Quantum Random Walks

Edwin Ng*
Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139
(Dated: May 5, 2012)

The notion of a quantum random walk has received notable attention recently from the field of quantum information, revealing many surprising and subtle properties that range from significant speedups in solving simple graph problems to complex behavior under decoherence. We present here an overview of the topic, starting with an introduction to classical walks and then generalizing to the discrete time and continuous time formulations of quantum walks. We focus primarily on the formalism used to describe these two models, while briefly discussing the more elementary properties via comparison with the classical case.

I. INTRODUCTION AND OVERVIEW

The concept of a classical random walk is pervasive in modern computer science, as they provide effective solutions to problems involving probabilistic sampling and searching. It is interesting to note, however, that classical walks have made their way into physics as well: random walks have been used to model statistical systems such as Brownian motion (diffusion) and the distribution of magnetic moments in a lattice.[1] In light of these applications, the properties of classical walks are naturally in accordance with the behavior of classical systems.

Quantum random walks, on the other hand, derive their properties from the laws of quantum mechanics. Although we will introduce quantum walks here in somewhat abstract terms and largely ignore the question of their physical realization, quantum walks are very intuitive models for many quantum systems—a particle propagating on a lattice, for example.

An early conception of a quantum walk appeared in Feynman's 1986 article[2] on the prospects of quantum computation, but it was not until 1993 that the usual formulation for quantum walks was published by Aharonov, Davidovich, and Zagury.[3] Several interesting results were already apparent, but attention to quantum walks finally took root in 1998 when Farhi and Gutmann[4] (and, later, Childs et al.[5]) introduced a series of graph problems which lent themselves to an attractive quantum walk solution. It is now known[6, 7] that the quantum random walk is universal, in the sense that it is equivalent to other proposed models of quantum computing, such as the circuit or adiabatic model.

As a result of these developments, however, we have at our disposal today two distinctly different formulations of quantum walks, divided according to the way the walk propagates. The first, due to Aharonov et al., is the *discrete time* random walk, which is more computational in flavor and features discrete, unitary dynamics. The second, used by Farhi et al., is the *continuous time* random walk, which is directly connected to the usual unitary time evolution through the Schrödinger equation.

* ngedwin@mit.edu

Generally speaking, the properties of quantum walks on the line and other simple structures are well understood. There are many important aspects of quantum walks—forming the bulk of present research—which we are unable to discuss here, as they require a lengthy introduction of concepts (e.g., hitting and mixing times) from the study of classical walks before we can generalize to the quantum case; simple proofs are also generally difficult to come by in studying quantum walks. Instead, our focus will be on the formalism that has been developed for quantum walks, giving some elementary results where possible. These results will already be clear evidence of the rich behavior of quantum walks.

The discussion in this paper will closely follow several recent introductory articles on quantum walks. In particular, many of the examples will be based on the ones in Kempe's article [8]. For a particularly complete bibliography on the subject, see [9].

We begin with a brief introduction to classical random walks before extending the discussion to the formulations of the discrete and continuous time quantum random walks.

II. CLASSICAL RANDOM WALKS

The simplest example of a classical random walk is the random walk on the line. Suppose we have a particle whose location at various discrete time steps can be described by an integer z_k, for the kth time step. The particle starts at position $z_0 = 0$, and at every time step, flips a biased coin, moving to the right with probability p and to the left with probability $1 - p$. That is, the particle moves according to

$$z_{k+1} = \begin{cases} z_k + 1 & \text{with probability } p \\ z_k - 1 & \text{with probability } 1 - p \end{cases}. \tag{1}$$

The process then iterates, giving us a sequence of random variables z_k. We note that the particle's motion at any time step does not depend on its motion in previous time step, simplifying our analysis.

Suppose we let the particle perform this walk for n time steps. We can then ask for the probability $\Pr(z_n{=}z)$ that the particle has moved exactly z steps to the right.

For this to happen, it must have followed some sequence of n steps with z more right steps than left steps. The probability of such a sequence is $p^{(n+z)/2}(1-p)^{(n-z)/2}$, and since any such sequence will do, we can choose any combination of $(n+z)/2$ out of n steps to use for our right moving ones. Thus,

$$\Pr(z_n{=}z) = \frac{n! \cdot p^{(n+z)/2}(1-p)^{(n-z)/2}}{\left(\frac{n+z}{2}\right)!\left(\frac{n-z}{2}\right)!}, \qquad (2)$$

where we specifically note that $\Pr(z_n{=}z) = 0$ when $n \pm z$ is odd. This is because we cannot move an odd number of steps in an even number of time steps and vice-versa, a fact reflected in the combinatorial prefactors. Some sample probability distributions are plotted in Figure 1.

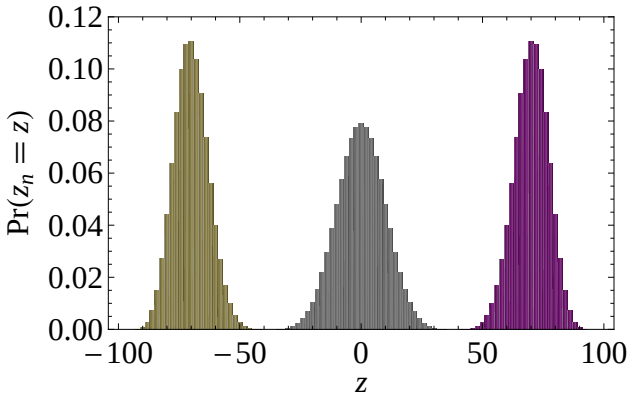

FIG. 1. Sample probability distributions of taking z steps to the right after $n = 100$ time steps in the biased classical random walk on a line. From left to right, the values of p are 0.15, 0.5, and 0.85. Note that only the even values of z are represented.

From Equation 2, we can deduce various properties of this random walk. The expected value of the variable z after n steps is given by

$$\mathrm{Ex}(z) = \sum_{z=-n}^{n} z \cdot \Pr(z_n{=}z) = n\,(2p-1)\,,$$

while its variance is

$$\mathrm{Var}(z) = \mathrm{Ex}(z^2) - \mathrm{Ex}^2(z) = 4np(1-p).$$

The upshot here is that $\sqrt{\mathrm{Var}(z)}$, which measures the essential "spreading" of the position of the particle (in terms of its root-mean-square displacement), goes as $O(\sqrt{n})$, up to constants. This is the signature of a classical random walk, and one of the first properties to change when we switch over to the quantum version.

Upon reflection, there are some obvious ways to generalize the walk on the line. Suppose that instead of a line, we use a *graph*, which is a structure $G = (V, E)$ where V is a set of *vertices* and $E \subseteq V \times V$ is the set of *edges* connecting them. If it is possible for a particle to move from v_1 to v_2, then $(v_1, v_2) \in E$, and we say that v_1 and v_2 are connected by a directed edge. To describe the walk itself, we could, for example, introduce transition probabilities p_j of moving along the edge e_j.

Of course, our random walk on the line is just a special case of a walk on a graph, since we can take V to be the set of all integers z and E to be the set of all ordered pairs $(z, z \pm 1)$. The probabilities are p for right moving edges of the form $(z, z+1)$ and $1-p$ for left moving edges of the form $(z, z-1)$. See Figure 2.a for an illustration.

Classically, nothing prevents us from considering random walks on very convoluted graphs, with intricate directed cycles, multiple edges between two vertices, and possibly even weighted edges. But the graphs which are most conducive to analysis (especially in the quantum case) are usually *undirected* and *regular*.

An undirected graph $G = (V, E)$ has edges without preferred direction, so that if $(v_1, v_2) \in E$, then $(v_2, v_1) \in E$ also. For example, in the general case of the biased classical walk analyzed above, the edges were directed, in order to differentiate between left and right moving edges. However, if we were to consider only the unbiased walk using a fair coin, then we can just place a single edge between any two neighboring vertexs, giving us an undirected graph.

In a regular graph, every vertex has the same number of neighbors. More precisely, if the *degree* of a vertex v is the number of edges in E that are connected to v, then every vertex in a regular graph has the same degree d. This is a nice property because we have a "default" transition probability scheme, using $p_j = 1/d$.

The most common examples of undirected regular graphs are the cycle graphs C_n, with n vertices on an undirected ring structure (as in Figure 2.b) and hypercubes, which are higher-dimensional analogues of the graph C_4. The infinite walk on the line with transition probabilities $p = 1/2$ would also count as an undirected and regular (though infinite) graph.

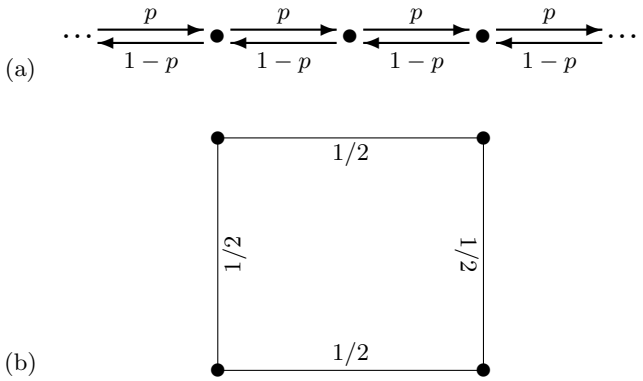

FIG. 2. Some examples of graphs. The graph in (a) is the infinite line graph used for the one-dimensional walk on the line. Shown in (b) is the cycle graph C_4, with equal transition probabilities from each vertex to its neighbor. Note that each line on C_4 represents *two* directed but symmetric edges, so we simply replace them by a single line without arrowheads.

III. DISCRETE TIME QUANTUM WALKS

As in the classical case, we formulate the discrete time quantum walk by starting with the walk on a line. Consider a quantum mechanical particle constrained to a discrete line, whose state $|\phi\rangle$ lies in the Hilbert space $\mathcal{H}_p$ with natural basis elements $|z\rangle$ for z an integer. Physically, this can be thought of as an electron tightly bound to an infinite periodic ion potential, for example.

However, rather than allowing for the unitary time evolution of this single particle according to Schrödinger's equation, we furthermore introduce a *coin system*, to play the role of dictating the transitions analogously to the classical coin. Specifically, the coin we use for the walk on the line is a two-state (i.e., spin-1/2) system $|\chi\rangle$ in the space $\mathcal{H}_c$ with natural basis elements $|\uparrow\rangle$ and $|\downarrow\rangle$.

Thus, the total wavefunction for a quantum walk on the line is a joint state $|\psi\rangle \in \mathcal{H}$, where $\mathcal{H} = \mathcal{H}_p \otimes \mathcal{H}_c$. Suppose now that we begin with the state $|\psi_0\rangle = |\chi_0\rangle|\phi_0\rangle$ on time step $k = 0$ (initially entangled states are possible but will not be considered here). How do we evolve the system to the next time step, to obtain $|\psi_1\rangle$ and more generally $|\psi_k\rangle$?

Following the laws of quantum mechanics, we are free to choose any unitary evolution on $\mathcal{H}$. We will define the evolution of the discrete time quantum walk to be

$$|\psi_k\rangle = U^k|\psi_0\rangle \quad \text{with} \quad U = S\,(C \otimes I)$$

where S is the *conditional shift* that takes a particular coin-particle state and moves it, advancing a single step. Before taking this step, however, we first apply the *coin flip operator* C on the coin, which mimics the flipping of the classical coin in Equation 1.

For a quantum walk on the line, the conditional shift which best represents the classical intuition is the one that moves the particle to the nearest neighbor depending on the state of the coin, realized by

$$S = |\uparrow\rangle\langle\uparrow| \otimes \sum_j |j+1\rangle\langle j| + |\downarrow\rangle\langle\downarrow| \otimes \sum_j |j-1\rangle\langle j|.$$

On the other hand, we have remarkable freedom in choosing C. One popular choice, motivated by the classical coin, is the *Hadamard coin operator*

$$C = \frac{|\uparrow\rangle\langle\uparrow| + |\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow| - |\downarrow\rangle\langle\downarrow|}{\sqrt{2}}, \tag{3}$$

which essentially takes any eigenstate of σ_z and sends it into an equal superposition of eigenstates of σ_x.

Thus far, everything we have described is strictly deterministic—the system simply evolves under a series of discrete unitary transfromations. The *random* aspect of the walk therefore lies in *measurement*. Suppose we measure the value of z on the particle after n time steps, and we denote the measurement result with the random variable z_n. Then the probability distribution, analogous to Equation 2, follows from the Born rule:

$$\Pr(z_n{=}z) = |\langle z|\psi_n\rangle|^2 = |\langle z|U^n|\psi_0\rangle|^2.$$

A plot of such a probability distribution is shown in Figure 3. The difference between Figure 1 and 3 is striking and serves to highlight many of the interesting properties of quantum walks.

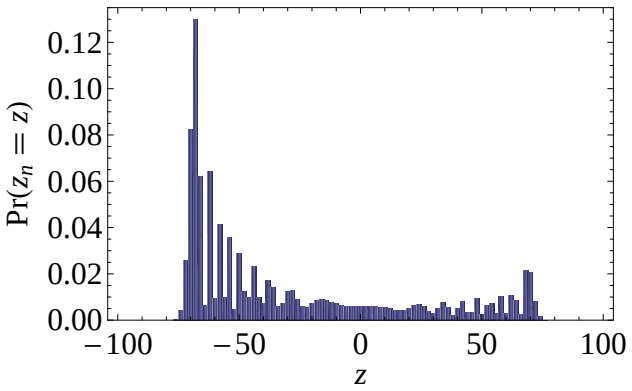

FIG. 3. A plot of the probability distribution for making a z measurement after $n = 100$ time steps in a quantum random walk on the line. The initial state is $|\uparrow\rangle|0\rangle$ and C is the Hadamard coin operator. Note that only the even values of z are represented.

Whereas the classical walk formed an even binomial-like distribution about its starting point, the quantum walk is strongly skewed. For the case in Figure 3, we have $\mathrm{Ex}(z) \approx -28.98$. Because we started with the state $|\uparrow\rangle|0\rangle$, the distribution is left-skewed; it turns out that starting with the state $|\downarrow\rangle|0\rangle$ leads to the same pattern but horizontally flipped, so that the distribution becomes right-skewed instead.

In fact, as explained in [8], the cause of the asymmetry is our choice of the coin. Because of the phase in the Hadamard coin operator in Equation 3, as long as we start the coin in a definite state of spin-up or down, we will end up biasing the walk. In order to remove the asymmetry, we need to use a superposition of these two walks, and furthermore ensure that they cannot interfere with each other. For example, we can start the coin in the up-eigenstate of the σ_y operator

$$|\chi_0\rangle = \frac{|\uparrow\rangle + i|\downarrow\rangle}{\sqrt{2}} = |+_y\rangle,$$

This initial state results in the symmetric distribution shown in Figure 4. Another solution is to change our coin operator; for example, $C = (I + i\sigma_x)/\sqrt{2}$ is a balanced coin operator that will not bias coins in the z basis.

Remarkably, even with a symmetric coin, the pattern still looks drastically different from the classical walk. Clearly, the qualitative reason for the patterns we see in Figure 4 are due to quantum interference during the determinstic steps of the walk. But these effects are finely controlled by a number of parameters in our definition of the quantum walk—our choice of initial state, the coin operator we use, and the number of steps we take before measurement. These all play a role in contributing to the rich structure we see in these distributions.

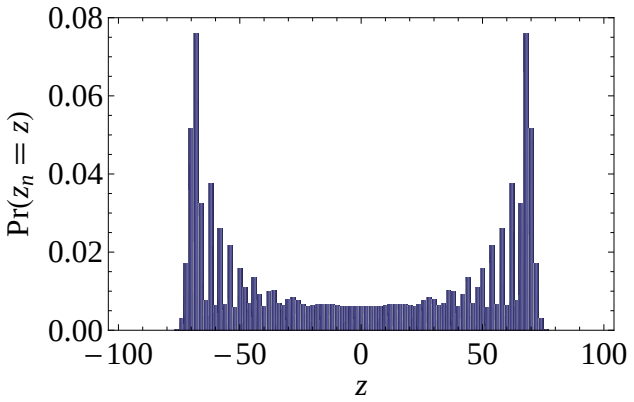

FIG. 4. A plot of the probability distribution for making a z measurement after $n = 100$ time steps in a quantum random walk on the line. The initial state is $|+_y\rangle|0\rangle$ and C is the Hadamard coin operator. Note that only the even values of z are represented.

One of the first properties of quantum walks to intrigue the community was its statistics, which also show sharp contrast to the classical case. A simple example of such a property is immediately evident from looking at the root-mean-square displacements of the walks in Figures 3 and 4. For the system in Figure 3, we have $\sqrt{\mathrm{Var}(z)} \approx 45.71$, while the symmetric walk reaches $\sqrt{\mathrm{Var}(z)} \approx 54.12$. Compared to an RMS displacement of merely 10 for the unbiased classical walk, the behavior of the quantum walk is clearly non-diffusive. Quantum walks can effectively "spread" much faster than their classical counterparts, all due to quantum interference!

In general, obtaining an exact probability distribution analogous to Equation 2 is extremely difficult for quantum random walks. Nevertheless, there have been a number of successful analytical methods, one of which uses a discrete path-integral approach to account for all possible paths and then analyze the asymptotic behavior.[10] For the Hadamard walk we considered here, it has been shown that the variance after n steps goes as $O(n^2)$ compared with $O(n)$ for the classical walk, so indeed, the quantum walk propagates quadratically faster.

Another case of interest where quantum walks differ from their classical counterparts is in the presence of an absorbing boundary. Suppose that instead of having the classical particle walk on an infinite line, we allow the walk to proceed freely only in the region $z < b$ for some integer b. If at any time step n, the particle reaches $z = b$, we stop the walk and say the particle has been absorbed. For this classical situation, it can be shown that the particle is doomed: the probability that the particle will be absorbed in finite time is one. Put another way, the absorbing classical walk is guaranteed to end at some finite time step n.

To see why this is true, consider a particle undergoing an unbiased classical walk on the line, starting at $z = 1$ and taking $b = 0$ for simplicity. Let us denote by p^* the probability of ever hitting b from $z = 1$. The contribution to p^* on the first step is $1/2$ if the particle moves to $z = 0$; otherwise with probability $1/2$, the particle moves to $z = 2$. The probability of absorption *given* that the latter event happens is the probability of ever reaching $z = 1$ from $z = 2$ times the probability of ever reaching $z = 0$ from $z = 1$. Both of these two probabilities is given again by p^*, so we have the recursion

$$p^* = \frac{1}{2} + \frac{1}{2} \cdot p^* \cdot p^* \quad \Rightarrow \quad p^* = 1.$$

Thus the probability of hitting the boundary is unity. It can be shown that this analysis works for any boundary, and the classical particle has zero probability of escape.

We can also formulate the concept of an absorbing barrier for the quantum case. At the end of each time step, we perform a *measurement* of the particle at $|b\rangle$. If the particle is measured to be at $|b\rangle$, then we localize the particle at $|b\rangle$ and stop the quantum walk. Otherwise, we project the particle's wavefunction onto the subspace orthogonal to $|b\rangle$ and iterate. One way to formalize this procedure is to say that if $|\psi\rangle$ is the state after the application of U, we apply

$$|\psi\rangle \mapsto \begin{cases} |b\rangle & \text{with prob. } \langle b|\psi\rangle \\ \dfrac{1 - \langle b|\psi\rangle|b\rangle}{\sqrt{1 - |\langle b|\psi\rangle|^2}} & \text{with prob. } 1 - \langle b|\psi\rangle \end{cases}.$$

Although this is no longer a unitary evolution, it is a well-defined process. We can then ask what the probability is that the particle will be measured to be at $|b\rangle$ after a finite number of time steps.

The analysis for the quantum walk is much more involved, but using generating functions and an eigenvalue technique, the probability of absorption in finite time has been calculated in [10] and [11] to be $p^* = 2/\pi \approx 0.637$ for a Hadamard coin operated walk on the line starting in the state $|\uparrow\rangle|1\rangle$ with an absorbing boundary at $|b\rangle = |0\rangle$. Thus, whereas the classical particle was doomed, the quantum particle has an escape probability of 0.363!

Finally, just as we can consider extensions of classical walks to a graph, we can do the same for a quantum walk. The most straightforward extension is to the d-regular graph, where every vertex has d neighbors. In this case, we have to change $\mathcal{H}_c$ into a Hilbert space of a d-level system (a "quantum die") and choose C to be a unitary operator on $\mathcal{H}_c$. Correspondingly, S should be modified to move to the jth neighbor when $|\chi\rangle$ is $|j\rangle\langle j|$.

The extension of the Hadamard operator to d dimensions is the discrete Fourier transform, effected by

$$D = \frac{1}{\sqrt{d}} \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \cdots & \omega^{d-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{(d-1)} & \omega^{2(d-1)} & \cdots & \omega^{(d-1)(d-1)} \end{pmatrix},$$

where $\omega = e^{2\pi i/d}$. This sends the jth face of the die into a superposition of all d faces, by assigning to each face a dth root of unity. This allows the particle to move into a superposition of all its nearest neighbors in one step.

IV. CONTINUOUS TIME QUANTUM WALKS

The continuous time quantum random walk is a different formulation of quantum walks first given by Farhi, et al. in [4] and takes its inspiration from a continuous time classical random walk on a graph $G = (V, E)$. Thus, let us first discuss the transition from a discrete to a continuous time *classical* walk on the graph, before moving onto the continuous time quantum walk.

The first step is to note that we can describe classical walks in a slightly different manner. Suppose that, instead of describing the position of the particle using a random variable, we instead use the probability p_j that the particle is at vertex v_j. Then a single step of the classical walk can be described by

$$p_j \mapsto \sum_k M_{jk}\, p_k$$

where M_{jk} gives the probability of travelling on the edge (v_j, v_k). However, we want to know if there is a way of making these transitions not in a discrete time step, but continuously in time.

To do this, we introduce the *adjacency matrix* $[A_{jk}]$ of the graph, where

$$A_{jk} = \begin{cases} 1 & \text{if } (v_j, v_k) \in E \\ 0 & \text{if } (v_j, v_k) \notin E \end{cases}.$$

Note that for a graph with equal transition probabilities, the matrix $[M_{jk}]$ is just A/d.

With this in hand, we next define the *Laplacian* of the graph to be[12]

$$L = A - D, \quad \text{where} \quad D_{jk} = \deg(v_j)\delta_{jk}$$

is the diagonal matrix of the degrees of each vertex. Note that for an undirected graph, L is a symmetric matrix.

Suppose now that instead of taking the transitions to be discrete steps, the particle has a *probability rate* γ of making a jump to neighboring vertices. The resulting walk is a continuous verison of the discrete walk with equal transition probabilities to all neighbors, and it is in fact generated by the Laplacian: if the probability vector is $\mathbf{p} = (p_1, p_2, \ldots, p_{|V|})$, then

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = \gamma L\mathbf{p} \quad \Rightarrow \quad \mathbf{p}(t) = \exp(\gamma L t)\,\mathbf{p}(0). \tag{4}$$

To justify Equation 4, note that the probability of a jump in a small time interval Δt is $\gamma\Delta t$, and $(A\mathbf{p})_j$ gives a measure of the probability that the particle hops onto the vertex v_j from its neighbors. But we also have to take into account the fact that with probability $1 - \gamma\Delta t$, the particle might not transition to v_j, even though it is at a neighboring vertex; hence we subtract $(D\mathbf{p})_j$, resulting in our definition of L. Taking the limit of $\Delta t \to 0$ gives Equation 4.

Already, some similarities to the formalism of quantum mechanics can be seen. For example, we can understand the name given to the Laplacian matrix in the following way. If the classical particle has probability one of being at vertex v_j, then abusing notation for a moment and writing $|j\rangle$ for the particle's state, we see that

$$L|j\rangle = |j-1\rangle - 2|j\rangle + |j+1\rangle$$

which (up to units) is a discrete version of the continuous Laplacian $\partial^2/\partial x^2$.[13] Furthermore, Equation 4 bears resemblance to the Schrödinger equation

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

and its solution (setting $\hbar = 1$ for convenience). We can even think of the probability vector components p_j as corresponding to the quantum probabilities $|\langle j|\psi\rangle|^2$.

Based on these and other considerations, Farhi et al. chose to quantize the continuous time classical walk by choosing the Hamiltonian $H = -\gamma L$, so that unitary time evolution is given by $U(t) = \exp(-i\gamma L t)$ with the familiar resulting dynamics

$$|\psi(t)\rangle = U(t)|\psi(0)\rangle. \tag{5}$$

Note that unlike the correspondence between the discrete and continuous versions of the classical random walk, which is treated by the theory of Markov processes, it is nontrivial to understand the relationships between the discrete and continuous time quantum walks. Because of the use of an auxiliary coin state in the discrete time quantum walk, it is not clear how to translate from one to the other. After several years, Childs in [14] finally details the correspondence, although the analysis is beyond the scope of this paper.

For completeness, we show in Figure 5 a plot of the probability $|\langle 10|\psi(t)\rangle|^2$ for a continuous time random walk on the line with initial state $|\psi(0)\rangle = |0\rangle$, according to Equation 5. Note that the probability becomes appreciable starting at $t \approx 10$ (i.e., linear in z), and then steadily decreases as the particle propagates past $z = 10$.

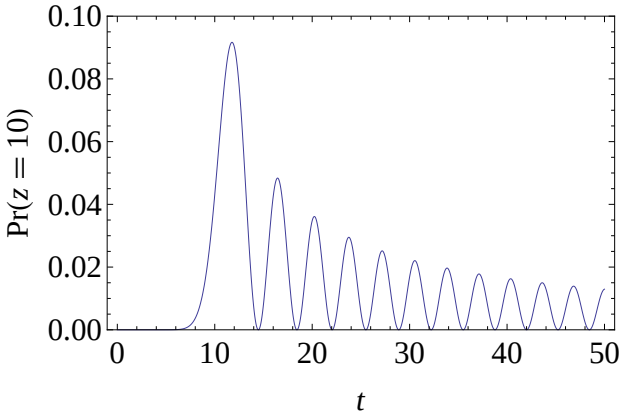

FIG. 5. A plot of the probability $|\langle 10|\psi(t)\rangle|^2$ for a continuous time random walk on the line with initial state $|0\rangle$, approximated by simulating only the region $z \in [-100, 100]$.

We now see that we in fact require G to be undirected or else L would not be symmetric under transpose and consequently H not Hermitian. However, as noted in [13], this formalism *can* handle irregular graphs, in contrast to the discrete case, for which a suitable coin system would be difficult to find.

Another advantage of the continuous quantum walk, and perhaps the more important one, is the fact that its formulation was presented alongside not just properties of the walk, but with an eye towards quantum algorithms. A first step towards this was put forth in [4], where it was shown that for some graphs, the probability of that the particle hit certain vertices was exponentially faster than in the classical walk. A simplified example, given in [13], involves the graph G_n, an example of which is shown in Figure 6 below.

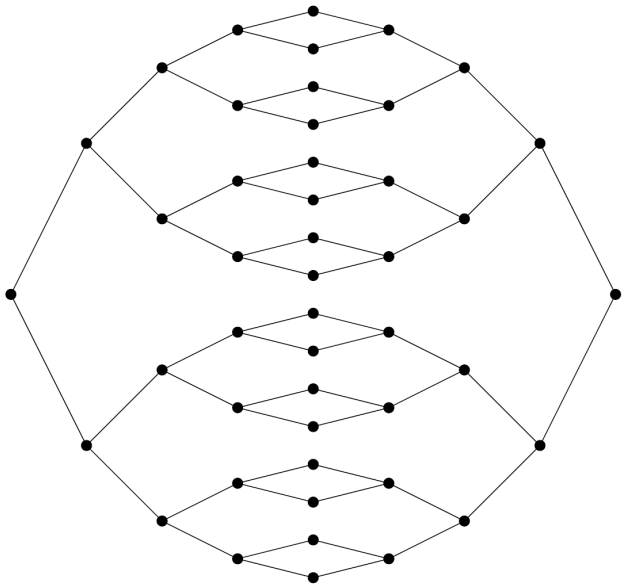

FIG. 6. The graph G_4, consisting of two binary trees of 4 levels glued together at their leaves. Adapted from [13].

Consider a classical particle, initially at the left-most vertex, taking a random walk on G_4. Clearly, the particle will propagate rather quickly towards the center while it is in the left half of the tree, since there are more edges taking it towards the right than towards the left. Once the particle reaches the center, however, it will tend to stay there, and in general, it can be shown that it will take time exponential in n for the particle to have any appreciable probability of reaching the right vertex.

On the other hand, [13] shows, through a detailed analysis and several numerical simulations, that the time it takes for the continuous time quantum walk is *linear*. That is, for the graph G_n, there will be an appreciable probability $|\langle 2n+1|\psi(t)\rangle|^2$ in time $t = O(n)$, where $|2n+1\rangle$ denotes the right-most vertex.

Nevertheless, this was not quite a success yet, because there exist classical algorithms—not based on random walks—which could still make it to the right-most vertex in quadratic time. However, it was then shown in [5] that, after modifying the question a bit and then adding many cycles to the middle of G_4, we can arrive at a situation where there there does not exist a classical algorithm to solve the problem in less than exponential time. This therefore signalled the birth of a new (albeit rather artificial) quantum algorithm.

V. CONCLUSIONS AND OUTLOOK

We have reviewed only the most basic properties of quantum random walks and have mostly focused on the formalism for defining them. More advanced research is typically done on properties such as the time to hit a certain vertex or the time it takes for a particle to revisit a vertex. These are all quantities of interest in quantifying the efficiency of a quantum walk.

Generally speaking, the quantum walks on the line and simple graphs like cycles and hypercubes are well understood. Less well understood are general higher-degree graphs, including those which are not regular. We have focused on undirected graphs in this paper for a reason: there is still no good formalism for handling directed graphs, since, for example, continuous time evolution on a directed graph is not as simple as just taking the Laplacian of the graph to be the Hamiltonian.

Finally, one of the most important topics in current research is that of decoherence, introduced in [15], and much of which is summarized in [16]. There is apparently very rich behavior to be found when we introduce decoherence into a walk, so that we get something in between the perfect quantum walk with strong interference, and the noninterfering classical random walk. For example, Lloyd et al. have published a recent result showing how the concepts of quantum walks apply to the photosynthetic efficiency of certain organisms, and in a way which depends crucially on the amount of decoherence there is in the system (the so-called "Goldilocks effect"—not too much and not too little).[17] Nevertheless, there are no general techniques for understanding decoherence, and many new results are still being developed today.

[1] J. McGreevy, "Probability, birthday problem, 1d random walk," (2011), 8.044 Recitation Notes.

[2] R.P. Feynman, "Quantum Mechanical Computers," Optics News (1985).

[3] Y. Aharonov, L. Davidovich, and N. Zagury, "Quantum random walks," Phys. Rev. A **48** (1993).

[4] E. Farhi and S. Gutmann, "Quantum computation and decision trees," Phys. Rev. A **58** (1998).

[5] A.M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D.A. Spielman, "Exponential algorithmic speedup by quantum walk," in *Proceedings of the thirty-fifth annual ACM Symposium on Theory of Computing* (STOC, 2003) pp. 58–68.

[6] A.M. Childs, "Universal Computation by Quantum Walk," Phys. Rev. Lett. **102** (2009).

[7] N.B. Lovett, S. Cooper, M. Everitt, M. Trevers, and V. Kendon, "Universal quantum computation using the discrete-time quantum walk," Phys. Rev. A **81** (2010).

[8] J. Kempe, "Quantum random walks: An introductory overview," Contemp. Phys. **50** (2009).

[9] S.E. Venagas-Andraca, "Quantum walks: a comprehensive review," (2012), (unpublished), arXiv:1201.4780v1.

[10] A. Ambainis, E. Bach, A. Nayak, A. Vishwanath, and J. Watrous, "One-dimensional quantum walks," in *Proceedings of the thirty-third annual ACM Symposium on Theory of Computing* (STOC, 2001) pp. 37–49.

[11] E. Bach, S. Coppersmith, M.P. Goldschen, R. Joynt, and J. Watrous, "One-dimensional quantum walks with absorbing boundaries," J. Comput. System Sci. **69**, 562–592 (2004).

[12] A.M. Childs and J. Goldstone, "Spatial search by quantum walk," Phys. Rev. A **70** (2004).

[13] A.M. Childs, E. Farhi, and S. Gutmann, "An example of the difference between quantum and classical random walks," Quantum Inf. Process. **1** (2002).

[14] A.M. Childs, "On the relationship between continuous- and discrete-time quantum walk," Commun. Math. Phys. **294**, 581–603 (2009).

[15] V. Kendon and B. Treganna, "Decoherence can be useful in quantum walks," Phys. Rev. A **67** (2003).

[16] V. Kendon, "Decoherence in quantum walks—a review," Math. Struct. in Comp. Science **17** (2007).

[17] S. Lloyd, M. Mohseni, A. Shabani, and H. Rabitz, "The quantum Goldilocks effect: on the convergence of timescales in quantum transport," (2011), Paper delivered at Royal Society Conference on Quantum Biology (June 2011), arXiv:1111.4982v1.