Quantum algorithms

Andrew Childs
Institute for Quantum Computing
University of Waterloo

11th Canadian Summer School on Quantum Information

8–9 June 2011

Based in part on slides prepared with Pawel Wocjan (UCF)

12th Canadian Summer School on Quantum Information

11–16 June 2012

Institute for Quantum Computing
University of Waterloo

http://cssqi2012.iqc.uwaterloo.ca

Outline

IV. Unstructured search

V. Quantum walk

VI. Adversary lower bounds

Part IV

Unstructured search

Unstructured search

Quantum computers can quadratically outperform classical computers at a very basic computational task, unstructured search

There is a set X containing N items, some of which are marked

We are given a Boolean black box $f : X \to \{0, 1\}$ that indicates whether a given item is marked

The problem is to decide if any item is marked, or alternatively, to find a marked item given that one exists

Unstructured search as a model for NP

Unstructured search can be thought of as a model for solving problems in NP by brute force search

If a problem is in NP, then we can efficiently recognize a solution, so one way to find a solution is to solve unstructured search

Of course, this may not be the best way to find a solution in general, even if the problem is NP-hard: we don't know if NP-hard problems are really "unstructured"

Classical vs. quantum query complexity

It is obvious that even a randomized classical algorithm needs $\Omega(N)$ queries to decide if any item is marked

But a quantum algorithm can do much better!

Phase oracle

We assume that we have a unitary operator U satisfying

$$U|x\rangle = (-1)^{f(x)}|x\rangle = \begin{cases} |x\rangle & x \text{ is not marked} \\ -|x\rangle & x \text{ is marked} \end{cases}$$

This can be created using one query to a standard reversible oracle via phase kickback

Target state

We consider the case where there is exactly one $x \in X$ element that is marked; call this element m

Our goal is to prepare the state $|m\rangle$

Initial state

We have no information about which item might be marked

Thus we take

$$|\psi\rangle := \frac{1}{\sqrt{N}} \sum_{x=1}^{N} |x\rangle$$

as the initial state

Rough idea behind Grover search

Start with the initial state $|\psi\rangle$

Implement a rotation that moves $|\psi\rangle$ toward $|m\rangle$

Realize the rotation with the help of two reflections

Visualization of a reflection in $\mathbb{R}^2$

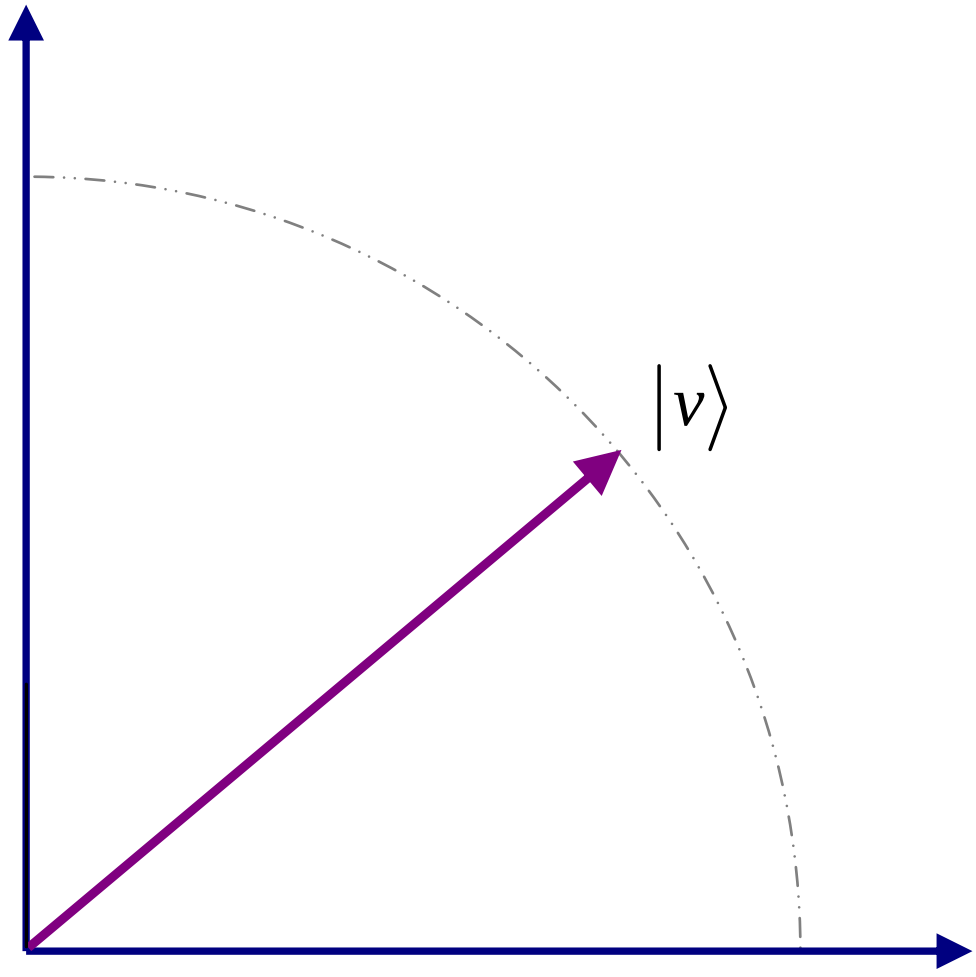

Visualization of a reflection in $\mathbb{R}^2$

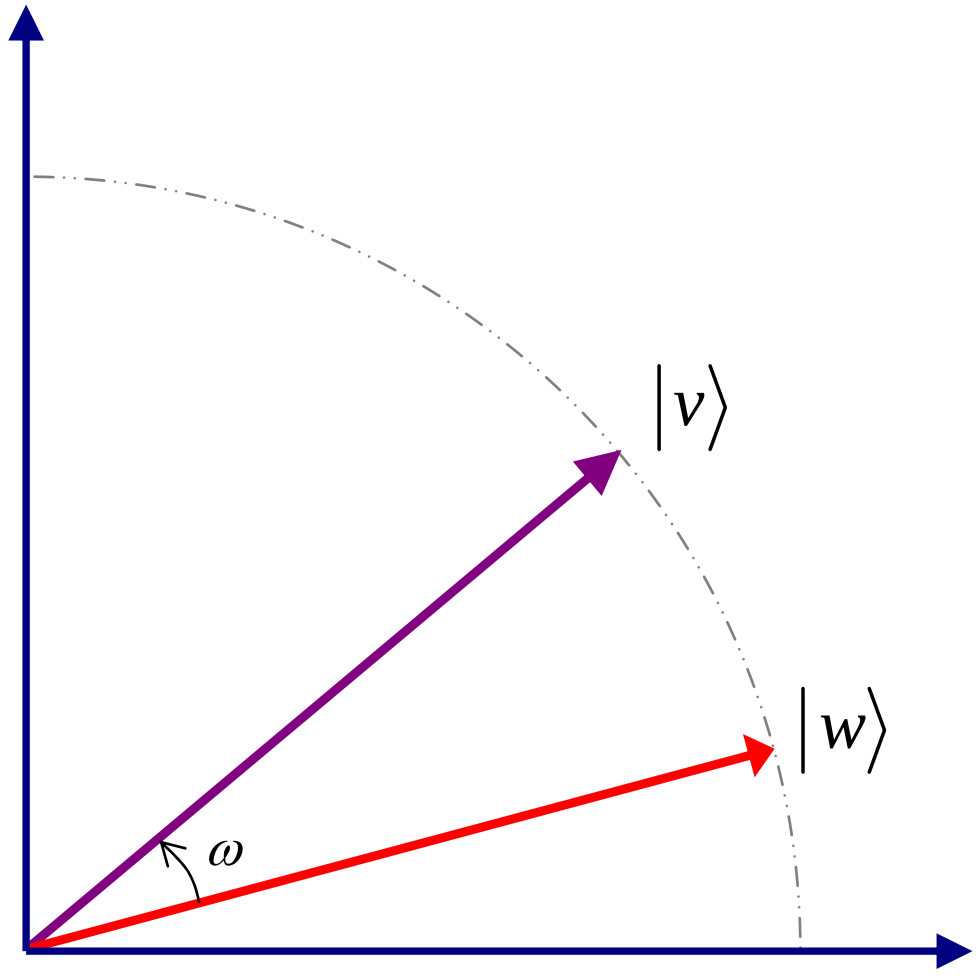

Visualization of a reflection in $\mathbb{R}^2$

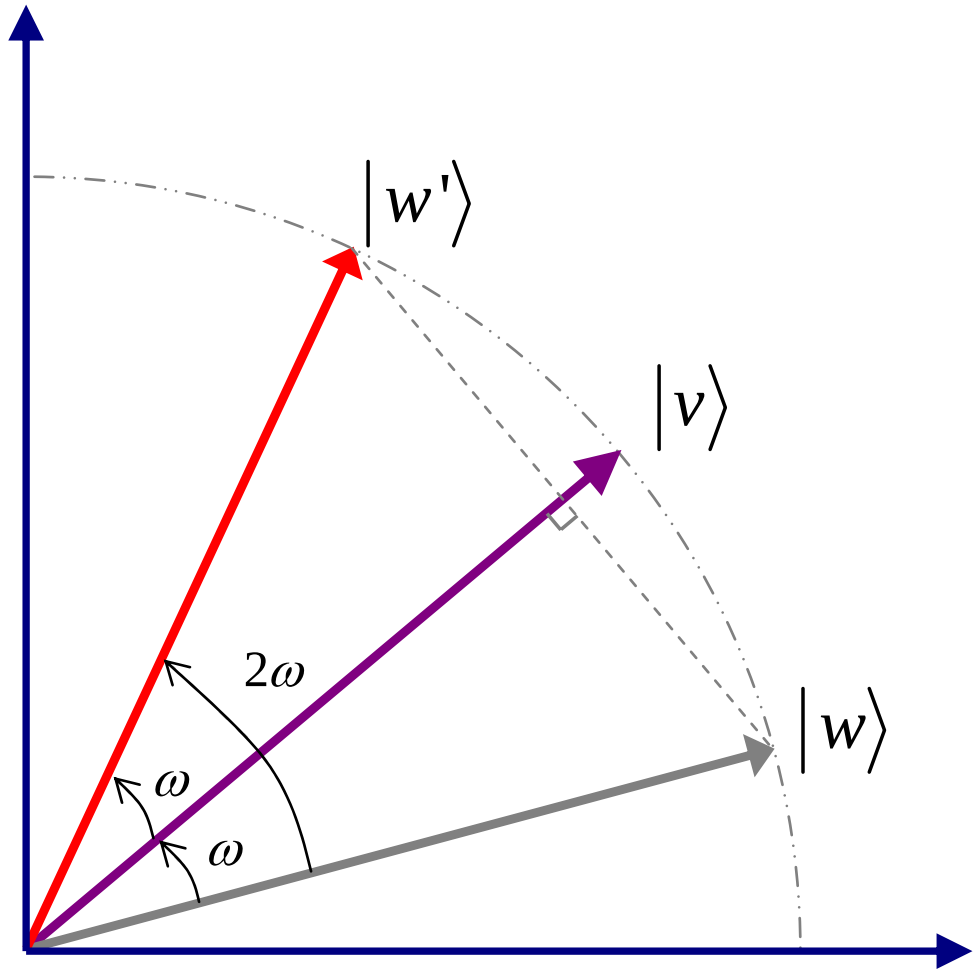

Reflections

$U = I - 2|m\rangle\langle m|$ is the reflection about the target state $|m\rangle$

$V := I - 2|\psi\rangle\langle\psi|$ is the reflection about the initial state $|\psi\rangle$:

$$\begin{aligned} V|\psi\rangle &= -|\psi\rangle \\ V|\psi^{\perp}\rangle &= |\psi^{\perp}\rangle \end{aligned}$$

for any state $|\psi^{\perp}\rangle$ orthogonal to $|\psi\rangle$

Structure of Grover's algorithm

The algorithm is as follows:

- start in $|\psi\rangle$,
- apply the Grover iteration $G := V\,U$ some number of times,
- make a measurement and hope that the outcome is m

Invariant subspace

Observe that span$\{|m\rangle, |\psi\rangle\}$ is a U- and V-invariant subspace, and both the inital and target states belong to this subspace

$\Rightarrow$ It suffices to understand the restriction of VU to this subspace

Let $\{|m\rangle, |\phi\rangle\}$ be an orthonormal basis for span$\{|m\rangle, |\psi\rangle\}$

The Gram-Schmidt process yields

$$|\phi\rangle = \frac{|\psi\rangle - \sin\theta |m\rangle}{\cos\theta}$$

where $\sin\theta := \langle m|\psi\rangle = 1/\sqrt{N}$

Invariant subspace

Now in the basis $\{|m\rangle, |\phi\rangle\}$, we have

$$|\psi\rangle = \sin\theta|m\rangle + \cos\theta|\phi\rangle \text{ where } \sin\theta = \langle m|\psi\rangle = 1/\sqrt{N}$$

$$U = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$$

$$\begin{aligned} V &= I - 2|\psi\rangle\langle\psi| \\ &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - 2\begin{pmatrix} \sin\theta \\ \cos\theta \end{pmatrix}\begin{pmatrix} \sin\theta & \cos\theta \end{pmatrix} \\ &= \begin{pmatrix} 1 - 2\sin^2\theta & -2\sin\theta\cos\theta \\ -2\sin\theta\cos\theta & 1 - 2\cos^2\theta \end{pmatrix} \\ &= -\begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} \end{aligned}$$

Grover iteration within the invariant subspace

$\Rightarrow$ We find

$$
\begin{aligned}
V\,U &= -\begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \\
&= -\begin{pmatrix} \cos 2\theta & \sin 2\theta \\ -\sin 2\theta & \cos 2\theta \end{pmatrix}
\end{aligned}
$$

This is a rotation up to a minus sign

Visualization of first Grover iteration

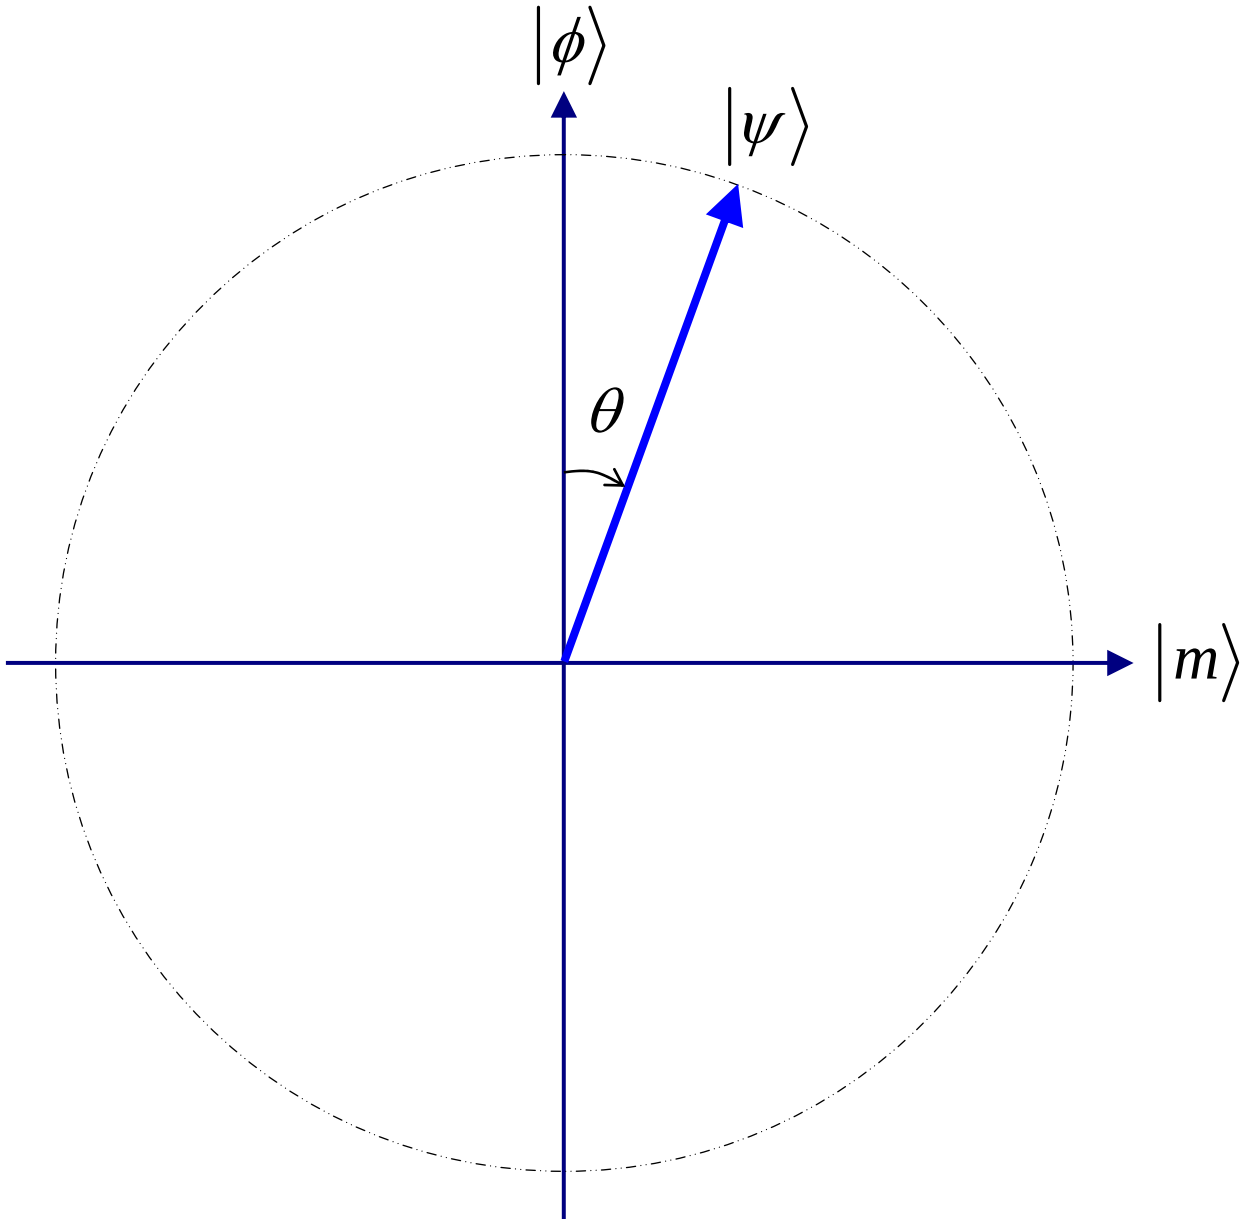

Visualization of first Grover iteration

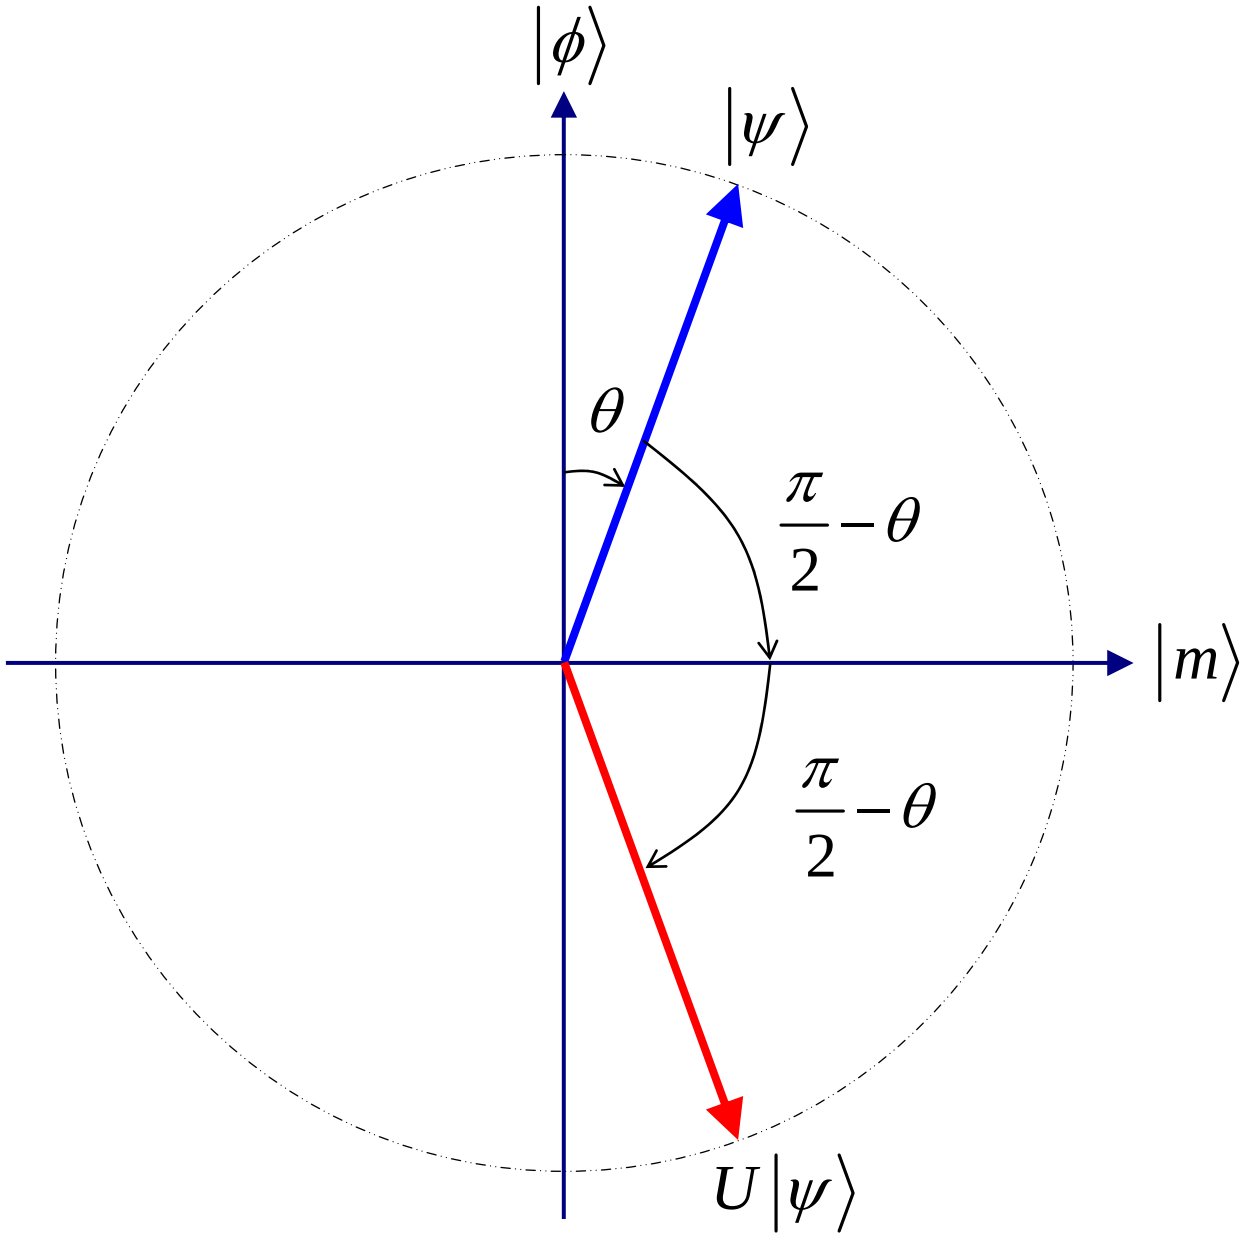

Visualization of first Grover iteration

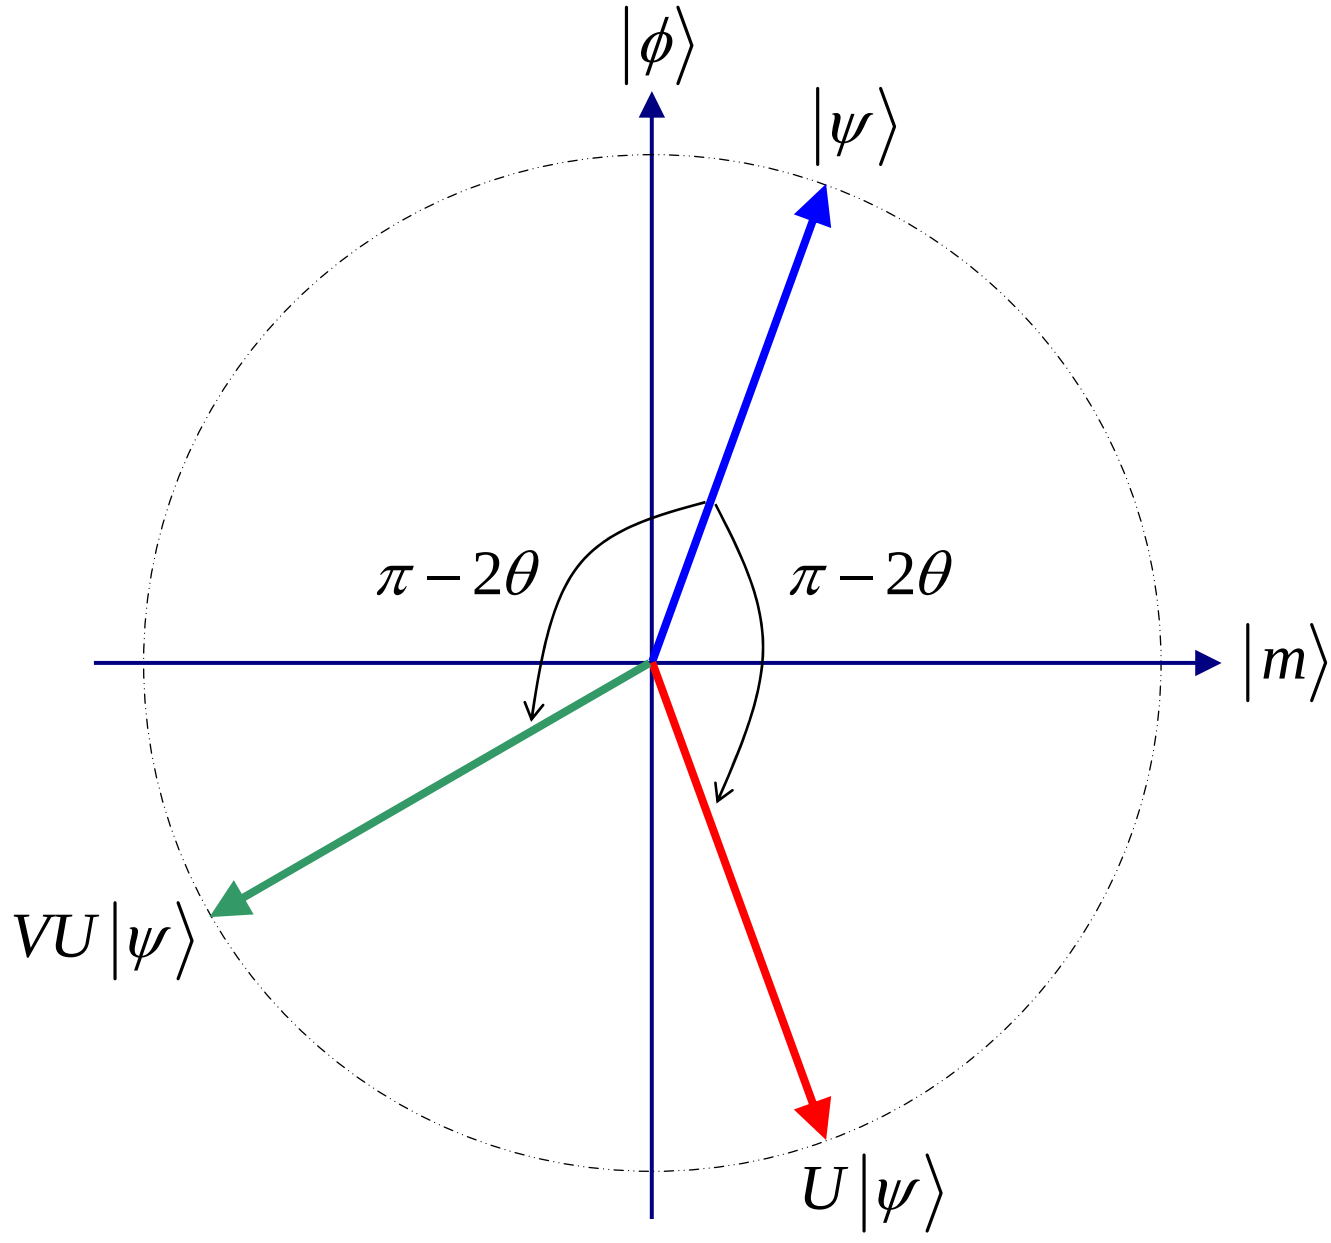

Visualization of first Grover iteration

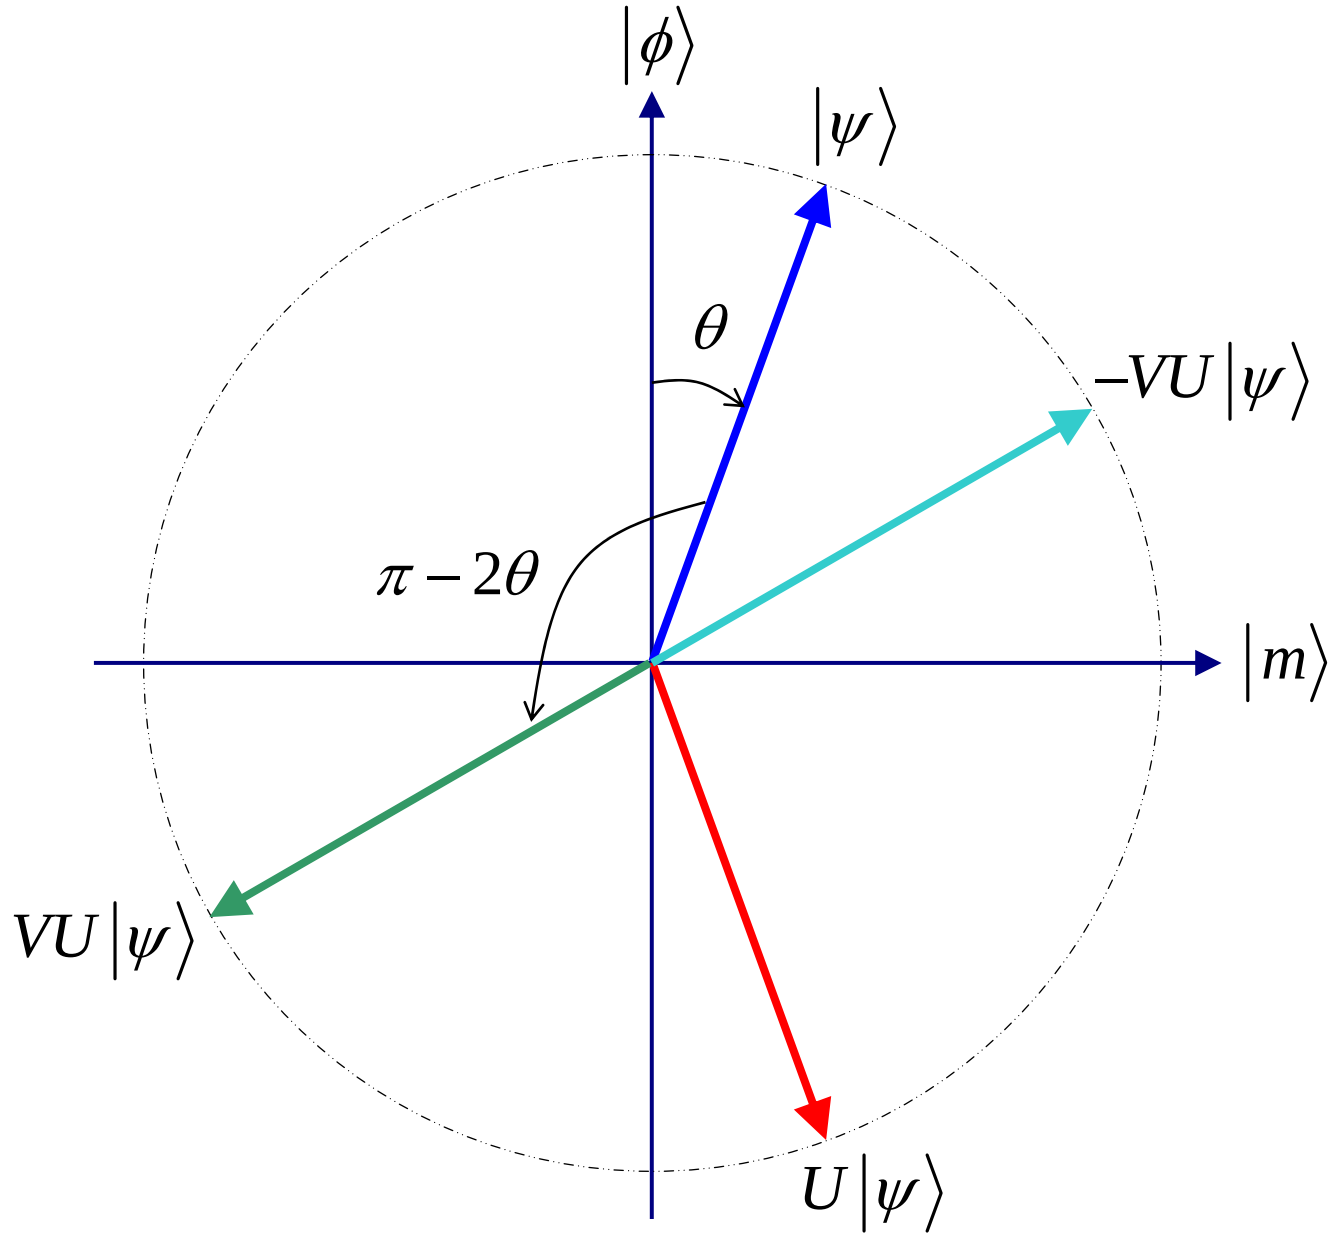

Visualization of first Grover iteration

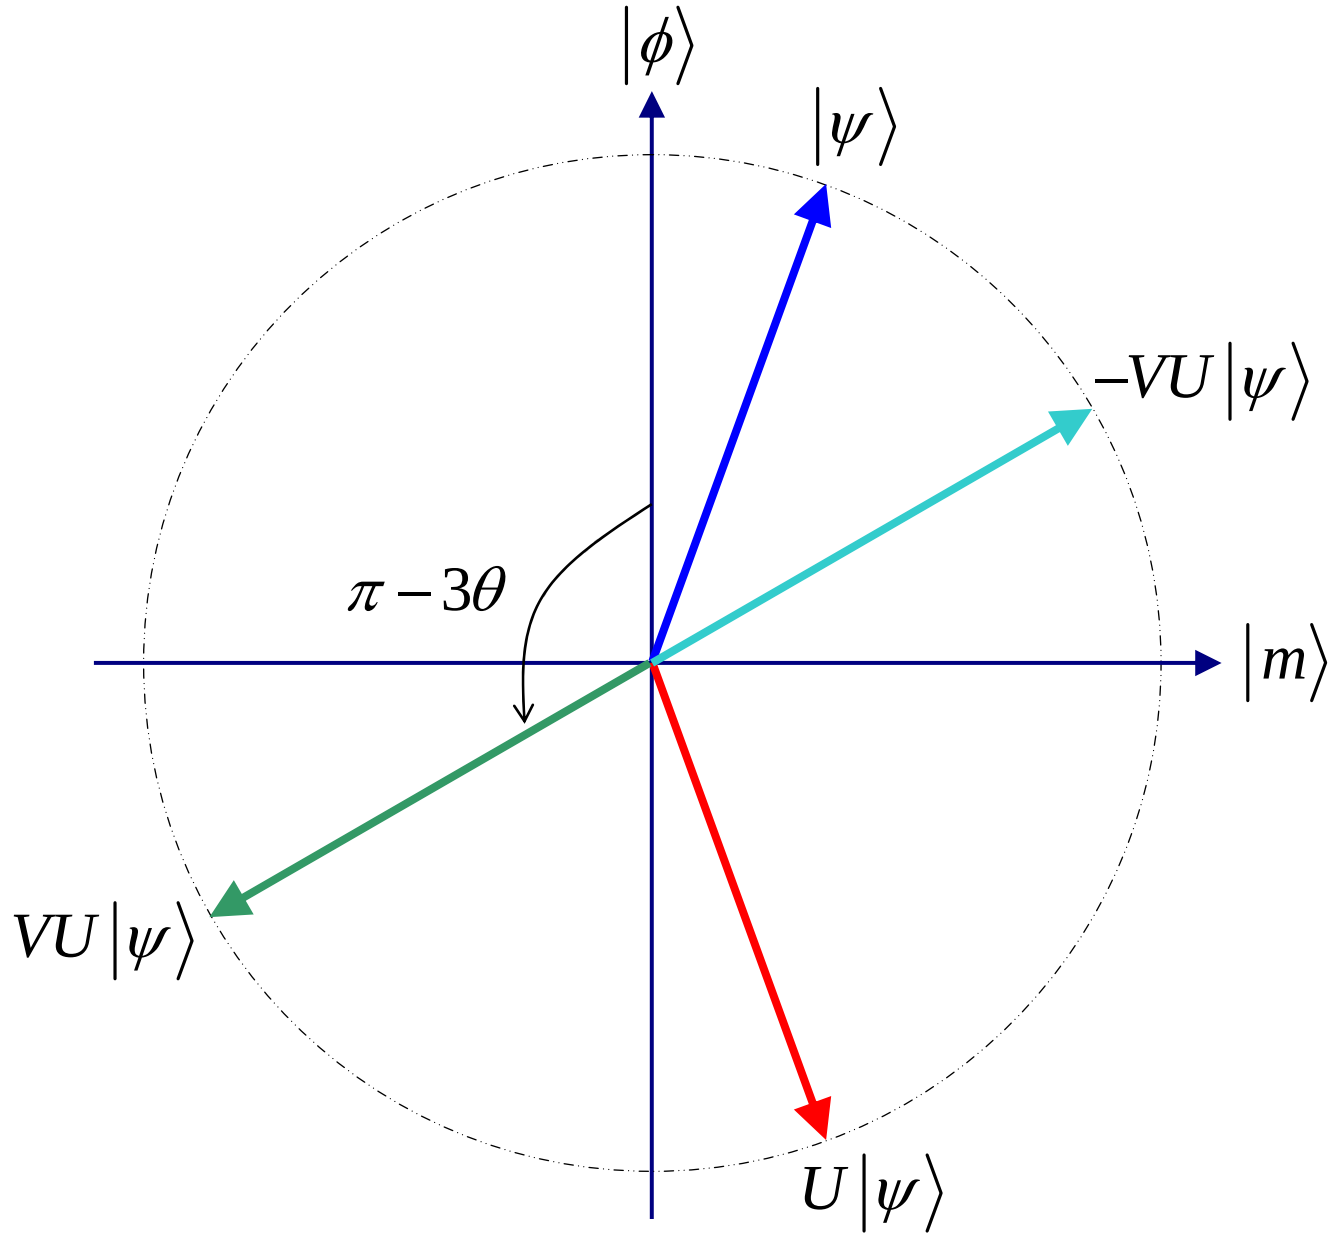

Visualization of first Grover iteration

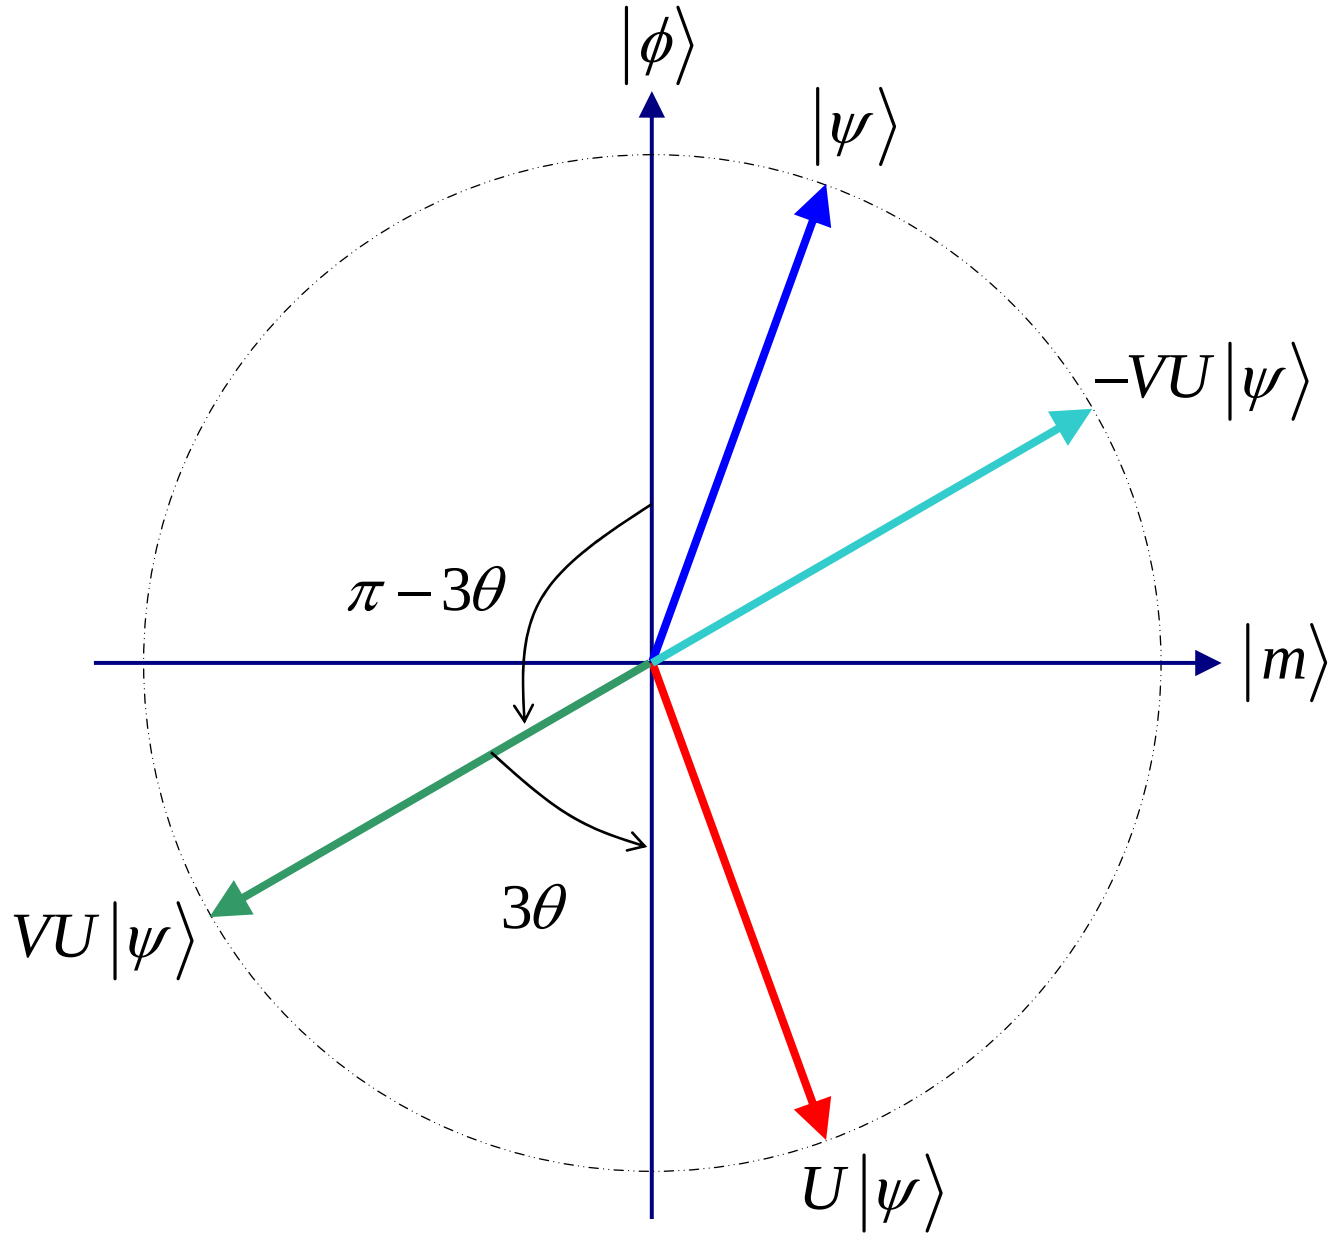

Visualization of first Grover iteration

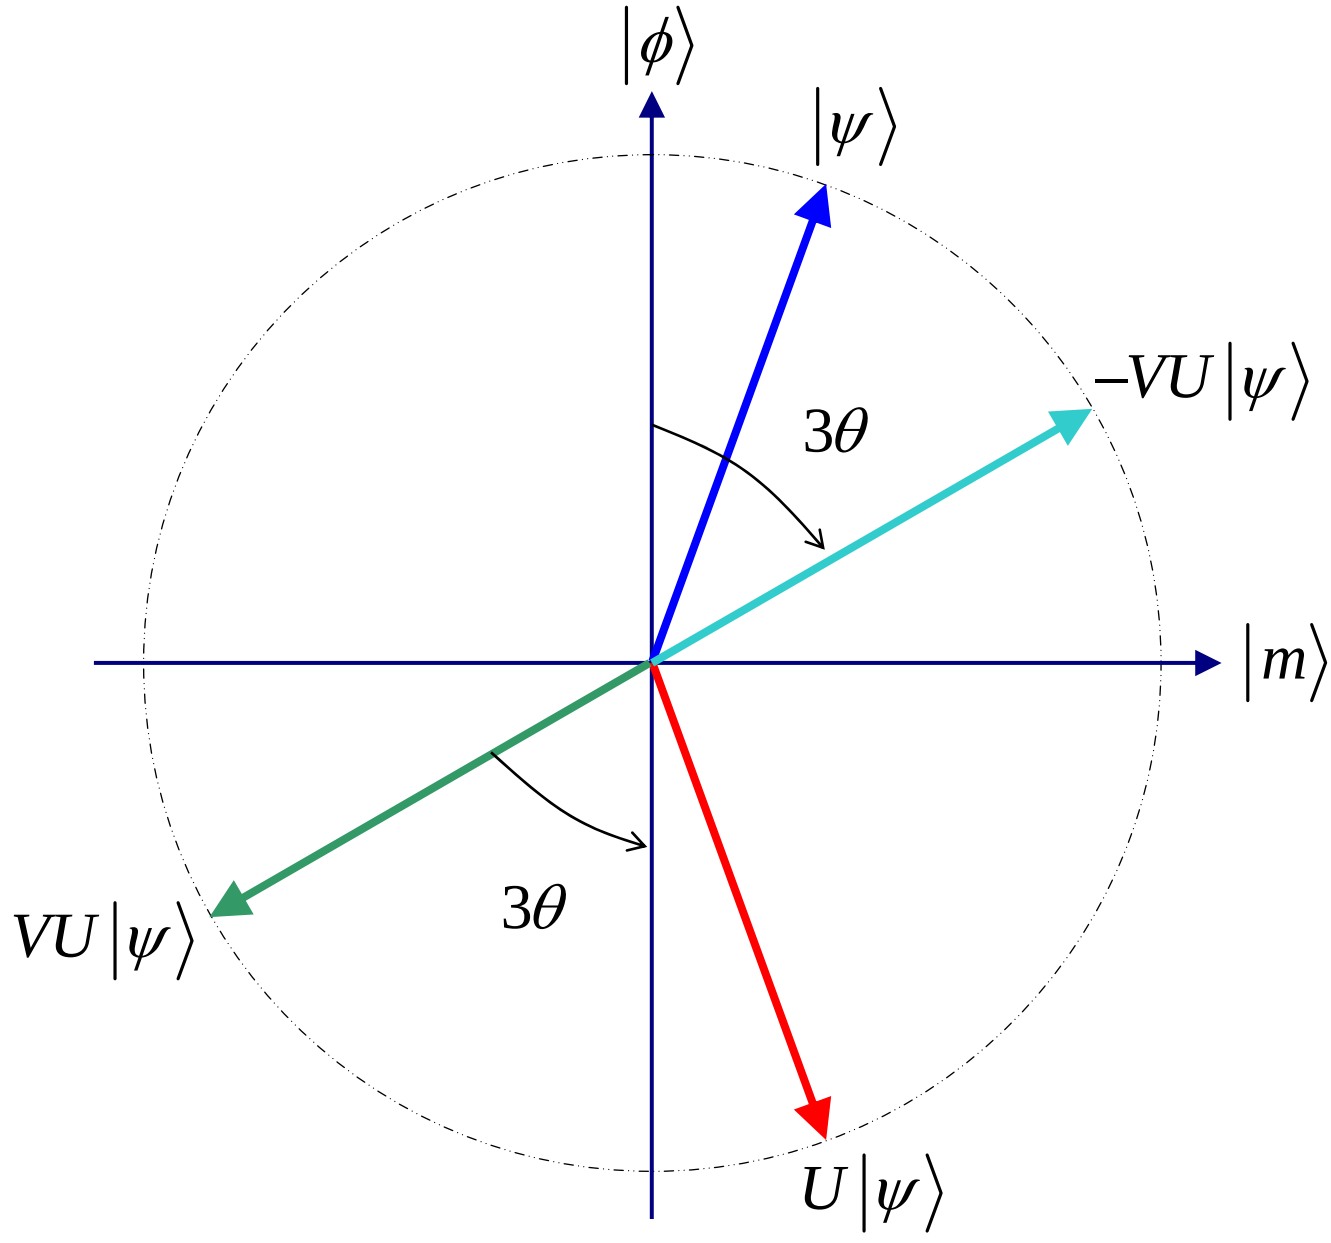

Visualization of first Grover iteration

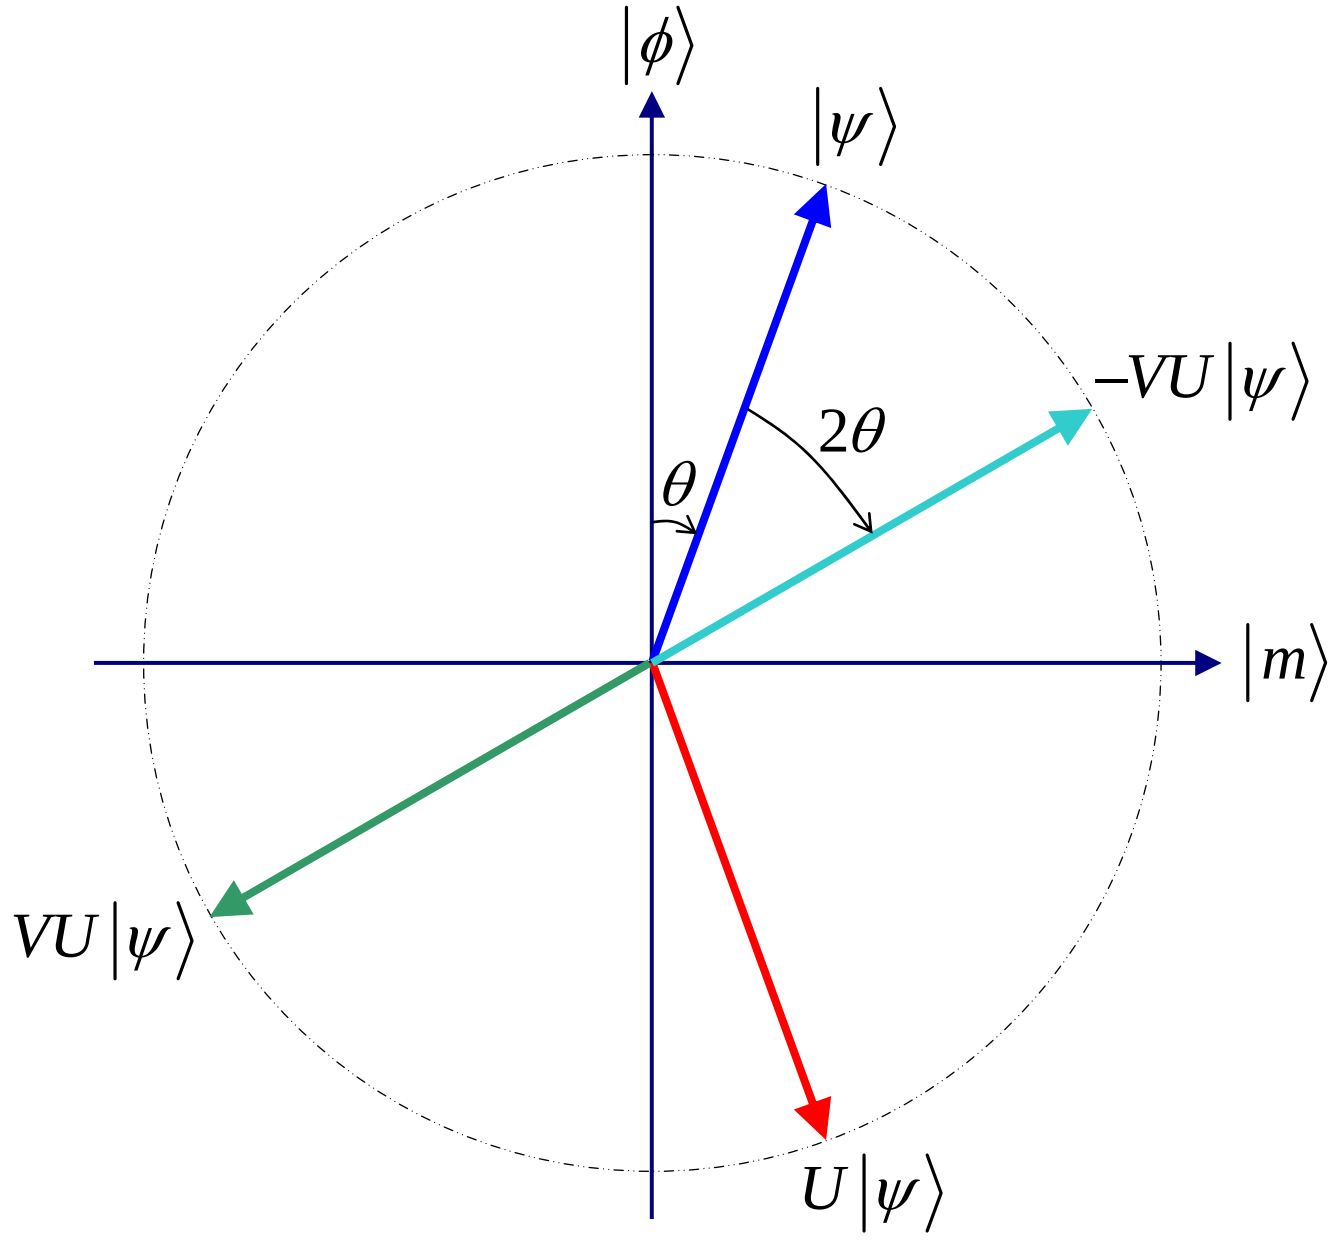

The Grover iteration is a rotation

Geometrically, U is a reflection around the $|m\rangle$ axis and V is a reflection around the $|\psi\rangle$ axis, which is almost but not quite orthogonal to the $|m\rangle$ axis

The product of these two reflections is a clockwise rotation by an angle 2θ, up to an overall minus sign

From this geometric picture, or by explicit calculation using trig identities, it is easy to verify that

$$(VU)^k = (-1)^k \begin{pmatrix} \cos 2k\theta & \sin 2k\theta \\ -\sin 2k\theta & \cos 2k\theta \end{pmatrix}$$

Complexity of Grover search

Recall that our initial state is $|\psi\rangle = \sin\theta|m\rangle + \cos\theta|\phi\rangle$

How large should k be before $(VU)^k|\psi\rangle$ is close to $|m\rangle$?

We start an angle θ from the $|\phi\rangle$ axis and rotate toward $|m\rangle$ by an angle 2θ per iteration

$$|\langle m|(VU)^k|\psi\rangle|^2 = \sin^2((2k+1)\theta)$$

$\Rightarrow$ To rotate by $\pi/2$, we need

$$\begin{aligned} \theta + 2k\theta &= \pi/2 \\ k &\approx \frac{\pi}{4}\theta^{-1} \approx \frac{\pi}{4}\sqrt{N} \end{aligned}$$

Grover search

Grover's algorithm solves a completely unstructured search problem with N possible solutions, yet finds a unique solution in only $O(\sqrt{N})$ queries!

While this is only a polynomial separation, it is very generic, and it is surprising that we can obtain a speedup for a search in which we have so little information to go on

Optimality of Grover's algorithm

It can also be shown that this quantum algorithm is optimal

Any quantum algorithm needs at least $\Omega(\sqrt{N})$ queries to find a marked item (or even to decide if some item is marked)

We will prove this in the last quantum algorithms lecture

Multiple solutions

Suppose there are M marked items

Then there is a two-dimensional invariant subspace span$\{|\mu\rangle, |\psi\rangle\}$ where

$$|\mu\rangle = \frac{1}{\sqrt{M}} \sum_{x \text{ marked}} |x\rangle$$

is the uniform superposition over all marked items

The Gram-Schmidt process yields the ONB $\{|\mu\rangle, |\phi\rangle\}$ where

$$|\phi\rangle = \frac{1}{\sqrt{N-M}} \sum_{x \text{ unmarked}} |x\rangle$$

is the uniform superposition of all non-solutions

Invariant subspace

Now in the basis $\{|\mu\rangle, |\phi\rangle\}$, we have

$$|\psi\rangle = \sin\theta|\mu\rangle + \cos\theta|\phi\rangle \text{ where } \sin\theta = \langle\mu|\psi\rangle = \sqrt{\frac{M}{N}}$$

$$VU = -\begin{pmatrix} \cos 2\theta & \sin 2\theta \\ -\sin 2\theta & \cos 2\theta \end{pmatrix}$$

Overshooting

The success probability is

$$\sin((2k+1)\theta) \text{ where } \sin\theta = \sqrt{\frac{M}{N}}$$

$\Rightarrow$ We need to apply VU

$$k \approx \frac{\pi}{4}\sqrt{\frac{N}{M}}$$

times

Due to the oscillatory behavior of the success probability, it is important not to overshoot: if the number of iterations is too large, the success probability will decrease

Quantum counting (1/2)

The eigenvalues of

$$-VU = \begin{pmatrix} \cos 2\theta & \sin 2\theta \\ -\sin 2\theta & \cos 2\theta \end{pmatrix}$$

are $e^{2i\theta}$ and $e^{-2i\theta}$

The initial state $|\psi\rangle$ is a superposition of the two eigenvectors corresponding to the above two eigenvalues

$\Rightarrow$ Using phase estimation, we can obtain an estimate $\tilde{\theta}$ such that

$$|\theta - \tilde{\theta}| \leq \epsilon$$

by invoking the controlled version of $-VU$

$$O(1/\epsilon) \text{ times}$$

Quantum counting (2/2)

The estimate $\tilde{\theta}$ of θ gives an estimate $\tilde{M}$ of M

Error:

$$\begin{aligned}\left|\frac{M}{N}-\frac{\tilde{M}}{N}\right| &= |\sin^2\theta-\sin^2\tilde{\theta}| \\ &= |\sin\theta+\sin\tilde{\theta}|\,|\sin\theta-\sin\tilde{\theta}| \\ &\approx 2\sqrt{\frac{M}{N}}\epsilon\end{aligned}$$

Equivalently, we get an approximation $\tilde{M}=M(1+O(\varepsilon))$ using $O(\frac{1}{\varepsilon}\sqrt{N/M})$ queries

Amplitude amplification

Suppose we have a classical (randomized) algorithm that produces a solution to some problem with probability p

Assume we can recognize correct solutions

Classical strategy: repeat $O(1/p)$ times

Quantum amplitude amplification uses only $O(1/\sqrt{p})$ repetitions

Exercise: Quantum search and state generation

Let $|\psi\rangle$ be an unknown quantum state. Consider quantum algorithms for preparing $|\psi\rangle$ given two different black boxes.

1. Suppose you are given the unitary $U := I - 2|\psi\rangle\langle\psi|$ as a black box. Consider a quantum algorithm that starts in some known state $|\phi\rangle$ and alternates between performing U and $V := I - 2|\phi\rangle\langle\phi|$. How many queries to U are required to prepare a state close to $|\psi\rangle$? Express your answer as a function of $|\langle\psi|\phi\rangle|$.
2. Now suppose you are given a reversible black box that, on input $x \in \{1, \ldots, N\}$, returns the amplitude $\alpha_x := \langle x|\psi\rangle$ of the state $|\psi\rangle$ in the computational basis state $|x\rangle$. (You may assume that the black box specifies the complex number α_x to arbitrary precision.) Describe an algorithm that prepares a state close to $|\psi\rangle$ using $O(\sqrt{N})$ queries. (Hint: Two queries to the black box can be used to perform the isometry $|x\rangle \mapsto |x\rangle(\alpha_x|0\rangle + \sqrt{1 - |\alpha_x|^2}|1\rangle)$.)

Part V

Quantum walk

Randomized algorithms

Randomness is an important tool in computer science

Black-box problems

- Huge speedups are possible (Deutsch-Jozsa: $2^{\Omega(n)}$ vs. $O(1)$)
- Polynomial speedup for some total functions (game trees: $\Omega(n)$ vs. $O(n^{0.754})$)

Natural problems

- Majority view is that derandomization should be possible (P=BPP)
- Randomness may give polynomial speedups (Schöning algorithm for k-SAT)
- Can be useful for algorithm design

Random walk

Graph $G = (V, E)$

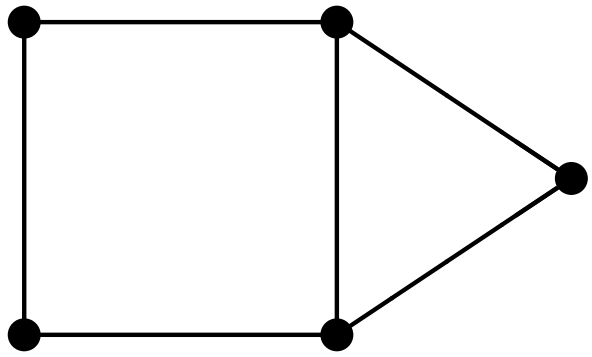

Two kinds of walks:

- Discrete time
- Continuous time

Random walk algorithms

Undirected s–t connectivity in log space

- Problem: given an undirected graph $G = (V, E)$ and $s, t \in V$, is there a path from s to t?
- A random walk from s eventually reaches t iff there is a path
- Taking a random walk only requires log space
- Can be derandomized (Reingold 2004), but this is nontrivial

Markov chain Monte Carlo

- Problem: sample from some probability distribution (uniform distribution over some set of combinatorial objects, thermal equilibrium state of a physical system, etc.)
- Create a Markov chain whose stationary distribution is the desired one
- Run the chain until it converges

Continuous-time quantum walk

Graph G

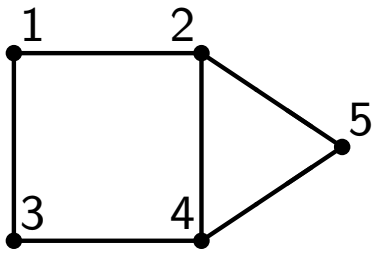

$$A = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 \end{pmatrix}$$

adjacency matrix

$$L = \begin{pmatrix} -2 & 1 & 1 & 0 & 0 \\ 1 & -3 & 0 & 1 & 1 \\ 1 & 0 & -2 & 1 & 0 \\ 0 & 1 & 1 & -3 & 1 \\ 0 & 1 & 0 & 1 & -2 \end{pmatrix}$$

Laplacian

Random walk on G

- State: probability $p_v(t)$ of being at vertex v at time t
- Dynamics: $\frac{\mathrm{d}}{\mathrm{d}t}\vec{p}(t) = -L\vec{p}(t)$

Quantum walk on G

- State: amplitude $q_v(t)$ to be at vertex v at time t (i.e., $|\psi(t)\rangle = \sum_{v \in V} q_v(t)|v\rangle$)
- Dynamics: $\mathrm{i}\frac{\mathrm{d}}{\mathrm{d}t}\vec{q}(t) = -L\vec{q}(t)$

Random vs. quantum walk on the line

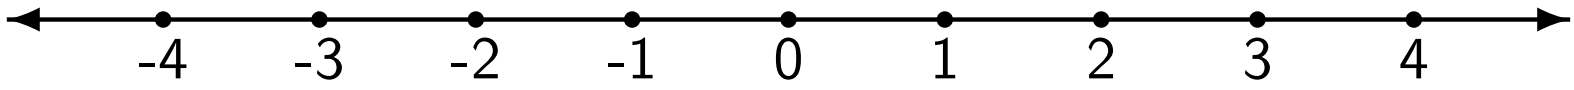

Classical:

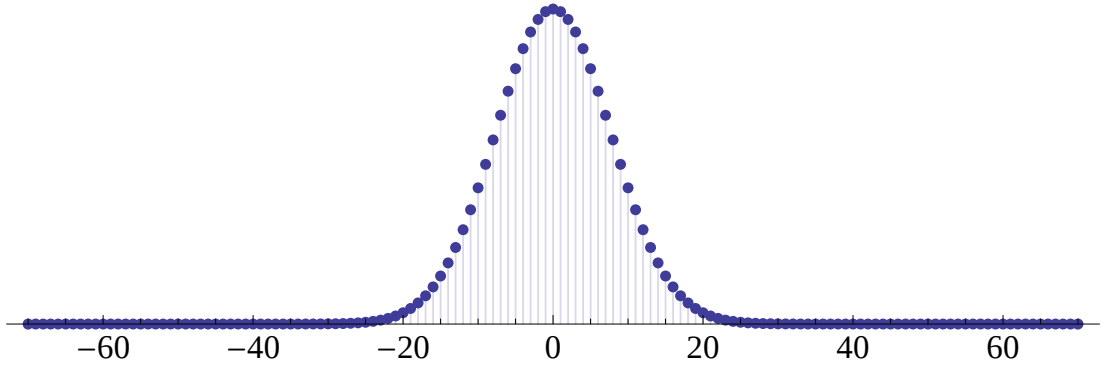

Quantum:

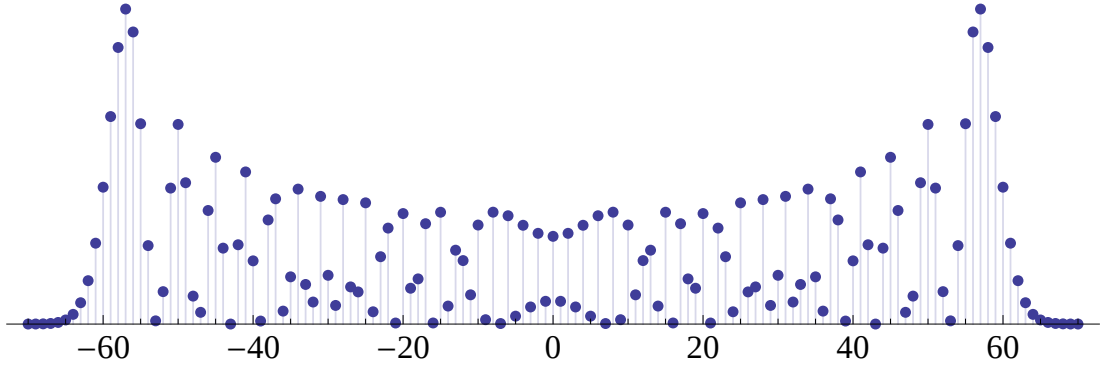

Random vs. quantum walk on the hypercube

$n = 3$:

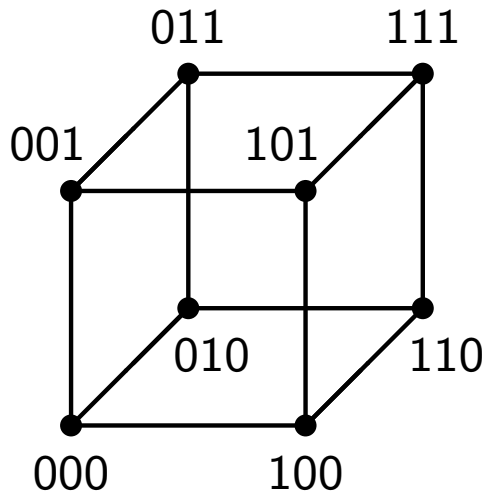

$V = \{0,1\}^n$

$E = \{(x,y) \in V \times V :$
x and y differ in
exactly one bit$\}$

Classical random walk: reaching $11\ldots1$ from $00\ldots0$ is exponentially unlikely

Quantum walk: with $A = \sum_{j=1}^{n} X_j$,

$$e^{-\mathrm{i}At} = \prod_{j=1}^{n} e^{-\mathrm{i}X_j t} = \bigotimes_{j=1}^{n} \begin{pmatrix} \cos t & -\mathrm{i}\sin t \\ -\mathrm{i}\sin t & \cos t \end{pmatrix}$$

Glued trees problem

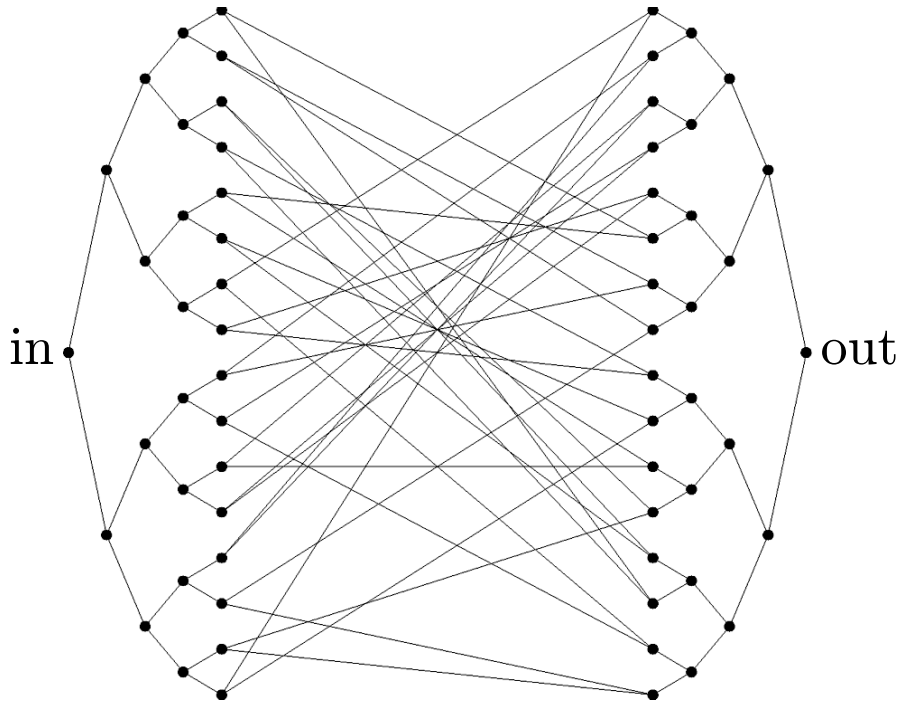

Black-box description of a graph

- Vertices have arbitrary labels
- Label of ‘in’ vertex is known
- Given a vertex label, black box returns labels of its neighbors
- Restricts algorithms to explore the graph locally

Glued trees problem: Classical query complexity

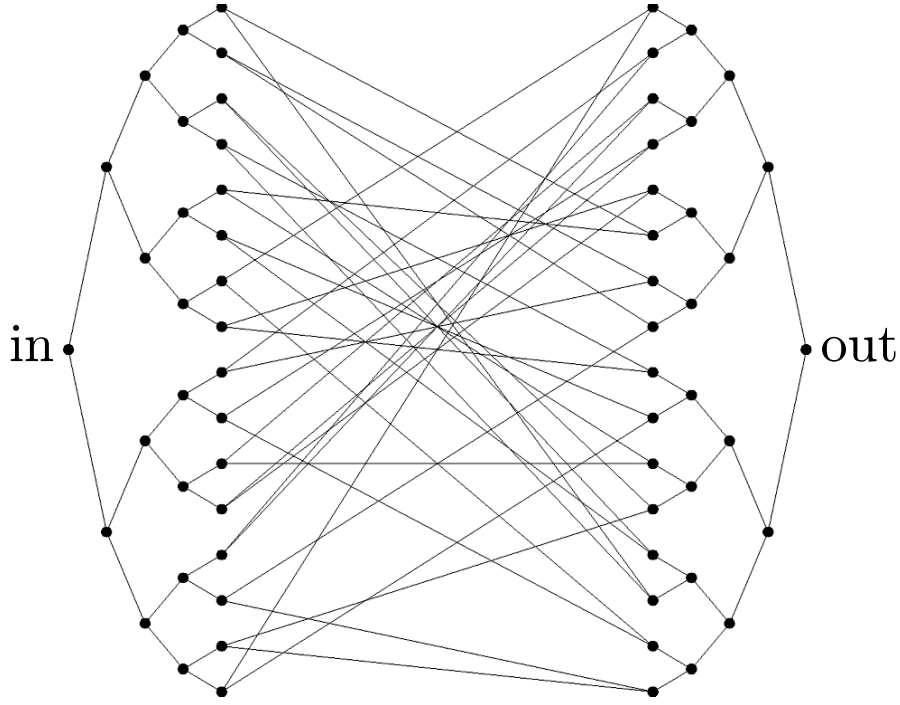

Let n denote the height of one of the binary trees

Classical random walk from 'in': probability of reaching 'out' is $2^{-\Omega(n)}$ at all times

In fact, the classical query complexity is $2^{\Omega(n)}$

Glued trees problem: Exponential speedup

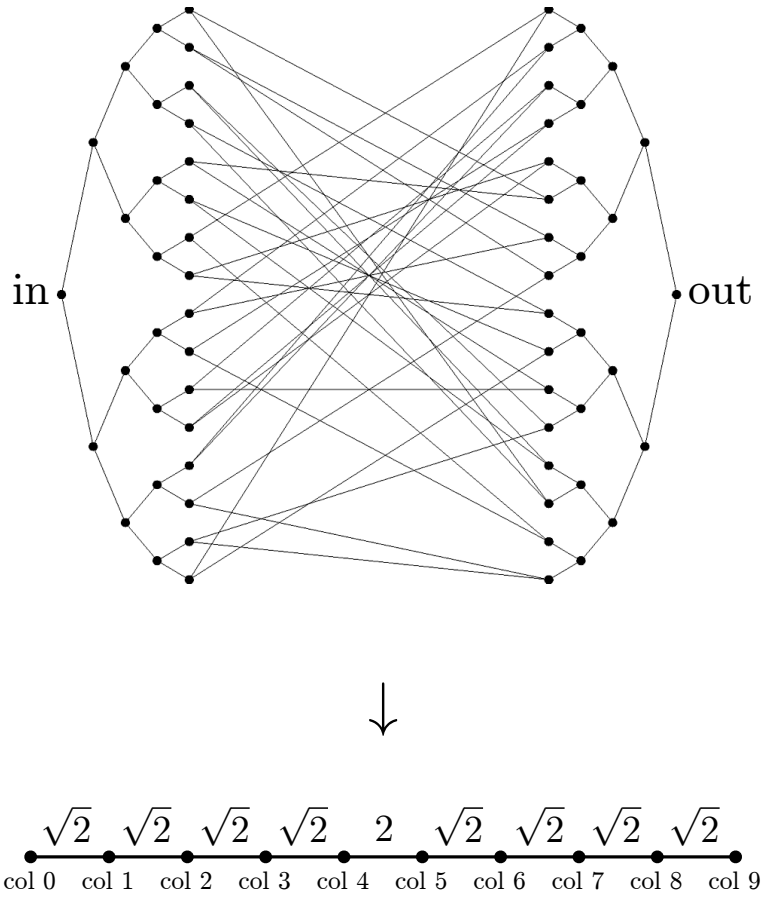

Column subspace

$$|\text{col } j\rangle := \frac{1}{\sqrt{N_j}} \sum_{v \in \text{column } j} |v\rangle$$

$$N_j := \begin{cases} 2^j & \text{if } j \in [0, n] \\ 2^{2n+1-j} & \text{if } j \in [n+1, 2n+1] \end{cases}$$

Reduced adjacency matrix

$$\langle \text{col } j | A | \text{col } j+1 \rangle = \begin{cases} \sqrt{2} & \text{if } j \in [0, n-1] \\ \sqrt{2} & \text{if } j \in [n+1, 2n] \\ 2 & \text{if } j = n \end{cases}$$

Discrete-time quantum walk: Need for a coin

Quantum analog of discrete-time random walk?

Unitary matrix $U \in \mathbb{C}^{|V|\times|V|}$ with $U_{vw} \neq 0$ iff $(v, w) \in E$

Consider the line:

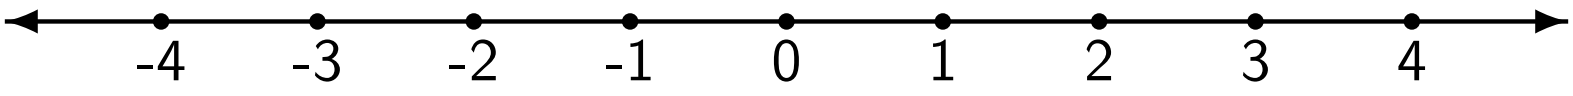

Define walk by $|x\rangle \mapsto \frac{1}{\sqrt{2}}(|x-1\rangle + |x+1\rangle)$?

But then $|x+2\rangle \mapsto \frac{1}{\sqrt{2}}(|x+1\rangle + |x+3\rangle)$, so this is not unitary!

In general, we must enlarge the state space.

Discrete-time quantum walk on a line

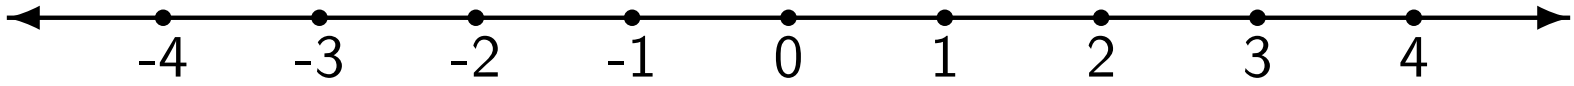

Add a "coin": state space $\text{span}\{|x\rangle \otimes |\leftarrow\rangle, |x\rangle \otimes |\rightarrow\rangle : x \in \mathbb{Z}\}$

Coin flip: $C := I \otimes H$

Shift:
$$S|x\rangle \otimes |\leftarrow\rangle = |x-1\rangle \otimes |\leftarrow\rangle$$
$$S|x\rangle \otimes |\rightarrow\rangle = |x+1\rangle \otimes |\rightarrow\rangle$$

Walk step: SC

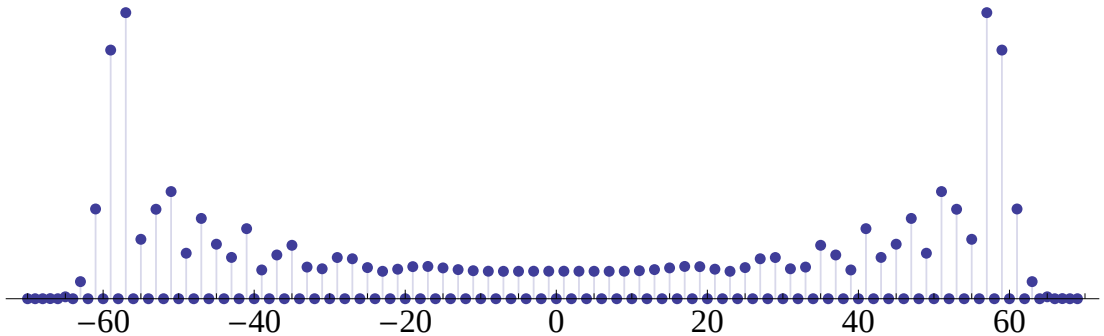

The Szegedy walk

State space: $\mathrm{span}\{|v\rangle \otimes |w\rangle, |w\rangle \otimes |v\rangle : (v, w) \in E\}$

Let W be a stochastic matrix (a discrete-time random walk)

$$\text{Define} \quad |\psi_v\rangle := |v\rangle \otimes \sum_{w \in V} \sqrt{W_{wv}} |w\rangle \quad (\text{note } \langle \psi_v | \psi_w \rangle = \delta_{v,w})$$

$$R := 2 \sum_{v \in V} |\psi_v\rangle\langle\psi_v| - I$$

$$S(|v\rangle \otimes |w\rangle) := |w\rangle \otimes |v\rangle$$

Then a step of the walk is the unitary operator $U := SR$

Spectrum of the walk

Let $T := \sum_{v \in V} |\psi_v\rangle\langle v|$, so $R = 2TT^\dagger - I$.

Theorem (Szegedy)

Let W be a stochastic matrix. Suppose the matrix

$$\sum_{v,w} \sqrt{W_{vw} W_{wv}} |w\rangle\langle v|$$

has an eigenvector $|\lambda\rangle$ with eigenvalue λ. Then

$$\frac{I - e^{\pm i \arccos \lambda} S}{\sqrt{2(1-\lambda^2)}} T|\lambda\rangle$$

are eigenvectors of $U = SR$ with eigenvalues

$$e^{\pm i \arccos \lambda}.$$

Proof of Szegedy's spectral theorem

Proof sketch.

Straightforward calculations give

$$TT^\dagger = \sum_{v \in V} |\psi_v\rangle\langle\psi_v| \qquad T^\dagger T = I$$

$$T^\dagger S T = \sum_{v,w \in V} \sqrt{W_{vw} W_{wv}} |w\rangle\langle v| = \sum_\lambda |\lambda\rangle\langle\lambda|$$

which can be used to show

$$U(T|\lambda\rangle) = ST|\lambda\rangle \qquad U(ST|\lambda\rangle) = 2\lambda ST|\lambda\rangle - T|\lambda\rangle.$$

Diagonalizing within the subspace span$\{T|\lambda\rangle, ST|\lambda\rangle\}$ gives the desired result. □

Exercise. Fill in the details

Random walk search algorithm

Given $G = (V, E)$, let $M \subset V$ be a set of *marked vertices*

Start at a random unmarked vertex

Walk until we reach a marked vertex:

$$W'_{vw} := \begin{cases} 1 & w \in M \text{ and } v = w \\ 0 & w \in M \text{ and } v \neq w \\ W_{vw} & w \notin M. \end{cases}$$

$$= \begin{pmatrix} W_M & 0 \\ V & I \end{pmatrix} \quad (W_M\text{: delete marked rows and columns of } W)$$

Question. How long does it take to reach a marked vertex?

Classical hitting time

Take t steps of the walk:

$$(W')^t = \begin{pmatrix} W_M^t & 0 \\ V(I + W_M + \cdots + W_M^{t-1}) & I \end{pmatrix}$$
$$= \begin{pmatrix} W_M^t & 0 \\ V\frac{I - W_M^t}{I - W_M} & I \end{pmatrix}$$

Convergence time depends on how close $\|W_M\|$ is to 1, which depends on the spectrum of W

Lemma

Let $W = W^T$ be a symmetric Markov chain. Let the second largest eigenvalue of W be $1 - \delta$, and let $\epsilon = |M|/|V|$ (the fraction of marked items). Then the probability of reaching a marked vertex is $\Omega(1)$ after $t = O(1/\delta\epsilon)$ steps of the walk.

Quantum walk search algorithm

Start from the state $\frac{1}{\sqrt{N-|M|}} \sum_{v \notin M} |\psi_v\rangle$

Consider the walk U corresponding to W':

$$\sum_{v,w \in V} \sqrt{W'_{v,w} W'_{w,v}} |w\rangle\langle v| = \begin{pmatrix} W_M & 0 \\ 0 & I \end{pmatrix}$$

Eigenvalues of U are $e^{\pm i \arccos \lambda}$ where the λ are eigenvalues of W_M

Perform phase estimation on U with precision $O(\sqrt{\delta\epsilon})$

- no marked items $\implies$ estimated phase is 0
- ϵ fraction of marked items $\implies$ nonzero phase with probability $\Omega(1)$

Further refinements give algorithms for *finding* a marked item

Grover's algorithm revisited

Problem

Given a black box $f : X \to \{0, 1\}$, *is there an* x *with* $f(x) = 1$*?*

Markov chain on $N = |X|$ vertices:

$$W := \frac{1}{N} \begin{pmatrix} 1 & \cdots & 1 \\ \vdots & \ddots & \vdots \\ 1 & \cdots & 1 \end{pmatrix} = |\psi\rangle\langle\psi|, \quad |\psi\rangle := \frac{1}{\sqrt{N}} \sum_{x \in X} |x\rangle$$

Eigenvalues of W are $0, 1 \implies \delta = 1$

Hard case: one marked vertex, $\epsilon = 1/N$

Hitting times

- Classical: $O(1/\delta\epsilon) = O(N)$
- Quantum: $O(1/\sqrt{\delta\epsilon}) = O(\sqrt{N})$

Element distinctness

Problem

Given a black box $f : X \to Y$*, are there distinct* x, x' *with* $f(x) = f(x')$*?*

Let $N = |X|$; classical query complexity is $\Omega(N)$

Consider a quantum walk on the Hamming graph $H(N, M)$

- Vertices: $\{(x_1, \ldots, x_M) : x_i \in X\}$
- Store the values $(f(x_1), \ldots, f(x_M))$ at vertex $(x_1, \ldots, x_M)$
- Edges between vertices that differ in exactly one coordinate

Element distinctness: Analysis

Spectral gap: $\delta = O(1/M)$

Fraction of marked vertices: $\epsilon \geq 2\binom{M}{2} N^{M-2}/N^M = \Theta(M^2/N^2)$

Quantum hitting time: $O(1/\sqrt{\delta\epsilon}) = O(N/\sqrt{M})$

Quantum query complexity:

- M queries to prepare the initial state
- 2 queries for each step of the walk (compute f, uncompute f)
- Overall: $M + O(N/\sqrt{M})$

Choose $M = N^{2/3}$: query complexity is $O(N^{2/3})$ (optimal!)

Quantum walk algorithms

Quantum walk search algorithms

- Spatial search
- Subgraph finding
- Checking matrix multiplication
- Testing if a black-box group is abelian

Evaluating Boolean formulas

Exponential speedup for a natural problem?

Exercise: Triangle finding (1/2)

The goal of the *triangle problem* is to decide whether an n-vertex graph G contains a triangle (a complete subgraph on 3 vertices). The graph is specified by a black box that, for any pair of vertices of G, returns a bit indicating whether those vertices are connected by an edge in G.

1. What is the classical query complexity of the triangle problem?
2. Say that an edge of G is a *triangle edge* if it is part of a triangle in G. What is the quantum query complexity of deciding whether a particular edge of G is a triangle edge?
3. Now suppose you know the vertices and edges of some m-vertex subgraph of G. Explain how you can decide whether this subgraph contains a triangle edge using $O(m^{2/3}\sqrt{n})$ quantum queries.

Exercise: Triangle finding (2/2)

4. Consider a quantum walk algorithm for the triangle problem. The walk takes place on a graph $\mathcal{G}$ whose vertices correspond to subgraphs of G on m vertices, and whose edges correspond to subgraphs that differ by changing one vertex. A vertex of $\mathcal{G}$ is marked if it contains a triangle edge. How many queries does this algorithm use to decide whether G contains a triangle? (Hint: Be sure to account for the S queries used to initialize the walk, the U queries used to move between neighboring vertices of $\mathcal{G}$, and the C queries used to check whether a given vertex of $\mathcal{G}$ is marked. If the walk has spectral gap δ and an ϵ-fraction of the vertices are marked, it can be shown that there is a quantum walk search algorithm with query complexity $S + \frac{1}{\sqrt{\epsilon}}(\frac{1}{\sqrt{\delta}}U + C)$.)
5. Choose a value of m that minimizes the number of queries used by the algorithm. What is the resulting upper bound on the quantum query complexity of the triangle problem?

Part VI

Adversary lower bounds

Query complexity

Task: Compute a function $f : S \to T$

$S \subseteq \Sigma^n$ is the set of possible inputs, where Σ is the *input alphabet*

- if $S = \Sigma^n$ then f is *total*
- if $S \subsetneq \Sigma^n$ then f is *partial*

Input $x \in S$ is specified by a black box:

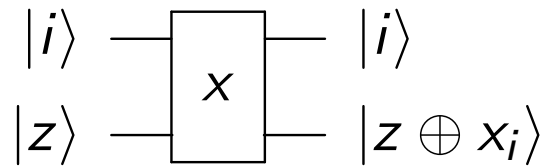

where $i \in \{1, \ldots, n\}$

Query algorithms

Structure of a quantum query algorithm:

- Initial state $|\psi\rangle$ does not depend on the oracle string x
- Alternate between queries to the black box O_x and non-query operations $U_1, U_2, \ldots, U_t$

$$|\psi_x^t\rangle := U_t O_x \ldots U_2 O_x U_1 O_x |\psi\rangle$$

- End with a measurement in the computational basis

Goal: Compute $f(x)$ using as few queries as possible

Query models

Three natural models for the query complexity of f:

- $D(f)$: deterministic query complexity
  (algorithm is classical and must always work correctly)
- $R_\epsilon(f)$: randomized query complexity with (two-sided) error probability at most ϵ
- $Q_\epsilon(f)$: quantum query complexity with (two-sided) error probability at most ϵ

For any constant ϵ,
$R_\epsilon(f) = \Theta(R_{1/3}(f))$ and $Q_\epsilon(f) = \Theta(Q_{1/3}(f))$
(repeat several times and take a majority vote)

Clearly $Q_\epsilon(f) \leq R_\epsilon(f) \leq D(f)$

Quantum queries: Boolean case

Consider $\Sigma = \{0, 1\}$

Bit flip oracle:

$$\hat{O}_x|i, b\rangle = |i, b \oplus x_i\rangle \qquad \text{for } i \in \{1, \ldots, n\},\ b \in \{0, 1\}$$

Phase flip oracle:

$$O_x|i, b\rangle = (-1)^{bx_i}|i, b\rangle \qquad \text{for } i \in \{1, \ldots, n\},\ b \in \{0, 1\}$$

Phase kickback: $O_x = (I \otimes H)\hat{O}_x(I \otimes H)$

Note: $O_x|i, 0\rangle = |i, 0\rangle$ for all i is wasteful; alternatively, use

$$O'_x|i\rangle = \begin{cases} (-1)^{x_i}|i\rangle & i \in \{1, \ldots, n\} \\ |i\rangle & i = 0 \quad (\text{i.e., } x_0 := 1) \end{cases}$$

But the ability to not query the phase oracle is essential!

Quantum queries: General case

Similar considerations hold when $|\Sigma| = d > 2$

Let $\Sigma = \mathbb{Z}_d$ without loss of generality

Addition oracle:

$$\hat{O}_x|i, b\rangle = |i, b + x_i \bmod d\rangle \qquad \text{for } i \in \{1, \ldots, n\},\ b \in \mathbb{Z}_d$$

Phase oracle:

$$O_x|i, b\rangle = e^{2\pi i b x_i / d}|i, b\rangle \qquad \text{for } i \in \{1, \ldots, n\},\ b \in \mathbb{Z}_d$$

Phase kickback:

$$O_x = (I \otimes F^\dagger)\hat{O}_x(I \otimes F)$$

where F is the Fourier transform over $\mathbb{Z}_d$

A quantum adversary

Lower bound strategy: Oracle is operated by a malicious adversary

Adversary creates a superposition over possible inputs: $\sum_{x \in S} a_x |x\rangle$

Each query is performed by the "super-oracle"

$$O := \sum_{x \in S} |x\rangle\langle x| \otimes O_x$$

After t steps, algorithm produces the state

$$|\psi^t\rangle := (I \otimes U_t) O \ldots (I \otimes U_2) O (I \otimes U_1) O \left(\sum_{x \in S} a_x |x\rangle \otimes |\psi\rangle \right)$$

$$= \sum_{x \in S} a_x |x\rangle \otimes |\psi_x^t\rangle$$

Getting entangled with the adversary

Intuition: To learn x, the state $|\psi^t\rangle$ must be very entangled

Reduced density matrix of the oracle:

$$\rho^t := \sum_{x,y \in S} a_x^* a_y \langle \psi_x^t | \psi_y^t \rangle \, |x\rangle\langle y|$$

Initial state ρ^0 is pure

Final state ρ^t must be mixed

Quantify how much more mixed the state can become with a single query

We could consider the von Neumann entropy of ρ^t, but this is cumbersome

Distinguishing quantum states

Fact

Given one of two pure states $|\psi\rangle, |\phi\rangle$, we can make a measurement that determines which state we have with error probability at most ϵ if and only if $|\langle\psi|\phi\rangle| \leq 2\sqrt{\epsilon(1-\epsilon)}$.

Exercise. Prove this

So it's convenient to consider measures that are linear in the inner products $\langle\psi_x^t|\psi_y^t\rangle$

Adversary matrices

The adversary bound uses a matrix $\Gamma \in \mathbb{R}^{|S| \times |S|}$

$\Gamma_{x,y}$ measures how hard it is to distinguish between x and y

We say Γ is an *adversary matrix* if

1. $\Gamma_{xy} = \Gamma_{yx}$,
2. $\Gamma_{xy} \geq 0$, and
3. if $f(x) = f(y)$ then $\Gamma_{xy} = 0$

Weight function

Given an adversary matrix Γ, we define a weight function

$$W_j := \sum_{x,y \in S} \Gamma_{xy} a_x^* a_y \langle \psi_x^j | \psi_y^j \rangle$$

We show:

1. W_0 is large
2. To compute f in t queries, W_t must be small
3. W_{j+1} cannot be too much smaller than W_j

Weight function: Initial value

The initial value of the weight function is

$$W_0 = \sum_{x,y \in S} \Gamma_{xy} a_x^* a_y \langle \psi_x^0 | \psi_y^0 \rangle$$
$$= \sum_{x,y \in S} a_x^* \Gamma_{xy} a_y$$

since $|\psi_x^0\rangle$ cannot depend on x

To make this as large as possible, take a to be a principal eigenvector of Γ

$\Rightarrow W_0 = \|\Gamma\|$

Weight function: Final value

If $f(x) \neq f(y)$ then the states $|\psi_x^t\rangle, |\psi_y^t\rangle$ must be distinguishable

To succeed with error probability at most ϵ, we need $|\langle\psi_x^t|\psi_y^t\rangle| \leq 2\sqrt{\epsilon(1-\epsilon)}$

Thus

$$\begin{aligned} W_t &= \sum_{x,y \in S} \Gamma_{xy} a_x^* a_y \langle\psi_x^t|\psi_y^t\rangle \\ &\leq \sum_{x,y \in S} \Gamma_{xy} a_x^* a_y 2\sqrt{\epsilon(1-\epsilon)} \\ &= 2\sqrt{\epsilon(1-\epsilon)}\|\Gamma\| \end{aligned}$$

Weight function: Making a query (1/5)

Change in weight function:

$$W_{j+1} - W_j = \sum_{x,y \in S} \Gamma_{xy} a_x^* a_y (\langle \psi_x^{j+1} | \psi_y^{j+1} \rangle - \langle \psi_x^j | \psi_y^j \rangle)$$

Change in state: $|\psi_x^{j+1}\rangle = U_{j+1} O_x |\psi_x^j\rangle$

Gram matrix elements:

$$\begin{aligned} \langle \psi_x^{j+1} | \psi_y^{j+1} \rangle &= \langle \psi_x^j | O_x^\dagger U_{j+1}^\dagger U_{j+1} O_y | \psi_y^j \rangle \\ &= \langle \psi_x^j | O_x O_y | \psi_y^j \rangle \end{aligned}$$

Therefore

$$W_{j+1} - W_j = \sum_{x,y \in S} \Gamma_{xy} a_x^* a_y \langle \psi_x^j | (O_x O_y - I) | \psi_y^j \rangle$$

Weight function: Making a query (2/5)

$$W_{j+1} - W_j = \sum_{x,y \in S} \Gamma_{xy} a_x^* a_y \langle \psi_x^j | (O_x O_y - I) | \psi_y^j \rangle$$

We have $O_x O_y |i, b\rangle = (-1)^{b(x_i \oplus y_i)} |i, b\rangle$

Let $P_0 = I \otimes |0\rangle\langle 0|$ and $P_i = |i, 1\rangle\langle i, 1|$

Then

$$O_x O_y - I = P_0 + \sum_{i=1}^{n} (-1)^{x_i \oplus y_i} P_i - I$$

$$= -2 \sum_{i \,:\, x_i \neq y_i}^{n} P_i$$

Weight function: Making a query (3/5)

$$O_x O_y - I = -2 \sum_{i\,:\, x_i \neq y_i}^{n} P_i$$

so

$$\begin{aligned}
|W_{j+1} - W_j| &= \sum_{x,y \in S} \Gamma_{xy} a_x^* a_y \langle \psi_x^j | (O_x O_y - I) | \psi_y^j \rangle \\
&= 2 \left| \sum_{x,y \in S} \sum_{i\,:\, x_i \neq y_i} \Gamma_{xy} a_x^* a_y \langle \psi_x^j | P_i | \psi_y^j \rangle \right| \\
&\leq 2 \sum_{x,y \in S} \sum_{i\,:\, x_i \neq y_i} \Gamma_{xy} |a_x^* a_y \langle \psi_x^j | P_i | \psi_y^j \rangle| \qquad (\triangle) \\
&\leq 2 \sum_{x,y \in S} \sum_{i\,:\, x_i \neq y_i} \Gamma_{xy} \| a_x P_i | \psi_x^j \rangle \| \cdot \| a_y P_i | \psi_y^j \rangle \| \qquad \text{(C-S)}
\end{aligned}$$

Weight function: Making a query (4/5)

For each $i \in \{1, \ldots, n\}$, define $\Gamma_i \in \mathbb{R}^{|S| \times |S|}$ by

$$(\Gamma_i)_{xy} = \begin{cases} \Gamma_{xy} & \text{if } x_i \neq y_i \\ 0 & \text{if } x_i = y_i, \end{cases}$$

and define vectors v_i with components $(v_i)_x = \|a_x P_i |\psi_x^j\rangle\|$

$$\begin{aligned} |W_{j+1} - W_j| &\leq 2 \sum_{x,y \in S} \sum_{i=1}^{n} (v_i)_x (\Gamma_i)_{xy} (v_i)_y \\ &= 2 \sum_{i=1}^{n} v_i^\dagger \Gamma_i v_i \\ &\leq 2 \sum_{i=1}^{n} \|\Gamma_i\| \cdot \|v_i\|^2 \end{aligned}$$

Weight function: Making a query (5/5)

$$|W_{j+1} - W_j| \leq 2 \sum_{i=1}^{n} \|\Gamma_i\| \cdot \|v_i\|^2$$

Since

$$\begin{aligned}
\sum_{i=1}^{n} \|v_i\|^2 &= \sum_{i=1}^{n} \sum_{x \in S} \|a_x P_i |\psi_x^j\rangle\|^2 \\
&\leq \sum_{x \in S} a_x^2 \| |\psi_x^j\rangle \|^2 \\
&= \sum_{x \in S} a_x^2 \\
&= 1,
\end{aligned}$$

we have

$$|W_{j+1} - W_j| \leq 2 \max_{i \in \{1,\ldots,n\}} \|\Gamma_i\|$$

Weight function: Putting everything together

Since $W_0 = \|\Gamma\|$, we have

$$W_t \geq \|\Gamma\| - 2t \max_{i \in \{1,\ldots,n\}} \|\Gamma_i\|$$

So $W_t \leq 2\sqrt{\epsilon(1-\epsilon)}\|\Gamma\|$ implies

$$t \geq \frac{1 - 2\sqrt{\epsilon(1-\epsilon)}}{2} \mathrm{Adv}(f)$$

where

$$\mathrm{Adv}(f) := \max_{\Gamma} \frac{\|\Gamma\|}{\max_{i \in \{1,\ldots,n\}} \|\Gamma_i\|}$$

with the maximum taken over all adversary matrices Γ

Example: Unstructured search (1/3)

Problem: Distinguish no marked item from unique marked item

$$S = \{000\ldots00, 100\ldots00, 010\ldots00, \ldots, 000\ldots01\}$$

Adversary matrix:

$$\Gamma = \begin{pmatrix} 0 & \gamma_1 & \cdots & \gamma_n \\ \gamma_1 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_n & 0 & \cdots & 0 \end{pmatrix} \qquad \gamma_1, \ldots, \gamma_n \geq 0$$

Symmetry: $\gamma_1 = \cdots = \gamma_n = 1$

Example: Unstructured search (2/3)

Consider

$$\Gamma^2 = \begin{pmatrix} n & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 1 & \cdots & 1 \end{pmatrix}$$

$\|\Gamma^2\| = n$, so $\|\Gamma\| = \sqrt{n}$

$$\|\Gamma_i\| = \|\Gamma_1\| = \left\| \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{pmatrix} \right\| = 1$$

Example: Unstructured search (3/3)

Our adversary matrix has $\|\Gamma\| = \sqrt{n}$, $\|\Gamma_i\| = 1$

So $\mathrm{Adv}(\mathrm{OR}) \geq \frac{\|\Gamma\|}{\|\Gamma_i\|} = \sqrt{n}$

Therefore $Q_\epsilon(\mathrm{OR}) \geq \frac{1-2\sqrt{\epsilon(1-\epsilon)}}{2}\sqrt{n}$

Thus Grover's algorithm is optimal up to a constant factor (recall that Grover's algorithm finds a unique marked item with probability $1 - o(1)$ in $(\frac{\pi}{4} + o(1))\sqrt{n}$ queries)

Other adversaries

The adversary method described above is a generalization of the method originally formulated by Ambainis, which considered only a relation between yes and no inputs and did not allow arbitrary positive weights.

More recently, it was realized that one can use negative weights and still obtain a lower bound, and that sometimes this bound can be dramatically better.

In fact, it was shown by Reichardt that the adversary bound allowing negative weights is essentially tight: up to constant factors, it characterizes quantum query complexity.

Exercise: Original formulation of the adversary method

Choose $X, Y \subset \{0,1\}^n$ such that $f(x) \neq f(y)$ for all $x \in X, y \in Y$. For any relation $R \subset X \times Y$, define

$$m := \min_{x \in X} |\{y \in Y : (x,y) \in R|$$

$$m' := \min_{y \in Y} |\{x \in X : (x,y) \in R|$$

$$\ell := \max_{\substack{x \in X \\ i \in \{1,\ldots,n\}}} |\{y \in Y : (x,y) \in R \text{ and } x_i \neq y_i\}|$$

$$\ell' := \max_{\substack{y \in Y \\ i \in \{1,\ldots,n\}}} |\{x \in X : (x,y) \in R \text{ and } x_i \neq y_i\}|.$$

Then define $\mathrm{Amb}(f) := \max_{X,Y,R} \sqrt{\frac{mm'}{\ell\ell'}}$.

Prove that $\mathrm{Adv}(f) \geq \mathrm{Amb}(f)$, and hence that $Q_\epsilon(f) \geq \frac{1-2\sqrt{\epsilon(1-\epsilon)}}{2}\mathrm{Amb}(f)$.

Exercise: Applying the adversary method

1. Define $\textsc{parity}\colon \{0,1\}^n \to \{0,1\}$ by $\textsc{parity}(x) = x_1 \oplus \cdots \oplus x_n$. Show that $Q(\textsc{parity}) = \Omega(n)$.
2. Define $\textsc{nand}^2\colon \{0,1\}^{n^2} \to \{0,1\}$ by

$$\textsc{nand}^2(x) = \textsc{nand}\big(\textsc{nand}(x_1,\ldots,x_n), \textsc{nand}(x_{n+1},\ldots,x_{2n}), \ldots, \textsc{nand}(x_{n^2-n+1},\ldots,x_{n^2})\big).$$

   Show that $Q(\textsc{nand}^2) = \Omega(n)$.
3. Let $x \in \{0,1\}^{\binom{n}{2}}$ specify the edges of a simple, undirected n-vertex graph, and define $\textsc{con}\colon \{0,1\}^{\binom{n}{2}} \to \{0,1\}$ by

$$\textsc{con}(x) = \begin{cases} 1 & \text{if the graph described by } x \text{ is connected} \\ 0 & \text{otherwise.} \end{cases}$$

   Show that $Q(\textsc{con}) = \Omega(n^{3/2})$.