Quantum Walk Algorithms

Follow a quantum walk's coin and position amplitudes and see how interference supports graph algorithms.

Undergraduate 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Distinguish coined and continuous-time walk models.
- Calculate a three-step Hadamard walk without discarding its coin.
- State graph-access and marked-vertex costs in a walk algorithm.

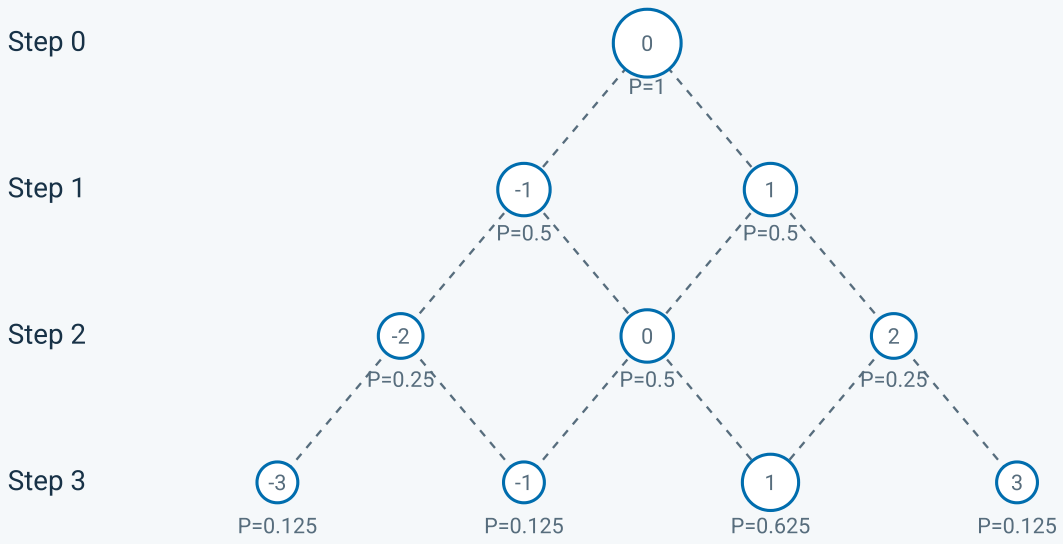
Replace random motion with coherent motion

A classical random walk moves through a graph according to transition probabilities. A quantum walk evolves complex amplitudes associated with positions, edges or additional coin states. Different paths can interfere when they reach the same final quantum configuration. This changes propagation and can support algorithms for graph search, collision problems and other structured tasks.

Review amplitude interference and unitary gates. Do not imagine a classical walker secretly trying all routes and then reading them all. The quantum state evolves coherently, and measurement still returns a limited record. The algorithm must make a useful property influence that record.

WORKED DIAGRAM · A10

A three-step walk keeps the signs



Apply H to the coin, then shift: coin 0 moves right and coin 1 left. Two coin-zero contributions reinforce at position +1; two coin-one contributions cancel at −1. Final position probabilities are 1/8, 1/8, 5/8, 1/8.

Two common models

A discrete-time coined walk uses a position register and an internal coin register. A unitary coin operation mixes directions, then a conditional shift moves the position. The full shift must be reversible. A continuous-time walk instead evolves under a Hamiltonian related to a graph's adjacency or Laplacian structure. These are different models with different implementation assumptions.

A walk is not yet a search algorithm. A search construction also specifies marked vertices, preparation, how marking changes the evolution, when to stop and what measurement counts as success. Different graph families can have very different behaviour. There is no theorem that an arbitrary quantum walk is always faster than every classical walk for every task.

Setup for a simple line walk

Let the position be an integer on a line, with a two-state coin. Coin zero moves one step right and coin one moves one step left. Apply H to the coin, then shift conditionally. Start at position zero with coin zero. We will calculate the distribution after three steps without measuring between them.

This is a learning instance, not a large-graph speedup claim. A finite implementation needs enough position labels to avoid unwanted wraparound during those three steps, or must explicitly define a cycle. The measurement asks only for final position; coin labels are summed as probabilities when they remain distinguishable.

Keep the full position-and-coin label. One walk step applies the coin first and then the conditional shift.

$$W = S(I \otimes H), \qquad S|x,0\rangle = |x+1,0\rangle, \quad S|x,1\rangle = |x-1,1\rangle.$$

x is an integer position; the second label is the coin bit. The rightmost operation acts first. The initial state is $|0,0\rangle$.

Work the first two steps

After one step, the state has amplitude one over square root of two on coin zero at position plus one and the same amplitude on coin one at position minus one. Position probabilities are therefore equal on the two neighbours.

After the second H-and-shift step, the nonzero configurations are coin zero at plus two, coin one at zero, coin zero at zero and coin one at minus two. Their amplitudes are one half, one half, one half and negative one half respectively. The position-zero probability is one half, obtained by adding the squared magnitudes of its two orthogonal coin components.

The minus sign has not yet changed the position distribution relative to a classical two-step fair walk. It is nevertheless essential for predicting the next step. A histogram at time two is an incomplete state description, just as equal qubit probabilities can hide a relative phase.

The first two exact states reveal the sign that will matter at the next step.

$$W|0,0\rangle = \frac{|1,0\rangle + |-1,1\rangle}{\sqrt{2}},$$
$$W^2|0,0\rangle = \frac{|2,0\rangle + |0,1\rangle + |0,0\rangle - |-2,1\rangle}{2}.$$

The two entries at position zero have different coin labels, so they are orthogonal configurations.

Checkpoint 1

PAUSE & THINK

At step two, position zero has amplitudes 1/2 in each coin state. What is P(0)?

- ☐ 1
- ☐ 0
- ☐ 1/2

Answer: 1/2. The coin labels remain distinguishable, so their probabilities add to 1/4+1/4=1/2.

Interference at the third step

Apply H and shift once more. Two contributions to coin zero at position plus one add, while two contributions to coin one at position minus one cancel. The final position probabilities are one eighth at plus three, five eighths at plus one, one eighth at minus one and one eighth at minus three.

A classical fair three-step walk instead has one eighth at either outer position and three eighths at either inner position. The difference comes from coherent addition and cancellation, not from a different classical probability assigned to each path. The asymmetric quantum result also depends on the initial coin state. A suitable symmetric complex coin preparation changes the symmetry of the distribution.

This example teaches two checks. Keep the coin in the state description until readout. Also preserve signed amplitudes through every step. Replacing the state by its position histogram after each step effectively inserts a loss of coherence and defines a different process.

At step three, two paths meet at the same full configuration and their amplitudes combine.

$$W^3|0,0\rangle = \frac{1}{2\sqrt{2}}\left(|3,0\rangle + |1,1\rangle + 2|1,0\rangle - |-1,0\rangle + |-3,1\rangle\right).$$

The amplitude on $|-1,1\rangle$ is zero through cancellation. The amplitude on $|1,0\rangle$ doubles through reinforcement.

For position-only readout, add the probabilities of its distinct coin states.

$$P(x) = |\alpha_{x,0}|^2 + |\alpha_{x,1}|^2, \qquad \big(P(-3),P(-1),P(1),P(3)\big) = \frac{1}{8}(1,1,5,1).$$

At $x = 1$, the coin-zero contribution is $1/2$ and the coin-one contribution is $1/8$, giving $5/8$.

How walks become algorithms

A graph algorithm can arrange for marked vertices to alter a reflection, phase or evolution generator. Interference then changes the probability of detecting a marked region. The useful runtime depends on the graph's spectral structure, marked fraction and costs of preparing, updating and checking the walk state.

Quantum-walk methods are especially valuable when local updates preserve useful information between checks. For example, a walk on subsets can update a small part of a stored candidate rather than rebuild it completely. Ambainis's element-distinctness algorithm uses this idea to obtain a query improvement for finding equal items. Its performance is a theorem about a specified oracle problem, not a claim that every database task shares that bound.

Costs and limitations

Graph access must be stated. Can you query neighbours coherently? Are vertex degrees bounded or supplied through an oracle? How expensive is a shift, state preparation or marked-vertex test? A mathematical adjacency matrix is not automatically a cheap hardware operation. Continuous-time algorithms similarly need a method for simulating the relevant Hamiltonian.

Classical baselines must match the task. Compare finding a marked vertex with finding a marked vertex, not with printing every graph path. Include error probability and output verification. Some graphs exhibit useful propagation advantages; others introduce localisation, small spectral gaps or costly preparation. The mechanism suggests possibilities, but the graph-specific proof establishes the algorithmic guarantee.

Reverse the walk as a coherence check

The conditional shift is a permutation: its inverse moves each coin label in the opposite direction. H is its own inverse. To undo one complete step, first undo the shift and then apply H, reversing the original order. Repeating that inverse sequence for all three steps returns the full ideal state to position zero and coin zero. Replacing intermediate amplitudes by a measured position distribution generally prevents that deterministic reversal. This provides a useful diagnostic distinction between coherent walk evolution and a stochastic simulation that merely updates probabilities after every step. An experiment can compare the inverse prediction, but hardware imperfections require their own error analysis.

Undo operations in reverse execution order to reverse the complete step.

$$W^\dagger = (I \otimes H)S^\dagger, \qquad (W^\dagger)^3 W^3|0,0\rangle = |0,0\rangle.$$

The inverse shift acts first, followed by H. This exact return tests the full coherent state, beyond its position histogram.

Checkpoint 2

PAUSE & THINK

Which execution order reverses a step that applied H then S?

- ☐ H then S†
- ☐ S† then H
- ☐ Measure then shift

Answer: S† then H. W=S(I⊗H), so W†=(I⊗H)S†. Execution runs from right to left.

Chapter summary

- A coined walk evolves amplitudes on joint position-and-coin configurations.
- Add amplitudes only when the full configurations agree; otherwise add readout probabilities.
- An algorithm needs graph access, marking, a stopping rule and a success condition in addition to a walk.

A common misconception

All amplitudes at the same position can be added before squaring. Different coin labels are orthogonal. For position-only measurement, square each coin amplitude and then add their probabilities.

Exercises

1. At time two in the worked walk, why is the probability at position zero one half rather than one or zero?
2. What changes if you measure the position after every step? Is the original three-step interference calculation still valid?

Worked answers

1. There are two coin components at position zero, each with magnitude one half. They are orthogonal coin states, so their probabilities add: one quarter plus one quarter. Their amplitudes cannot be added as if they represented the same full configuration.
2. Intermediate measurement removes coherence between different measured positions and conditions the remaining state on a record. Later evolution is therefore a different experiment. The unmeasured path amplitudes used in the three-step derivation no longer describe it.

Check your understanding

PAUSE & THINK

At step two, position zero has amplitude 1/2 in each of two orthogonal coin states. What is its probability?

- [ ] 0
- [ ] 1/2
- [ ] 1

Answer: 1/2. The coin states are distinguishable configurations, so add their squared magnitudes: 1/4+1/4=1/2. Adding the amplitudes would be wrong.

Sources

Ambainis's overview of quantum walks explains the algorithmic models. His element-distinctness paper provides a primary example of a walk-based query advantage.