and it is easy to show that they are separated from the rest of the spectrum by a constant amount. By taking (say) $T = 5n/\Delta = O(n^4)$, we can ensure that the probability to reach the EXIT is $\Omega(1/n)$. Thus there is an efficient quantum algorithm to traverse the glued trees graph.

Classical lower bound

It remains to show that this problem is difficult for a classical computer. A formal proof of this fact can be given using a sequence of reductions to problems that are essentially no easier than the original one, but that restrict the nature of the allowed algorithms. Here we will simply sketch the main ideas.

First, note that if we name the vertices at random using strings of about, say, $2\log|V|$ bits, then there are exponentially many more possible names than there are actual vertices. Since the probability that a randomly guessed name corresponds to a vertex of the graph is exponentially small, we can essentially restrict our attention to algorithms that query a connected set of vertices, starting from the ENTRANCE (the only vertex whose name is known initially).

Next, suppose we consider the algorithm to succeed not only if it reaches the EXIT, but also if it manages to find a cycle in the graph. This only makes it easier for the algorithm to succeed, but not significantly so, since it turns out to be hard even to find a cycle.

Now we can restrict our attention to the steps the algorithm takes before it finds a cycle. Notice that for such steps, the names supplied by the black box provide no information whatsoever about the structure of the graph: they could just as well be simulated by a sequence of random responses. Therefore, we can think of an algorithm as simply producing a rooted binary tree and embedding it into the glued trees graph at random. To show that the algorithm fails, it suffices to show that under such a random embedding, the probability of any rooted binary tree giving rise to a cycle or reaching the exit is small. By a fairly straightforward probabilistic argument, one can show that even for exponentially large trees (say, having at most $2^{n/6}$ vertiecs), the probability of the embedded tree giving rise to a cycle or reaching the exit is exponentially small. Thus any classical algorithm for solving the black box glued trees traversal problem must make exponentially many queries to succeed with more than exponentially small probability.