Quantum simulations of classical random walks and undirected graph connectivity

John Watrous*
Département IRO
Université de Montréal
Montréal (Québec), Canada
watrous@iro.umontreal.ca

October 2, 2006

Abstract

It is not currently known if quantum Turing machines can efficiently simulate probabilistic computations in the space-bounded case. In this paper we show that space-bounded quantum Turing machines can efficiently simulate a limited class of random processes: random walks on undirected graphs. By means of such simulations, it is demonstrated that the undirected graph connectivity problem for regular graphs can be solved by one-sided error quantum Turing machines that run in logspace and halt absolutely. It follows that symmetric logspace is contained in the quantum analogue of randomized logspace, i.e., $\mathrm{SL} \subseteq \mathrm{QR}_H\mathrm{L}$.

1 Introduction

This paper addresses the problem of space-efficient quantum simulations of probabilistic computations. Although quantum Turing machines are known to be at least as powerful as probabilistic Turing machines with respect to polynomial-time, bounded-error computations [3], known methods for simulating bounded-error probabilistic computations with quantum machines are either very space-inefficient (as in the case of [3]) or very time-inefficient (as in the case of [12]) in the worst case. The apparent difficulty in simulating probabilistic computations with space-bounded quantum machines by means of the most straightforward technique (i.e., directly simulating coin-flips with appropriately defined quantum transformations) lies in the problem of reusing the space required for each coin-flip, of which there may be a number exponential in the space-bound.

We prove in this paper that quantum Turing machines can simulate a limited class of random processes—random walks on regular, undirected graphs—in a time-efficient and space-efficient manner. A random walk on a regular, undirected graph $G = (V, E)$ of degree d is a Markov chain defined as follows: the states of the Markov chain correspond to the vertices of G, and the transition probability from vertex u to vertex v is defined to be $1/d$ in case v is adjacent to u, and zero otherwise.

The study of random walks has had a number of interesting applications in complexity theory. From the perspective of this paper, the most important such application is due to Aleliunas, Karp, Lipton, Lovász and Rackoff [1], who used random walks to show that the undirected graph connectivity problem

*Part of this work was performed while the author was at the University of Wisconsin–Madison Computer Sciences Department under the support of NSF grant CCR-95-10244.

can be solved in $\mathrm{R}_H\mathrm{L}$ (sometimes written RL^{poly} or simply RL). Since this problem is complete for symmetric logspace (SL) with respect to logspace reductions [8], the relation $\mathrm{SL} \subseteq \mathrm{R}_H\mathrm{L}$ follows. The most space-efficient deterministic algorithm for USTCON requires space $O((\log n)^{4/3})$ [2]. We focus on the variant of this problem in which the graph in question is regular of a fixed degree d:

d-Regular Undirected Graph Connectivity (d-USTCON)

Instance:	A regular, undirected graph $G = (V, E)$ of degree d and $s, t \in V$.
Question:	Are s and t connected in G?

For $d \geq 3$, d-USTCON is SL-complete, as a straightforward reduction shows USTCON $\leq_m^{\log}$ d-USTCON.

By considering suitable quantum variants of random walks on graphs we prove d-USTCON $\in \mathrm{QR}_H\mathrm{L}$, which is the quantum analogue of $\mathrm{R}_H\mathrm{L}$. This is done in two steps. First we show d-USTCON can be solved with one-sided error by logspace quantum Turing machines that halt absolutely, but which have considerably worse acceptance probability than 1/2 for positive instances. We then demonstrate that $\mathrm{QR}_H\mathrm{L}$ is sufficiently robust with respect to acceptance probabilities to yield d-USTCON $\in \mathrm{QR}_H\mathrm{L}$. This implies the following containment.

Theorem 1 $\mathrm{SL} \subseteq \mathrm{QR}_H\mathrm{L}$.

Symmetric logspace is closed under complementation [10], which, together with Theorem 1, implies $\mathrm{SL} \subseteq \mathrm{QR}_H\mathrm{L} \cap \text{co-}\mathrm{QR}_H\mathrm{L} =: \mathrm{ZQ}_H\mathrm{L}$.

The remainder of this paper has the following organization. First, in Section 2 we review relevant facts concerning space-bounded quantum computation. In Section 3, we define a number of quantum operators and prove a lemma regarding these operators that will be useful in Section 4, which contains the construction of quantum Turing machines for simulating classical random walks on d-regular graphs. In Section 5, we address the issue of robustness of $\mathrm{QR}_H\mathrm{L}$ that, along with the machine constructed in Section 4, allows us to deduce the main theorem. Section 6 contains some concluding remarks.

2 Space-bounded quantum Turing machines

We begin by briefly reviewing some relevant facts concerning space-bounded quantum computation; for further information see [12]. For background on quantum computation more generally, we refer the reader to [3] and [4], and for classical space-bounded computation see [11].

The model of computation we use is the quantum Turing machine (QTM). Our QTMs have three tapes: a read-only input tape, a work tape, and a write only output tape. The input and work tape alphabets are denoted Σ and Γ, respectively, and the output is assumed to be in binary. Since our attention is restricted to decision problems, the only part of the output we care about is the first output bit: 1 indicates acceptance and 0 indicates rejection.

The behavior of a QTM is determined by a transition function, along with an observation of the output tape that is assumed to take place after each computation step. The computation continues as long as no output symbols are observed, and acceptance or rejection is determined (in a probabilistic sense) by the first output bit observed. There are strict conditions the transition function of a QTM must satisfy, as the evolution between observations must correspond to a norm-preserving operator on the Hilbert space spanned by classical configurations of the machine—see [3, 12] for further discussion.

A QTM M runs in logspace if there exists a function $f(n) = O(\log n)$ such that, for every input x, the position of the work tape head of M is never outside the range $[-f(|x|), f(|x|)]$ with nonzero amplitude during its computation on input x.

A QTM M halts absolutely if, for each input x, there exists $t = t(x)$ such that the probability that M halts (i.e., a 0 or a 1 is observed written to the output tape) during the first t steps of its computation on x is 1. If a QTM M halts absolutely, it must do so in time at most exponential in its space bound. In particular, a logspace QTM that halts absolutely necessarily runs in polynomial time.

The class $\mathrm{QR}_H\mathrm{L}$ consists of all languages A for which there exists a QTM M that, on each input x, runs in logspace, halts absolutely, and satisfies the following.

If $x \in A$, then M accepts x with probability at least $1/2$.

If $x \notin A$, then M accepts x with probability 0.

Substituting PTM for QTM in this definition yields the class $\mathrm{R}_H\mathrm{L}$. It is not currently known if $\mathrm{QR}_H\mathrm{L}$ and $\mathrm{R}_H\mathrm{L}$ are different, nor if one is contained in the other. In Section 5 we show that the $1/2$ in the above definition for $\mathrm{QR}_H\mathrm{L}$ may be replaced by any function $f(|x|)$ satisfying $f(|x|) \geq 1/g(|x|)$ and $f(|x|) \leq 1 - 2^{-g(|x|)}$ for some polynomial $g(|x|) > 0$.

The quantum Turing machines we construct will be described using pseudo-code in a manner typical for classical Turing machine descriptions. Computations will be composed of transformations of two types: *quantum transformations* and *reversible transformations* (both necessarily inducing norm-preserving operators on the associated Hilbert space). Quantum transformations will consist of a single step, so it will be trivial to argue that each quantum transformation can be performed as claimed. For reversible transformations, we rely on the result of Lange, McKenzie and Tapp [7], which implies that any logspace deterministic computation can be simulated reversibly in logspace. However, because the interference patterns produced by a given QTM depend greatly upon the precise lengths of the various computation paths comprising that machine's computation, we must take care to insure that these lengths are predictable in order to correctly analyze the computation. In the remainder of this section, we discuss reversible transformations somewhat more formally, and state a theorem based on the main result of [7] that will simplify this task.

For a given space-bound f and work tape alphabet Γ, define $W_{f(|x|)}(\Gamma)$ to be the set of all mappings of the form $w : \mathbb{Z} \to \Gamma$ taking the value # (blank) outside the interval $[-f(|x|), f(|x|)]$ (i.e., those mappings representing the possible contents of the work tape of a machine on input x having work tape alphabet Γ and running in space f). By a *reversible transformation*, we mean a one-to-one and onto mapping of the form $\Phi : W_{f(|x|)}(\Gamma) \to W_{f(|x|)}(\Gamma)$ for some f, x and Γ. For a given machine M having internal state set Q and work tape alphabet a superset of Γ, define $c(q, w)$ to be that configuration of M for which the work tape contents are described by w, the input and work tape heads are scanning the squares indexed by 0, the internal state is q, and no output has been written to the output tape. Now, we say that a deterministic Turing machine M on input x performs the transformation Φ on $W_{f(|x|)}(\Gamma)$ if the following conditions are satisfied:

1. The work tape alphabet of M is a superset of Γ.
2. The state set of M includes two distinguished states q_0 and q_f (the initial state and final state).
3. If M on input x is placed in a configuration $c(q_0, w)$ for any $w \in W_{f(|x|)}(\Gamma)$, then there exists some positive integer $t = t(x, w)$ such that if M is run for precisely t steps, it will then be in configuration $c(q_f, \Phi(w))$. Furthermore, at no time prior to step number t is the internal state of M equal to q_f.

Naturally, we say that t is the number of steps required for M on x to perform Φ. If the work tape head of M never leaves the region indexed by numbers in the range $[-g(|x|), g(|x|)]$ during this process, we say that M on x performs transformation Φ in space g.

Theorem 2 *Let $f(n) = O(\log n)$ and let M be a deterministic Turing machine that, on each input x, performs reversible transformation Φ_x on $W_{f(|x|)}(\Gamma)$ in space $O(\log |x|)$. Then there exists a reversible Turing machine M' that, on each input x, performs Φ_x on $W_{f(|x|)}(\Gamma)$ in space $O(\log |x|)$. Furthermore, the number of steps required for M' to perform Φ_x depends only on x and not on the particular argument of Φ_x.*

This theorem is based on a result due to Lange, McKenzie and Tapp [7], with added consideration payed to the number of steps required for transformations. See [12] (alternately [13]), along with [7] for a proof.

3 Quantum operators

In this section we define some operators and prove a key lemma that will be used in the analysis of the machines in the next section.

Throughout this subsection, assume $G = (V, E)$ is an undirected, regular graph of degree d that is not necessarily connected. The Hilbert space upon which the operators we define act is $\mathcal{H} = \ell_2(V \times V)$, i.e., the classical states of our space consists of all ordered pairs of vertices of G. Let $n = |V|$, $m = |E|$, and for each $u \in V$ define $S(u) = \{v \in V : \{u, v\} \in E\}$ and $B(u) = S(u) \cup \{u\}$. Each operator we consider is linear: we define the action of operators on the basis $\{|u, v\rangle : u, v \in V\}$ and extend to $\mathcal{H}$ by linearity.

First, define F as follows:

$$F|u,v\rangle = \begin{cases} |u,v\rangle - \dfrac{2}{d+1}\displaystyle\sum_{v' \in B(u)} |u,v'\rangle & v \in B(u) \\ |u,v\rangle & v \notin B(u). \end{cases}$$

We now verify that F is both unitary and hermitian. Define

$$|\psi_u\rangle = \frac{1}{\sqrt{d+1}} \sum_{v \in B_u} |u,v\rangle$$

for each $u \in V$. Note that $\{|\psi_u\rangle : u \in V\}$ is an orthonormal set. We may rewrite F as follows:

$$\begin{aligned} F &= \sum_{u \in V} \sum_{v \in B(u)} \left(|u,v\rangle - \frac{2}{d+1} \sum_{v' \in B(u)} |u,v'\rangle \right) \langle u,v| + \sum_{u \in V} \sum_{v \notin B(u)} |u,v\rangle \langle u,v| \\ &= \sum_{u,v \in V} |u,v\rangle \langle u,v| - \frac{2}{d+1} \sum_{u \in V} \left(\sum_{v,v' \in B(u)} |u,v'\rangle \langle u,v| \right) \\ &= I - 2 \sum_{u \in V} |\psi_u\rangle \langle \psi_u| . \end{aligned}$$

Consequently, each vector $|\psi_u\rangle$ is an eigenvector of F with eigenvalue -1, and every vector orthogonal to $\{|\psi_u\rangle : u \in V\}$ is an eigenvector of F with eigenvalue 1. From this it follows that F is both unitary

and hermitian: $F = F^\dagger = F^{-1}$. The operator F is related to the operator D defined on $\ell_2(\{0, \ldots, d\})$ as follows:

$$D\,|a\rangle = |a\rangle - \frac{2}{d+1}\sum_{b=0}^{d} |b\rangle\,.$$

Up to a sign change, this is the "diffusion" operator used in Grover's searching technique [6].

Next, define X as follows.

$$X = \sum_{u,v\in V} |v,u\rangle\,\langle u,v|\,.$$

The operator X simply exchanges the vertices u and v. Clearly $X = X^\dagger = X^{-1}$; X is unitary and hermitian.

Finally, define a projection P on $\mathcal{H}$ as

$$P = \sum_{u\in V} |u,u\rangle\,\langle u,u|\,.$$

Lemma 3 *Let $G = (V, E)$ be a regular graph of degree $d \geq 2$, let F, X and P be as defined above, define $Q = P\,F\,X\,F\,P$, and let $k \geq \frac{d(d+1)^2 n^2 \log(1/\epsilon)}{8}$ for given $\epsilon > 0$. For each $u \in V$, let $G_u = (V_u, E_u)$ denote the connected component of G containing u, and write $n_u = |V_u|$. Then for every $u \in V$ we have*

$$\left\| Q^k\,|u,u\rangle - \frac{1}{n_u}\sum_{v\in V_u} |v,v\rangle \right\| < \epsilon.$$

Proof. First, we note that

$$\begin{aligned}
Q\,|u,u\rangle &= P\,F\,X\,F\,P\,|u,u\rangle \\
&= P\,F\,X\,F\,|u,u\rangle \\
&= P\,F\,X\left(|u,u\rangle - \frac{2}{d+1}\sum_{v\in B(u)} |u,v\rangle\right) \\
&= P\,F\left(|u,u\rangle - \frac{2}{d+1}\sum_{v\in B(u)} |v,u\rangle\right) \\
&= P\left(|u,u\rangle - \frac{2}{d+1}\sum_{v'\in B(u)} |u,v'\rangle - \frac{2}{d+1}\sum_{v\in B(u)}\left(|v,u\rangle - \frac{2}{d+1}\sum_{v'\in B(v)} |v,v'\rangle\right)\right) \\
&= \left(1 - \frac{2}{d+1}\right)^2 |u,u\rangle + \left(\frac{2}{d+1}\right)^2 \sum_{v\in S(u)} |v,v\rangle \qquad (1)
\end{aligned}$$

for each $u \in V$, and clearly $Q\,|u,v\rangle = 0$ for $u \neq v$.

For given $u \in V$ we have that $v \notin V_u$ implies $\left\langle v,v \,\middle|\, Q^l \,\middle|\, u,u\right\rangle = 0$ for $l = 1$, and a simple induction shows that this holds for any $l \geq 1$. For each u, define T_u to be a projection operator as follows:

$$T_u = \sum_{v\in V_u} |v,v\rangle\,\langle v,v|\,.$$

Defining $Q_u = T_u Q T_u$, we therefore have $Q_u^l \left|u,u\right\rangle = Q^l \left|u,u\right\rangle$ for $t \geq 0$. Note that Q_u is hermitian: $Q_u^\dagger = (T_u P F X F P T_u)^\dagger = T_u P F X F P T_u = Q$, following from the fact that T_u, P, F, and X are hermitian.

Let A denote the adjacency matrix of G_u and let f_A denote the characteristic polynomial of A. By (1), we determine that f_{Q_u}, the characteristic polynomial of Q_u, satisfies

$$\begin{aligned} f_{Q_u}(z) &= z^{(n^2-n_u)} \det\left(z\, I_{n_u} - \left(1 - \frac{2}{d+1}\right)^2 I_{n_u} - \left(\frac{2}{d+1}\right)^2 A \right) \\ &= z^{(n^2-n_u)} \left(\frac{2}{d+1}\right)^{2n_u} \det\left(\frac{(d+1)^2 z - (d-1)^2}{4}\, I - A \right) \\ &= z^{(n^2-n_u)} \left(\frac{2}{d+1}\right)^{2n_u} f_A\left(\frac{(d+1)^2 z - (d-1)^2}{4} \right). \end{aligned}$$

Letting $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{n_u}$ be the eigenvalues of A, we see that Q_u has eigenvalues

$$\mu_j = \frac{4\lambda_j + (d-1)^2}{(d+1)^2},$$

for $j = 1, \ldots, n_u$, as well as eigenvalues $\mu_j = 0$ for $j = n_u + 1, \ldots, n^2$. Note that the eigenvalues of A (and hence the eigenvalues of Q_u) are real since A is symmetric. Since G_u is connected and regular of degree d, we have $\lambda_1 = d$, $\lambda_j < d$ for $j = 2, \ldots, n_u$, and $\lambda_{n_u} \geq -d$ (see, e.g., [5], page 14). Furthermore, it follows from [9] that

$$\lambda_j \;\leq\; d - \frac{2}{d n_u^2},$$

for $j = 2, \ldots, n_u$. Hence $\mu_1 = 1$, and

$$\mu_j \in \left[1 - \frac{8d}{(d+1)^2},\; 1 - \frac{8}{d(d+1)^2 n_u^2}\right]$$

for $j = 2, \ldots, n_u$. In particular, we have that $\mu_2, \ldots, \mu_{n_u}$ are bounded in absolute value by

$$1 - \frac{8}{d(d+1)^2 n_u^2}.$$

Define

$$\left|\phi_1\right\rangle = \frac{1}{\sqrt{n_u}} \sum_{u \in V_u} \left|u,u\right\rangle.$$

It may be verified that $\left|\phi_1\right\rangle$ is an eigenvector of Q_u corresponding to the eigenvalue $\mu_1 = 1$. As Q_u is hermitian, we may choose eigenvectors $\left|\phi_2\right\rangle, \ldots, \left|\phi_{n^2}\right\rangle$ corresponding to eigenvalues $\mu_2, \ldots, \mu_{n^2}$ in such a way that $\{\left|\phi_1\right\rangle, \ldots, \left|\phi_{n^2}\right\rangle\}$ is an orthonormal basis of $\mathcal{H}$. Now, let $c_j = \left\langle \phi_j \right| u,u\rangle$ for $j = 1, \ldots, n^2$. We have $\left|u,u\right\rangle = \sum_{j=1}^{n^2} c_j \left|\phi_j\right\rangle$, and thus

$$Q_u^l \left|u,u\right\rangle = \sum_{j=1}^{n_u} c_j \mu_j^l \left|\phi_j\right\rangle,$$

for $l \geq 1$. Consequently,

$$\left\| Q_u^l \left| u,u \right\rangle - \frac{1}{n_u} \sum_{v \in V_u} \left| v,v \right\rangle \right\|^2 = \left\| \sum_{j=2}^{n_u} c_j \mu_j^l \left| \phi_j \right\rangle \right\|^2 = \sum_{j=2}^{n_u} |c_j|^2 |\mu_j|^{2l} \leq \left(1 - \frac{8}{d(d+1)^2 n_u^2} \right)^{2l}, \quad (2)$$

again for $l \geq 1$.

Now, since $k \geq \frac{d(d+1)^2 n^2 \log(1/\epsilon)}{8}$ for given $\epsilon > 0$, we have

$$\left(1 - \frac{8}{d(d+1)^2 n_u^2} \right)^k \leq \left(1 - \frac{8}{d(d+1)^2 n^2} \right)^k < \epsilon,$$

for every u, following from the fact that $(1 - 1/x)^x < 1/e$ for $x \geq 1$. Thus

$$\left\| Q_u^k \left| u,u \right\rangle - \frac{1}{n_u} \sum_{v \in V_u} \left| v,v \right\rangle \right\| < \epsilon$$

follows by (2). Since $Q^k \left| u,u \right\rangle = Q_u^k \left| u,u \right\rangle$, this completes the proof. ■

By taking $\epsilon = \frac{1}{2n}$ in Lemma 3, we obtain the following corollary.

Corollary 4 *Let $G = (V,E)$ be a regular graph of degree $d \geq 2$ with $s,t \in V$, let Q be as defined in Lemma 3, and let $k \geq \lceil d(d+1)^2 n^2 \log(2n)/8 \rceil$. If s and t are connected in G, then*

$$\left| \left\langle t,t \right| Q^k \left| s,s \right\rangle \right|^2 \geq \frac{1}{4n^2},$$

and otherwise $\left| \left\langle t,t \right| Q^k \left| s,s \right\rangle \right|^2 = 0$.

4 Quantum Turing machine construction and analysis

We now construct, for each fixed degree $d \geq 2$, a logspace QTM solving d-USTCON that operates with one-sided error and halts absolutely. Although the QTMs we construct have somewhat poor probabilities of acceptance for positive instances of d-USTCON, it will be demonstrated in the next section that these machines may be modified to yield logspace QTMs for d-USTCON having arbitrarily small one-sided error while still halting absolutely.

Lemma 5 *For $d \geq 2$, there exists a quantum Turing machine M that runs in logspace, halts absolutely, and operates as follows. For any input encoding (G,s,t), where $G = (V,E)$ is a regular, undirected graph of degree d, $s,t \in V$, and s is connected to t in G, M accepts with probability greater than $\frac{1}{4n^2}$, and for all other inputs M accepts with probability zero.*

Proof. The work tape of M will consist of four tracks, one for each of the following variables: u, v, b and c. Each variable will contain an integer, with the exception of v, which will store either an integer or a single symbol in the set $\{0,\ldots,d\}$. Integers are assumed to be encoded as strings over the alphabet $\{0',1'\}$, taken to be disjoint from $\{0,\ldots,d\}$. We make the assumption that each integer has exactly one encoding and that 0 is encoded by the empty string. Note that this implies u, v, b and c are all initially set to 0, as the work tape initially contains only blanks. Vertices of G are assumed to be labeled by

1. Reject if the input does not encode (G, s, t) for G undirected and regular of degree d.
2. Set $u = u + s$ and $v = v + s$.
3. Loop with starting/stopping condition "b=0":
   - i. If $v \in B(u)$, replace v with the symbol in $\{0, \ldots, d\}$ corresponding to its index in $B(u)$ modulo $d + 1$.
   - ii. If $v \in \{0, \ldots, d\}$, perform transformation D on v.
   - iii Invert step i.
   - iv. Exchange u and v.
   - v. If $v \in B(u)$, replace v with the symbol in $\{0, \ldots, d\}$ corresponding to its index in $B(u)$ modulo $d + 1$.
   - vi. If $v \in \{0, \ldots, d\}$, perform transformation D on v.
   - vii. Invert step v.
   - viii. In case $u \neq v$, increment c modulo $d(d+1)^2n^3 + 1$.
   - ix. Increment b modulo $d(d+1)^2n^3$.
4. If $c = 0$ and $u = t$, then *accept*, else *reject*.

Figure 1: Description of quantum Turing machine M for Lemma 5.

integers having length at most logarithmic in the input size, and each vertex has a unique label. When u or v contains an integer, this integer is to be interpreted as the label of a vertex.

The execution of M is described in Figure 1. For each of the steps in Figure 1 we may define an appropriate reversible or quantum transformation corresponding to the action described. Each transformation is to maintain the invariant that all tracks contain strings having no embedded blanks and having leftmost symbol stored in the work tape square indexed by 0. The quantum transformations are steps ii and vi—these transformations require a single step and involve only the symbol in square 0 of the track corresponding to v. The remaining transformations are reversible transformations—it is straightforward to note that each such transformation may be performed by a DTM running in space $O(\log n)$ in the manner described in Section 2 for a suitable space-bound $f(n) = O(\log n)$. (It is for this reason that we increment c modulo $d(d+1)^2n^3+1$ instead of simply incrementing c in step viii, although the same effect results; each transformation must be defined on a bounded region of the work tape). We note that the quantity $d(d+1)^2n^3$ is somewhat arbitrary in steps viii and ix—any quantity at least $\lceil d(d+1)^2n^2 \log(2n)/8 \rceil$ suffices. The loop may be implemented reversibly, in the manner described in [12]. By Theorem 2, it follows that each reversible step in Figure 1 may be performed reversibly in logspace, requiring time depending only on the input (G, s, t) and not on the particular contents of the work tape of M when the step is performed. Two consequences of this are (i) M runs in logspace, and (ii) each step in Figure 1 may be viewed as requiring unit time, insofar as the analysis of the machine is concerned.

Now let us analyze the computation of M on a given input (G, s, t). When describing superpositions of M, we will restrict our attention to the variables u, v, b and c; since we will only care about superpositions between the transformations described above, all other aspects of M (specifically, tape head positions and internal state) are deterministic. It will be most convenient to express such superpositions in terms of classical states of the form $|u, v\rangle |c\rangle |b\rangle$ for $u, v \in V$, $c, b \in \mathbb{Z}$, which may be interpreted as being equivalent to classical states the form $|u, v, c, b\rangle$.

Assume that M does not reject during step 1, so that G is indeed regular of degree d and undirected—otherwise M of course functions as required. After step 2 is performed, the superposition of M is $|s,s\rangle|0\rangle|0\rangle$. Now the loop in step 3 is performed. It may be verified that after one iteration of the loop, the superposition of M is $(Q|s,s\rangle)|0\rangle|1\rangle + |\xi_{1,1}\rangle|1\rangle|1\rangle$, where Q as defined in Section 3 and $|\xi_{1,1}\rangle$ is some vector (that we don't care about). More generally, after $j < d(d+1)^2n^3$ iterations of the loop, the superposition is

$$\left(Q^j|s,s\rangle\right)|0\rangle|j\rangle + \sum_{c\geq 1}|\xi_{c,j}\rangle|c\rangle|j\rangle\,,$$

and after $k = d(d+1)^2n^3$ iterations, the superposition is

$$\left(Q^k|s,s\rangle\right)|0\rangle|0\rangle + \sum_{c\geq 1}|\xi_{c,0}\rangle|c\rangle|0\rangle\,.$$

At this point, the loop terminates, so that in step 4 the probability of accepting is $\left|\langle t,t|\,Q^k|s,s\rangle\right|^2$. By Lemma 3, we conclude that M accepts (G,s,t) with probability at least $\frac{1}{4n^2}$ in case s is connected to t, and probability 0 otherwise. ■

5 Amplification of Acceptance Probability

It is well-known that $\mathrm{R}_H\mathrm{L}$ is robust with respect to the probability with which positive instances are accepted: the 1/2 in the definition of $\mathrm{R}_H\mathrm{L}$ may be replaced by any function $f(|x|)$ satisfying $f(|x|) \geq 1/g(|x|)$ and $f(|x|) \leq 1 - 2^{-g(|x|)}$ for $g(|x|) > 0$ a polynomial. It is not immediate that an analogous fact holds for $\mathrm{QR}_H\mathrm{L}$; repeated simulation a given QTM computation requires that the simulated machine be in its initial configuration at the start of each simulation, but resetting this machine to its initial configuration constitutes an irreversible action that cannot be performed by the quantum machine performing the simulation. It is, however, not difficult to show that the analogous fact for $\mathrm{QR}_H\mathrm{L}$ does hold by using a method described in [12] for showing related facts regarding non-halting versions of space-bounded classes. Since the proof in the present case is similar, we will just sketch the proof.

Lemma 6 *Let M be a QTM that runs in logspace, halts absolutely, and accepts each input x with probability $p(x)$. Then for any polynomial f, there exists a QTM M_f that runs in logspace, halts absolutely, and accepts each input x with probability*

$$1 - (1-p(x))(1-2p(x))^{2f(|x|)}.$$

Proof. [Sketch] Given M and f as in the statement of the theorem, we let M_f be a quantum Turing machine functioning as described in Figure 2. We may assume without loss of generality that M halts after precisely $g(|x|)$ steps along all computation paths for g bounded by a polynomial and computable in logspace [12], from which it follows that we may take M_f to run in logspace and halt absolutely.

During the simulation of M, M_f will store an encoding of some configuration of M (including the first square on the output tape of M) on its work tape. Let us denote by E an operator corresponding to performing step i on the state of M stored by M_f; since M does not produce output during the simulation, we may take E to be unitary. Initially, the state of M represented by M_f is $|c_0\rangle$, for c_0 the initial configuration of M. After performing step i for the first time, this state is mapped to

1. Repeat the following $f(|x|)+1$ times:
   i. Simulate the computation of M on x.
   ii. Accept if M accepts x.
   iii. Invert step i.
   iv. If the current configuration of M is not the initial configuration, multiply the current amplitude by -1.
2. Reject.

Figure 2: Description of quantum Turing machine M_f for Lemma 6.

$E|c_0\rangle = |\psi\rangle$. Write $|\psi\rangle = |\psi_{acc}\rangle + |\psi_{acc}^{\perp}\rangle$, where $|\psi_{acc}\rangle$ denotes the projection of $|\psi\rangle$ onto the space spanned by accepting configurations of M and $|\psi_{acc}^{\perp}\rangle$ is orthogonal to $|\psi_{acc}\rangle$. During step ii, M_f accepts x with probability $p(x) = \| |\psi_{acc}\rangle \|^2$, and otherwise the superposition of M_f collapses so that $|\psi_{acc}^{\perp}\rangle$ is the state of M represented.

Now we consider the sequence of steps iii, iv, i, ii, which are at this point performed $f(|x|)$ times. It can be checked that each iteration of the sequence of steps ii, iv, i has the effect of mapping $|\psi_{acc}^{\perp}\rangle$ to $(2-2p(x))|\psi_{acc}\rangle + (1-2p(x))|\psi_{acc}^{\perp}\rangle$ (following from the fact that $\langle c_0 \,|\, E^{\dagger} \,|\, \psi_{acc}\rangle = p(x)$). The j-th time step i is performed, for $j \geq 2$, thus results in acceptance with probability $p(x)(1-2p(x))^{2(j-2)}(2-2p(x))^2$. We therefore have that M_f accepts with probability

$$p(x) + \sum_{j=2}^{f(|x|)+1} p(x)(1-2p(x))^{2(j-2)}(2-2p(x))^2 \;=\; 1-(1-p(x))(1-2p(x))^{2f(|x|)}$$

as claimed. ■

By Lemma 5 and Lemma 6, we have d-USTCON $\in \mathrm{QR}_H\mathrm{L}$.

Theorem 1 now follows in straightforward fashion, relying again on Theorem 2; given a particular language $A \in \mathrm{SL}$ we have a logspace many-one reduction to d-USTCON, and we may replace various reversible transformations of our machine for d-USTCON with appropriately defined reversible transformations based on compositions of the reduction with the replaced transformation. Details will appear in the final version of this paper.

6 Concluding Remarks

In this paper we have shown that logspace quantum Turing machines can simulate a limited class of probabilistic computations in a time-efficient manner. This leaves open the question of whether probabilistic computations can be simulated efficiently by space-bounded quantum machines in general (e.g., is $\mathrm{R}_H\mathrm{L}$ contained in $\mathrm{QR}_H\mathrm{L}$?) We hope our techniques will provide insight into this problem.

We have defined in this paper quantum processes that attempt to mimic classical random walks on graphs. There are a number of ways in which to define *quantum walks on graphs* having properties quite different from classical random walks. It may be interesting to consider possible applications of such processes to quantum complexity theory.

References

[1] R. Aleliunas, R. Karp, R. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the time complexity of maze problems. In *Proceedings of the 20th Annual Symposium on Foundations of Computer Science*, pages 218–223, 1979.

[2] R. Armoni, A. Ta-Shma, A. Wigderson, and S. Zhou. $SL \subseteq L^{4/3}$. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 230–239, 1997.

[3] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997.

[4] A. Berthiaume. Quantum computation. In L. Hemaspaandra and A. Selman, editors, *Complexity Theory Retrospective II*, pages 23–50. Springer, 1997.

[5] N. Biggs. *Algebraic Graph Theory*. Cambridge University Press, 1974.

[6] L. Grover. A fast quantum mechanical algorithm for database search. In *28th Annual ACM Symposium on the Theory of Computing*, pages 212–219, 1996.

[7] K. Lange, P. McKenzie, and A. Tapp. Reversible space equals deterministic space (extended abstract). In *Proceedings of the 12th Annual IEEE Conference on Computational Complexity*, pages 45–50, 1997. To appear in Journal of Computer and System Sciences.

[8] H. Lewis and C. Papadimitriou. Symmetric space-bounded computation. *Theoretical Computer Science*, 19:161–187, 1982.

[9] L. Lovász and P. Winkler. Mixing of random walks and other diffusionson a graph. In Peter Rowlinson, editor, *Surveys in Combinatorics*, volume 218 of *London Mathematical Society Lecture Note Series*, pages 119–154. Cambridge University Press, 1995.

[10] N. Nisan and A. Ta-Shma. Symmetric logspace is cloased under complement. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 140–146, 1995.

[11] M. Saks. Randomization and derandomization in space-bounded computation. In *Proceedings of the 11th Annual IEEE Conference on Computational Complexity*, pages 128–149, 1996.

[12] J. Watrous. Space-bounded quantum complexity. *Journal of Computer and System Sciences*. To appear. A preliminary version appeared in *Proceedings of the 13th Annual IEEE Conference on Computational Complexity*, 1998, pages 210–227.

[13] J. Watrous. *Space-Bounded Quantum Computation*. PhD thesis, University of Wisconsin – Madison, 1998.