Quantum algorithms

Andrew Childs
Institute for Quantum Computing
University of Waterloo

USEQIP

2 June 2011

Based on slides prepared with Pawel Wocjan (University of Central Florida)
Original full version at www.math.uwaterloo.ca/∼amchilds/talks/qalg.pdf

Outline

I. Quantum circuits
II. Elementary quantum algorithms
III. The QFT and phase estimation
IV. Factoring
V. Quantum search
VI. Beyond Shor and Grover

Part I

Quantum circuits

Quantum circuits

Quantum circuits are generalizations of Boolean circuits

input transformation output (probabilistic)

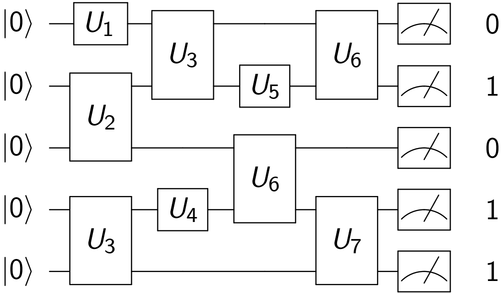

Quantum circuit model

To quantify complexity, a quantum algorithm must be implemented by a quantum circuit, i.e., a sequence of elementary gates

$$\text{Quantum circuit model} \quad = \quad \begin{array}{c}\text{Quantum mechanics}\\ + \\ \text{Notion of complexity}\end{array}$$

A universal gate set

Every unitary can be implemented exactly by quantum circuits using only

- CNOT gates (acting on adjacent qubits) and
- arbitrary single qubit gates

The gate complexity $\kappa(U)$ of a unitary $U \in \mathcal{U}(\mathcal{H})$ is minimal number of elementary gates needed to implement U

Example: quantum Fourier transform has gate complexity $O(n^2)$

(Approximately) universal gate sets

For every $\epsilon \in (0,1)$ and every unitary U, there is a unitary V such that

$$\|U - V\| \leq \epsilon \quad \text{where} \quad \|U - V\| = \sup_{|\psi\rangle} \|(U - V)|\psi\rangle\|$$

where V is implemented by a quantum circuits using only

- CNOT gates (acting on adjacent qubits)
- the single-qubit gates $H, R(\frac{\pi}{4})$ where

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad R(\theta) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{pmatrix}$$

There are other universal gate sets

Gate complexity of unitaries

The gate complexity $\kappa_\epsilon(U)$ of a unitary U is the minimal number of gates (from a universal gate set) need to implement a unitary V with $\|U - V\| \leq \epsilon$

The Solovay-Kitaev theorem implies that

$$\kappa_\epsilon(U) = O\Big(\kappa(U) \cdot \log^c \big(\kappa(U)/\epsilon\big)\Big)$$

for some small constant c

Counting arguments show that most n-qubit unitaries have gate complexity exponential in n

Structure of quantum algorithms

An efficient quantum algorithm consists of

- preparing the initial state $|0\rangle^{\otimes n}$,
- applying a quantum circuit of polynomially many in n gates from some universal gate set, and
- measuring all qubits in the computational basis

These steps can be repeated polynomially many times to collect statistics, followed by efficient classical post-processing

Reversible computing

The classical AND gate is irreversible because if the output is 0 then we cannot determine which of the three possible pairs was the actual input

x_1	x_2	$x_1 \wedge x_2$
0	0	0
0	1	0
1	0	0
1	1	1

But it is easy to simulate the AND gate with one Toffoli gate

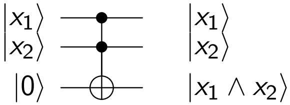

Problem of garbage

To simulate irreversible circuits with Toffoli gates, we keep the input and intermediate results to make everything reversible

Consider the function $y = x_1 \wedge x_2 \wedge x_3$

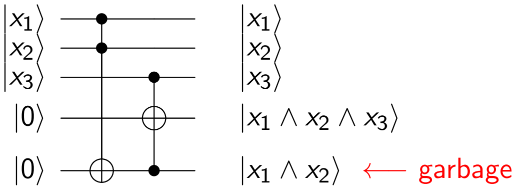

It is important to not leave any garbage; otherwise, we could not make use of quantum parallelism and constructive interference effects

Reversible garbage removal

It is always possible to reversibly remove (uncompute) the garbage

In the case $y = x_1 \wedge x_2 \wedge x_3$, this can be done with the circuit

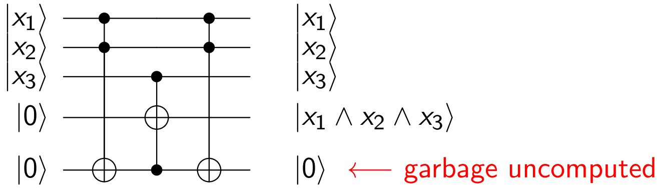

Simulating irreversible circuits

Let $f : \{0,1\}^n \to \{0,1\}$ be any boolean function

Assume this function can be computed classically using only t classical elementary gates such as AND, OR, NAND

We can implement a unitary U_f on $(\mathbb{C}^2)^{\otimes n} \otimes \mathbb{C}^2 \otimes (\mathbb{C}^2)^{\otimes w}$ such that

$$U_f \left(|x\rangle_{\mathrm{in}} \otimes |y\rangle_{\mathrm{out}} \otimes |0\rangle_{\mathrm{work}}^{\otimes w} \right) = |x\rangle \otimes |y \oplus f(x)\rangle \otimes |0\rangle^{\otimes w}$$

U_f is built from polynomially many in t Toffoli gates and the size w of the workspace register is polynomial in t

During the computation the qubits of the workspace register are changed, but at the end they reversibly reset to $|0\rangle^{\otimes w}$

Part II

Elementary quantum algorithms

Black box problems

Standard computational problem: determine a property of some input data

- Example: Find the prime factors of N

Alternate model: Input is provided by a *black box* (or *oracle*)

- Query: On input x, black box returns $f(x)$
- Determine a property of f using as few queries as possible
- The minimum number of queries is the *query complexity*
- Example: Given a black box for $f : \{1, 2, \ldots, N\} \to \{0, 1\}$, is there some x such that $f(x) = 1$?
- Why black boxes?
  - Facilitates proving lower bounds
  - Can lead to algorithms for standard problems

Black boxes for reversible/quantum computing

Black box 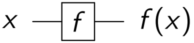

is not reversible

Reversible version:

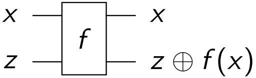

Given a circuit that computes f non-reversibly, we can implement the reversible version with little overhead

Quantum version:

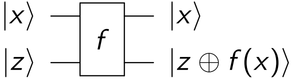

A reversible circuit is a quantum circuit

Deutsch's problem

Problem

- Given: a black-box function $f : \{0,1\} \to \{0,1\}$
- Task: determine whether f is *constant* or *balanced*

x	$f_1(x)$
0	0
1	0

x	$f_2(x)$
0	1
1	1

constant: $f(0) = f(1)$

x	$f_3(x)$
0	0
1	1

x	$f_4(x)$
0	1
1	0

balanced: $f(0) \neq f(1)$

How many queries are needed?

- Classically: 2 queries are necessary and sufficient
- Quantumly: ?

Toward a quantum algorithm for Deutsch's problem

Quantum black box for f:

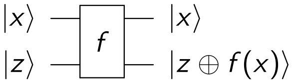

Compute f in superposition:

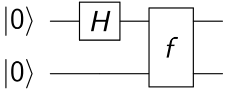

$$
\begin{aligned}
|0\rangle \otimes |0\rangle &\mapsto \frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes |0\rangle \\
&\mapsto \frac{1}{\sqrt{2}}(|0\rangle \otimes |f(0)\rangle + |1\rangle \otimes |f(1)\rangle)
\end{aligned}
$$

Can't extract more than one bit of information about f

Phase kickback

Quantum black box for f:

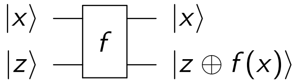

Phase kickback:

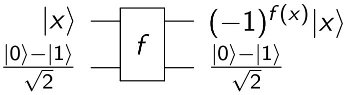

$$\begin{aligned}
|x\rangle \otimes \tfrac{1}{\sqrt{2}}(|0\rangle - |1\rangle) &= \tfrac{1}{\sqrt{2}}(|x\rangle \otimes |0\rangle - |x\rangle \otimes |1\rangle) \\
&\mapsto \tfrac{1}{\sqrt{2}}(|x\rangle \otimes |f(x)\rangle - |x\rangle \otimes |1 \oplus f(x)\rangle) \\
&= |x\rangle \otimes \tfrac{1}{\sqrt{2}}(|f(x)\rangle - |\overline{f(x)}\rangle) \\
&= \underbrace{(-1)^{f(x)}}_{\text{not necessarily global}} |x\rangle \otimes \tfrac{1}{\sqrt{2}}(|0\rangle - |1\rangle)
\end{aligned}$$

Quantum algorithm for Deutsch's problem

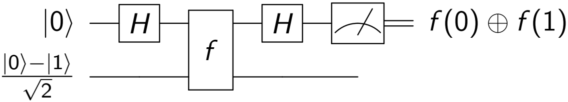

$$
\begin{aligned}
|0\rangle \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}} &\mapsto \frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}} \\
&\mapsto \frac{(-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}} \\
&= (-1)^{f(0)} \frac{|0\rangle + (-1)^{f(0)\oplus f(1)}|1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}} \\
&\mapsto (-1)^{f(0)} |f(0) \oplus f(1)\rangle \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}}
\end{aligned}
$$

1 quantum query vs. 2 classical queries!

The Deutsch-Jozsa problem

Problem

- Given: a black-box function $f : \{0,1\}^n \to \{0,1\}$
- Promise: f is either
  constant ($f(x)$ is independent of x)
  or *balanced* ($f(x) = 0$ for exactly half the values of x)
- Task: determine whether f is constant or balanced

How many queries are needed?

- Classically: $2^n/2 + 1$ queries to answer with certainty
- Quantumly: ?

Phase kickback for a Boolean function of n bits

Black box function:

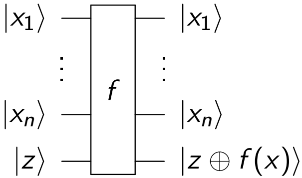

Phase kickback:

$$|x_1\rangle \otimes \cdots \otimes |x_n\rangle \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}} \mapsto (-1)^{f(x)} |x_1\rangle \otimes \cdots \otimes |x_n\rangle \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

Quantum algorithm for the Deutsch-Jozsa problem

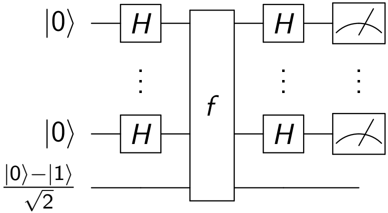

$$|0\rangle^{\otimes n} \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}} \mapsto \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}}\right)^{\otimes n} \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

$$= \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

$$\mapsto \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle \otimes \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

Hadamard transform

What do the final Hadamard gates do?

$$\begin{aligned} H|x\rangle &= \frac{1}{\sqrt{2}}(|0\rangle + (-1)^x|1\rangle) \\ &= \frac{1}{\sqrt{2}} \sum_{y \in \{0,1\}} (-1)^{xy}|y\rangle \end{aligned}$$

$$\begin{aligned} H^{\otimes n}(|x_1\rangle \otimes \cdots \otimes |x_n\rangle) &= \bigotimes_{i=1}^{n} \left(\frac{1}{\sqrt{2}} \sum_{y_i \in \{0,1\}} (-1)^{x_i y_i}|y_i\rangle \right) \\ &= \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{x \cdot y}|y\rangle \end{aligned}$$

Quantum D-J algorithm: Finishing up

$$\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle \overset{H^{\otimes n}}{\mapsto} \frac{1}{2^n} \sum_{x,y \in \{0,1\}^n} (-1)^{f(x)} (-1)^{x \cdot y} |y\rangle$$

- If f is constant, the amplitude of $|y\rangle$ is

$$\pm \frac{1}{2^n} \sum_{x \in \{0,1\}^n} (-1)^{x \cdot y} = \pm \begin{cases} 1 & \text{if } y = 0 \dots 0 \\ 0 & \text{otherwise} \end{cases}$$

  so we definitely measure $0 \dots 0$

- If f is balanced, the amplitude of $|0 \dots 0\rangle$ is

$$\sum_{x \in \{0,1\}^n} (-1)^{f(x)} = 0$$

  so we measure some nonzero string

The Deutsch-Jozsa problem: Quantum vs. classical

Above quantum algorithm uses only one query.

Need $2^n/2 + 1$ classical queries to answer with certainty.

What about randomized algorithms? Success probability arbitrarily close to 1 with a constant number of queries.

Can we get a separation between randomized and quantum computation?

Simon's problem

Problem

- Given: a black-box function $f : \{0,1\}^n \to \{0,1\}^m$
- Promise: there is some $s \in \{0,1\}^n$ such that $f(x) = f(y)$ if and only if $x = y$ or $x = y \oplus s$
- Task: determine s

One classical strategy:

- query a random x
- repeat until we find $x_i \neq x_j$ such that $f(x_i) = f(x_j)$
- output $x_i \oplus x_j$

By the birthday problem, this uses about $\sqrt{2^n}$ queries.

It can be shown that this strategy is essentially optimal.

Quantum algorithm for Simon's problem

Quantum black box: $|x\rangle \otimes |y\rangle \mapsto |x\rangle \otimes |y \oplus f(x)\rangle$

$$(x \in \{0,1\}^n,\ y \in \{0,1\}^m)$$

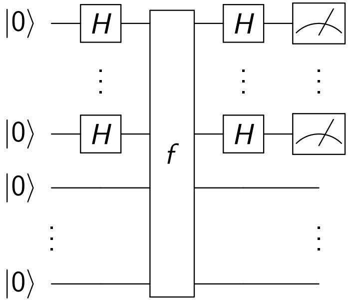

Repeat many times and post-process the measurement outcomes

Quantum algorithm for Simon's problem: Analysis I

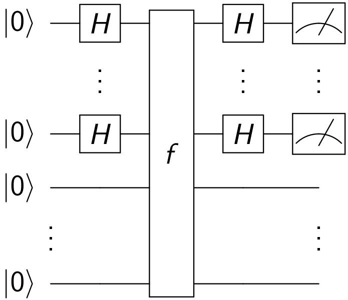

$$|0\rangle^{\otimes n} \otimes |0\rangle^{\otimes m}$$

$$\mapsto \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle \otimes |0\rangle^{\otimes m}$$

$$\mapsto \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle \otimes |f(x)\rangle$$

$$= \frac{1}{\sqrt{2^{n-1}}} \sum_{x \in R} \frac{|x\rangle + |x \oplus s\rangle}{\sqrt{2}} \otimes |f(x)\rangle$$

for some $R \subset \{0,1\}^n$

Quantum algorithm for Simon's problem: Analysis II

Recall $H^{\otimes n}|x\rangle = \sum_{y \in \{0,1\}^n} (-1)^{x \cdot y} |y\rangle$

$$H^{\otimes n}\left(\frac{|x\rangle + |x \oplus s\rangle}{\sqrt{2}}\right) = \frac{1}{\sqrt{2^{n+1}}} \sum_{y \in \{0,1\}^n} [(-1)^{x \cdot y} + (-1)^{(x \oplus s) \cdot y}]|y\rangle$$

$$= \frac{1}{\sqrt{2^{n+1}}} \sum_{y \in \{0,1\}^n} (-1)^{x \cdot y}[1 + (-1)^{s \cdot y}]|y\rangle$$

Two cases:

- if $s \cdot y = 0 \bmod 2$, $1 + (-1)^{s \cdot y} = 2$
- if $s \cdot y = 1 \bmod 2$, $1 + (-1)^{s \cdot y} = 0$

Measuring gives a random y orthogonal to s (i.e., $s \cdot y = 0$)

Quantum algorithm for Simon's problem: Post-processing

Measuring gives a random y orthogonal to s ($s \cdot y = 0$)

Repeat k times, giving vectors $y_1, \ldots, y_k \in \{0,1\}^n$; solve a system of k linear equations for $s \in \{0,1\}^n$:

$$y_1 \cdot s = 0, \quad y_2 \cdot s = 0, \quad \ldots, \quad y_k \cdot s = 0$$

How big should k be to give a unique (nonzero) solution?

- Clearly $k \geq n-1$ is necessary
- It can be shown that $k = O(n)$ suffices

$O(n)$ quantum queries, $O(n^3)$ quantum gates

Compare to $\Omega(2^{n/2})$ classical queries (even for bounded error)

Recap

We have seen several examples of quantum algorithms that outperform classical computation:

- Deutsch's problem: 1 quantum query vs. 2 classical queries
- Deutsch-Jozsa problem: 1 quantum query vs. $2^{\Omega(n)}$ classical queries (deterministic)
- Simon's problem: $O(n)$ quantum queries vs. $2^{\Omega(n)}$ classical queries (randomized)

Quantum algorithms for more interesting problems build on the tools used in these examples.

Part III

The QFT and phase estimation

Quantum phase estimation

Problem

We are given a unitary U and an eigenvector $|\psi\rangle$ of U with unknown eigenvalue

We seek to estimate its eigenphase $\varphi \in [0, 1)$ such that

$$U|\psi\rangle = e^{2\pi i \varphi}|\psi\rangle$$

Phase kickback for U

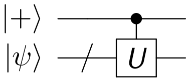

$$\frac{|0\rangle + |1\rangle}{\sqrt{2}} \otimes |\psi\rangle \mapsto \frac{|0\rangle + e^{2\pi i\varphi}|1\rangle}{\sqrt{2}} \otimes |\psi\rangle$$

The eigenstate $|\psi\rangle$ in the target register emerges unchanged

$\Rightarrow$ It suffices to focus on the control register

The state $|0\rangle + |1\rangle$ of the control qubit is changed to $|0\rangle + e^{2\pi i\varphi}|1\rangle$

Hadamard test

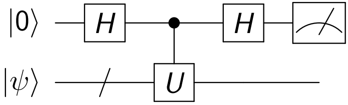

$$
\begin{aligned}
& \frac{|0\rangle + e^{2\pi i\varphi}|1\rangle}{\sqrt{2}} \\
\mapsto \quad & \frac{1}{2}\big((|0\rangle + |1\rangle) + e^{2\pi i\varphi}(|0\rangle - |1\rangle)\big) \\
= \quad & \frac{1}{2}\big((1 + e^{2\pi i\varphi})|0\rangle + (1 - e^{2\pi i\varphi})|1\rangle)\big)
\end{aligned}
$$

Hadamard test

$$\frac{1}{2}\left((1+e^{2\pi i\varphi})|0\rangle + (1-e^{2\pi i\varphi})|1\rangle)\right)$$

The probability of obtaining 0 is

$$\begin{aligned}
\Pr(0) &= |\langle 0|\varphi\rangle|^2 \\
&= |\frac{1}{2}(1+e^{2\pi i\varphi})|^2 \\
&= \frac{1}{4}|e^{\pi i\varphi}+e^{-\pi i\varphi}|^2 \\
&= \frac{1}{4}|2\cos(\pi\varphi)|^2 \\
&= \cos^2(\pi\varphi)
\end{aligned}$$

Phase kickback due to higher powers of U

For arbitrary k, we obtain

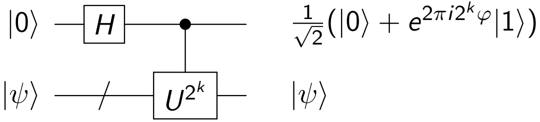

since

$$U^{2^k}|\psi\rangle = e^{2\pi i 2^k \varphi}|\psi\rangle$$

Phase kickback part of phase estimation

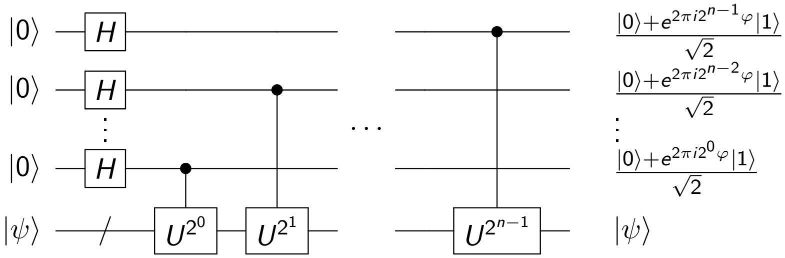

We set

$$|\varphi\rangle := \frac{|0\rangle + e^{2\pi i 2^{n-1}\varphi}|1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle + e^{2\pi i 2^{n-2}\varphi}|1\rangle}{\sqrt{2}} \otimes \cdots \otimes \frac{|0\rangle + e^{2\pi i 2^{0}\varphi}|1\rangle}{\sqrt{2}}$$

Binary fractions

Assume that the eigenphase φ is an exact n-bit binary fraction, i.e.,

$$\varphi = 0.x_1 x_2 \dots x_n = \sum_{i=1}^{n} \frac{x_i}{2^i}$$

For $k \in \{0, \dots, n-1\}$, we have

$$2^k \varphi = x_1 x_2 \dots x_k . x_{k+1} \dots x_n$$

$$\begin{aligned}
e^{2\pi i 2^k \varphi} &= e^{2\pi i (x_1 x_2 \dots x_k . x_{k+1} \dots x_n)} \\
&= e^{2\pi i (x_1 x_2 \dots x_k + 0.x_{k+1} \dots x_n)} \\
&= e^{2\pi i (x_1 x_2 \dots x_k)} \cdot e^{2\pi i (0.x_{k+1} \dots x_n)} \\
&= e^{2\pi i (0.x_{k+1} \dots x_n)}
\end{aligned}$$

Phase kickback part of phase estimation

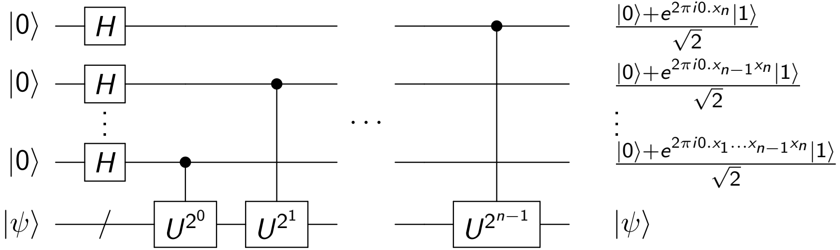

Quantum Fourier transform

The quantum Fourier transform F is defined by

$$F\big(|x_n\rangle \otimes |x_{n-1}\rangle \otimes \cdots \otimes |x_1\rangle\big)$$

$$= \quad \frac{|0\rangle + e^{2\pi i 0.x_n}|1\rangle}{\sqrt{2}} \otimes \frac{|0\rangle + e^{2\pi i 0.x_{n-1}x_n}|1\rangle}{\sqrt{2}} \otimes \cdots \otimes \frac{|0\rangle + e^{2\pi i 0.x_1 x_2 \dots x_n}|1\rangle}{\sqrt{2}}$$

Use inverse quantum Fourier transform $F^\dagger$ to obtain the bits of the eigenphase

Quantum circuit for phase estimation

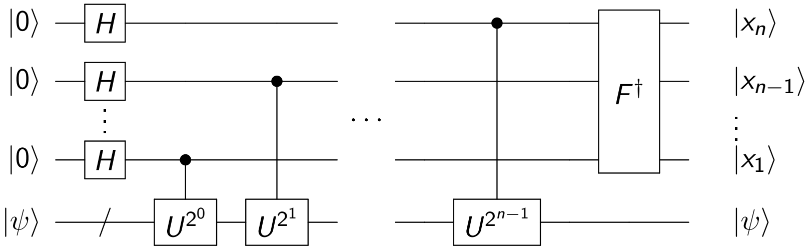

Inverse quantum Fourier transform for 3 bits

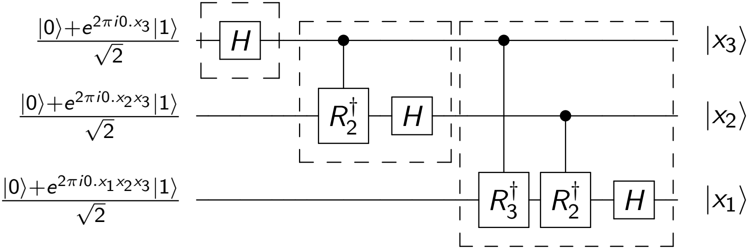

The phase shift R_k is defined by

$$R_k := \begin{bmatrix} 1 & 0 \\ 0 & e^{2\pi i/2^k} \end{bmatrix}$$

Summary of phase estimation circuit

We use phase kick back due to the controlled U^{2^k} gate to prepare the state

$$\frac{|0\rangle + e^{2\pi i 0.x_{k+1}x_{k+2}\ldots x_n}|1\rangle}{\sqrt{2}}$$

Using the previously determined bits $x_{k+2}, \ldots, x_n$, we change this state to

$$\frac{|0\rangle + e^{2\pi i 0.x_{k+1}0\ldots 0}|1\rangle}{\sqrt{2}} = \frac{|0\rangle + (-1)^{x_k}|1\rangle}{\sqrt{2}}$$

We apply the Hadamard gate to obtain

$$|x_{k+1}\rangle$$

The controlled phase shifts enable us to reduce the problem of determining each bit to distinguishing between $|+\rangle$ and $|-\rangle$ (deterministic Hadamard test)

Special case: exact n-bit binary fraction

Assume that φ is an exact n-bit binary fraction, i.e., $\varphi = 0.x_1 \ldots x_{n-1}x_n$

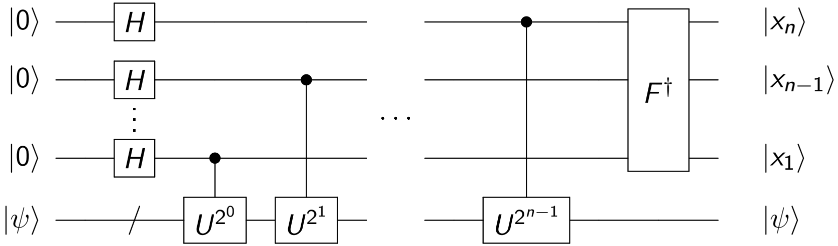

$\Rightarrow$ The measurment of the qubits yields the bits $x_n, x_{n-1}, \ldots, x_1$ deterministically

General case: arbitrary eigenphases

Let φ be arbitrary

Unless φ is an exact n-bit fraction, the application of the inverse quantum Fourier transform

$$F^{\dagger}|\varphi\rangle$$

produces a superposition of n-bit strings

Probability of obtaining a certain estimate

Lemma

Let $x = \sum_{k=1}^{n} x_i 2^{n-i}$ *and* $\varphi_x := 0.x_1 x_2 \ldots x_n = \frac{x}{2^n}$ *be the corresponding n-bit fraction*

The probability of obtaining the estimate φ_x *when the true eigenphase is* φ *is*

$$\Pr(x) = \frac{1}{2^{2n}} \frac{\sin^2\left(2^n \pi \left(\varphi - \varphi_x\right)\right)}{\sin^2\left(\pi \left(\varphi - \varphi_x\right)\right)}$$

This distribution is peaked around the true value

Examples of probability distributions for different φ

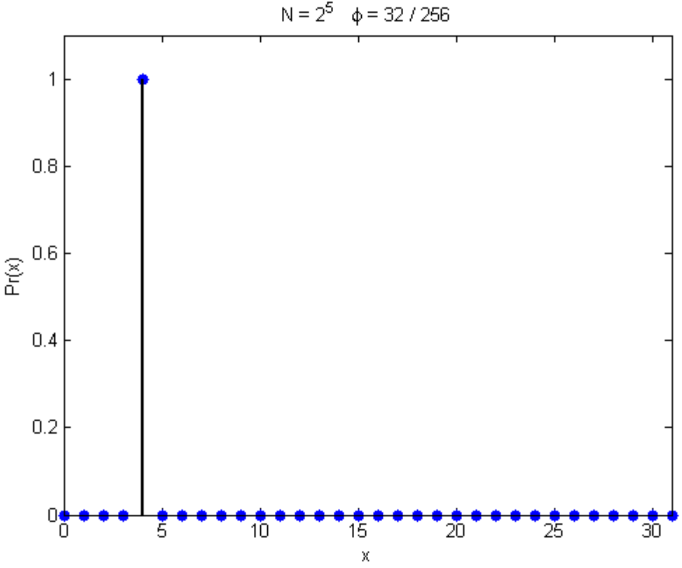

Examples of probability distributions for different φ

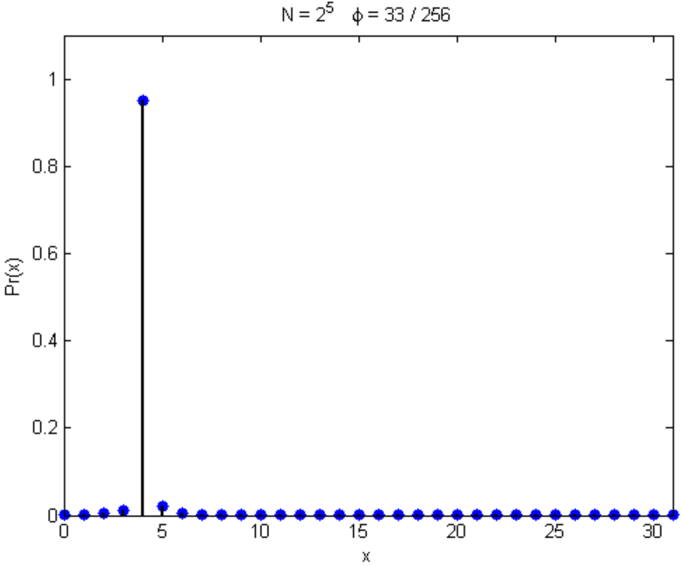

Examples of probability distributions for different φ

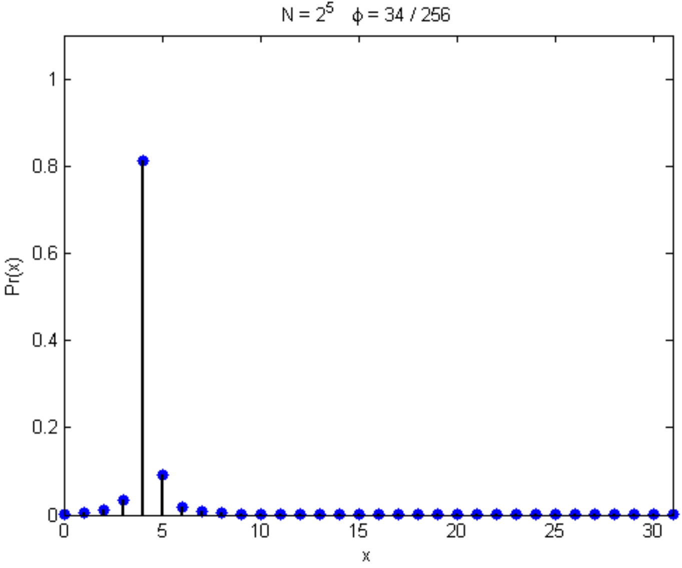

Examples of probability distributions for different φ

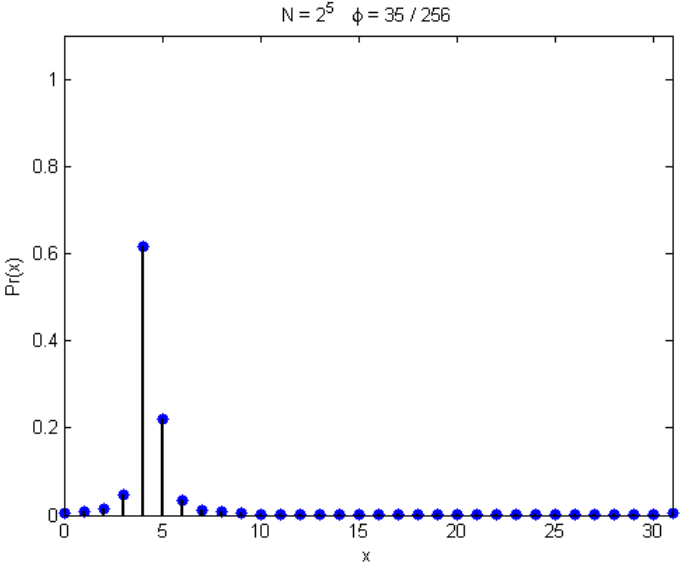

Examples of probability distributions for different φ

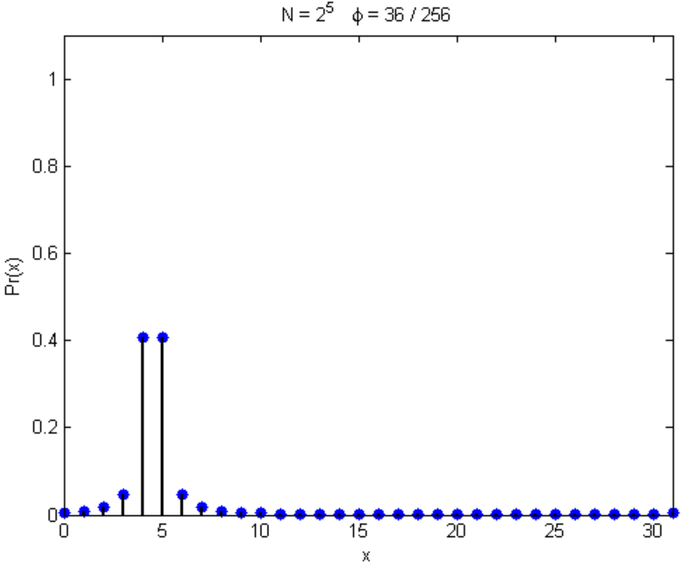

Examples of probability distributions for different φ

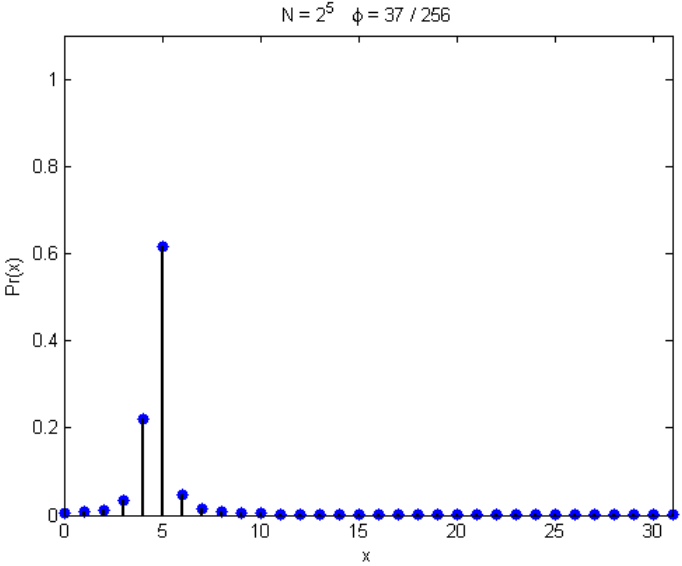

Examples of probability distributions for different φ

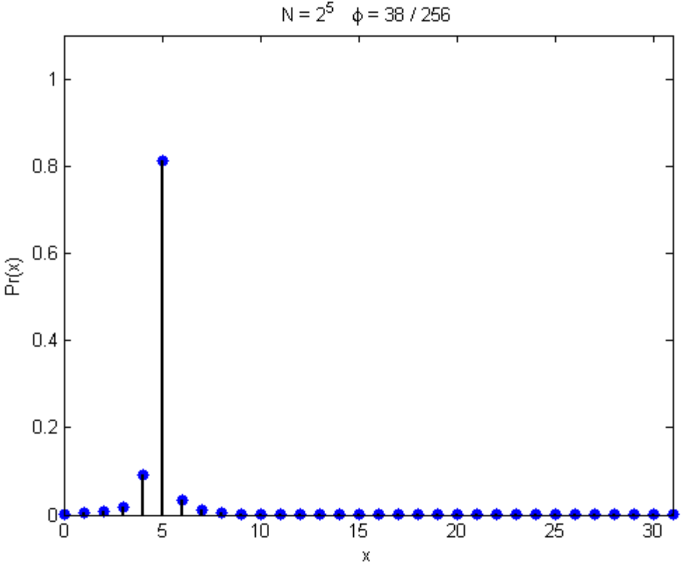

Examples of probability distributions for different φ

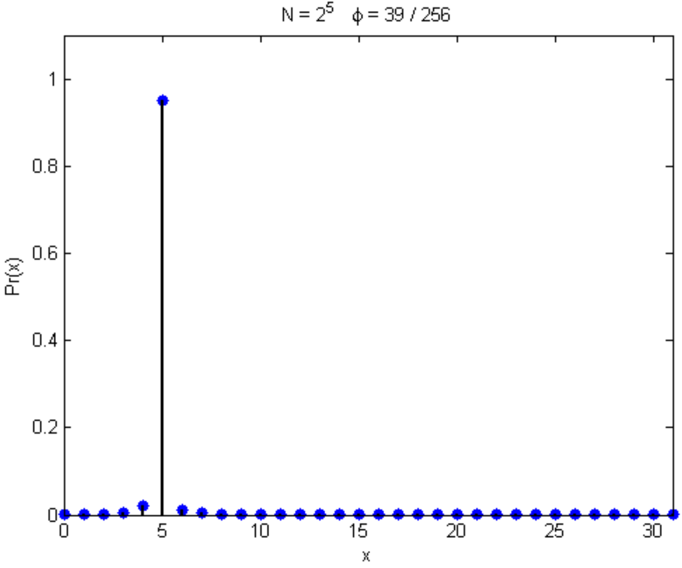

Examples of probability distributions for different φ

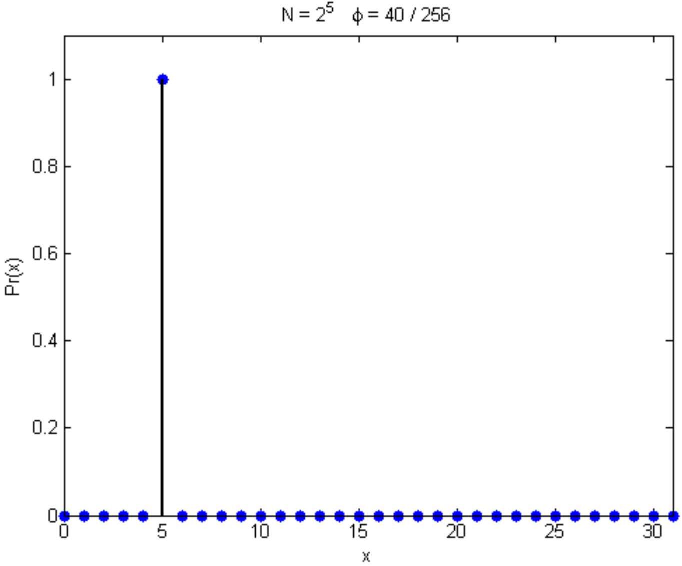

Lower bound on success probability

Theorem

Let x be such that $\frac{x}{2^n} \leq \varphi < \frac{x+1}{2^n}$

The probability of returning one of the two closest n-bit fractions φ_x and φ_{x+1} is at least $\frac{8}{\pi^2}$

Summary of phase estimation

We are given a unitary U and an eigenvector $|\psi\rangle$ of U with unknown eigenphase φ

We obtain an estimate $\hat{\varphi}$ such that

$$\Pr\left(|\hat{\varphi} - \varphi| \leq \frac{1}{2^n}\right) \geq \frac{8}{\pi^2}$$

To do this, we need invoke each of the controlled U, $U^2, \ldots, U^{2^{n-1}}$ gates once

We can boost the success probability to $1 - \epsilon$ by repeating the above algorithm $O(\log(1/\epsilon))$ times and outputting the median of the outcomes

Phase estimation applied to superpositions of eigenstates

We are given a unitary U with eigenvectors $|\psi_i\rangle$ and corresponding eigenphases φ_i

Let

$$|\psi\rangle = \sum_i \alpha_i |\psi_i\rangle$$

What happens if we apply phase estimation to $|0\rangle^{\otimes n} \otimes |\psi\rangle$?

After the n phase kickbacks due to U^{2^0}, U^{2^1}, $\ldots$ $U^{2^{n-1}}$, we obtain

$$\sum_i \alpha_i |\varphi_i\rangle \otimes |\psi_i\rangle$$

After applying the inverse quantum Fourier transform, we obtain

$$\sum_i \alpha_i |\tilde{x}_i\rangle \otimes |\psi_i\rangle$$

where $|\tilde{x}_i\rangle$ denotes a superpositions of n-bit estimates of φ_i

Part IV

Factoring

The fundamental theorem of arithmetic

Theorem

Every positive integer larger than 1 can be factored as a product of prime numbers, and this factorization is unique (up to the order of the factors).

$$N = 2^{n_2} \times 3^{n_3} \times 5^{n_5} \times 7^{n_7} \times \cdots$$

Examples

$$15 = 3 \times 5$$

$$239815173914273 = 15485863 \times 15486071$$

$$\begin{array}{l} 3107418240490043721350750 \\ 0358885679300373460228427 \\ 2754572016194882320644051 \\ 8081504556346829671723286 \\ 7824379162728380334154710 \\ 7310850191954852900733772 \\ 4822783525742386454014691 \\ 736602477652346609 \end{array} = \begin{array}{c} 1634733645809253848 4 \\ 4313388386509085984 1 \\ 7836700330923121811 1 \\ 0852389333100104508 1 \\ 5121211816751157 9 \\ \times \\ 1900871281664822113 1 \\ 2685157393541397547 1 \\ 8967899685154936666 3 \\ 8539088027103802104 4 \\ 9895719126146557 1 \end{array}$$

Why care about factoring?

"The problem of distinguishing prime numbers from composite numbers and of resolving the latter into their prime factors is known to be one of the most important and useful in arithmetic. It has engaged the industry and wisdom of ancient and modern geometers to such an extent that it would be superfluous to discuss the problem at length... Further, the dignity of the science itself seems to require that every possible means be explored for the solution of a problem so elegant and so celebrated."

– Carl Friedrich Gauss, Disquisitiones Arithmeticæ (1801)

More practically: The presumed hardness of factoring is the basis of much of modern cryptography (RSA cryptosystem)

Order finding

Definition

Given $a, N \in \mathbb{Z}$ with $\gcd(a, N) = 1$, the *order* of a modulo N is the smallest positive integer r such that $a^r \equiv 1 \pmod{N}$.

Problem

- Given: $a, N \in \mathbb{Z}$ with $\gcd(a, N) = 1$
- Task: find the order of a modulo N

Spectrum of a cyclic shift

Let P be a cyclic shift modulo r: $P|x\rangle = |x + 1 \bmod r\rangle$

Claim. For any $k \in \mathbb{Z}$, the state $|u_k\rangle := \frac{1}{\sqrt{r}} \sum_{x=0}^{r-1} e^{-2\pi i k x / r} |x\rangle$ is an eigenstate of P.

Proof.
$$
\begin{aligned}
U|u_k\rangle &= \frac{1}{\sqrt{r}} \sum_{x=0}^{r-1} e^{-2\pi i k x / r} |x + 1 \bmod r\rangle \\
&= \frac{1}{\sqrt{r}} \sum_{x=0}^{r-1} e^{2\pi i k / r} e^{-2\pi i k (x+1) / r} |x + 1 \bmod r\rangle \\
&= e^{2\pi i k / r} \frac{1}{\sqrt{r}} \sum_{x=1}^{r} e^{-2\pi i k x / r} |x \bmod r\rangle \\
&= e^{2\pi i k / r} |u_k\rangle \qquad \square
\end{aligned}
$$

The multiplication-by-a map

Define U by $U|x\rangle = |ax\rangle$ for $x \in \mathbb{Z}_N$.

Computing U:

$$\begin{aligned} |x, 0\rangle &\mapsto |x, ax\rangle && \text{(reversible multiplication by } a\text{)} \\ &\mapsto |ax, x\rangle && \text{(swap)} \\ &\mapsto |ax, 0\rangle && \text{(uncompute reversible division by } a\text{)} \end{aligned}$$

High powers of U can be implemented efficiently using repeated squaring

Spectrum of the multiplication-by-a map

Define U by $U|x\rangle = |ax\rangle$ for $x \in \mathbb{Z}_N$.

Claim. Let r be the order of a modulo N. For any $k \in \mathbb{Z}$, the state

$$|u_k\rangle := \frac{1}{\sqrt{r}} \sum_{x=0}^{r-1} e^{-2\pi i k x / r} |a^x \bmod N\rangle$$

is an eigenstate of U with eigenvalue $e^{2\pi i k/r}$.

Proof.
Same as for the cyclic shift, due to the isomorphism

$$x \bmod r \quad \leftrightarrow \quad a^x \bmod N$$

□

Order finding and phase estimation

$U|u_k\rangle = e^{2\pi i k/r}|u_k\rangle$

Phase estimation of U on $|u_k\rangle$ can be used to approximate k/r.

Problems:

1. We don't know r, so we can't prepare $|u_k\rangle$.
2. We only get an approximation of k/r.
3. Even if we knew k/r exactly, k and r could have common factors.

Solutions:

1. Estimate k/r for a superposition of the $|u_k\rangle$.
2. Use the continued fraction expansion.
3. Show that $gcd(k, r) = 1$ with reasonable probability.

Estimating k/r in superposition

A useful identity:

$$\sum_{k=0}^{r-1} e^{2\pi \mathrm{i} kx/r} = \begin{cases} r & \text{if } x = 0 \\ 0 & \text{otherwise} \end{cases}$$

Consider

$$\frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} |u_k\rangle = \frac{1}{r} \sum_{k,x=0}^{r-1} e^{-2\pi \mathrm{i} kx/r} |a^x \bmod N\rangle$$
$$= |a^0 \bmod N\rangle = |1\rangle$$

Phase estimation:

$$|0\rangle \otimes |1\rangle = \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} |0\rangle \otimes |u_k\rangle \mapsto \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} |\widetilde{k/r}\rangle \otimes |u_k\rangle$$

Measurement gives an approximation of k/r for a random k

Continued fractions

Problem

Given samples x of the form $\lfloor k\frac{2^n}{r}\rfloor$, $\lceil k\frac{2^n}{r}\rceil$ ($k \in \{0, 1, \ldots, r-1\}$), determine r.

Continued fraction expansion:

$$\frac{x}{2^n} = \frac{1}{a_1 + \frac{1}{a_2 + \frac{1}{a_3 + \cdots}}}$$

Gives an efficiently computable sequence of rational approximations

Theorem

If $2^n \geq N^2$, then k/r is the closest convergent of the CFE to $x/2^n$ among those with denominator smaller than N.

Since $r < N$, it suffices to take $n = 2\log_2 N$

Common factors

If $\gcd(k, r) = 1$, then the denominator of k/r is r

Fact

The probability that $\gcd(k, r) = 1$ *for a random* $k \in \{0, 1, \ldots, r-1\}$ *is*

$$\frac{\phi(r)}{r} = \Omega\left(\frac{1}{\log \log r}\right)$$

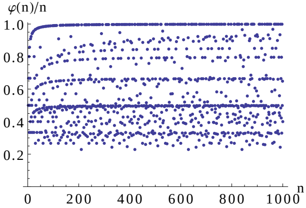

Thus $\Omega(\log \log N)$ repetitions suffice to give r with constant probability

Alternatively, find two (or more) denominators and take their least common multiple; then $O(1)$ repetitions suffice

Factoring $\rightarrow$ finding a nontrivial factor

Suppose we want to factor the positive integer N.

Since primality can be tested efficiently, it suffices to give a procedure for finding a nontrivial factor of N with constant probability.

```
function factor(N)

if N is prime
     output N
else
     repeat
          x=find_nontrivial_factor(N)
     until success
     factor(x)
     factor(N/x)
end if
```

We can assume N is odd, since it is easy to find the factor 2.

We can also assume that N contains at least two distinct prime powers, since it is easy to check if it is a power of some integer.

Reduction of factoring to order finding

Factoring N reduces to order finding in $\mathbb{Z}_N^\times$ [Miller 1976].

Choose $a \in \{2, 3, \ldots, N-1\}$ uniformly at random.

If $\gcd(a, N) \neq 1$, then it is a nontrivial factor of N.

If $\gcd(a, N) = 1$, let r denote the order of a modulo N.

Suppose r is even. Then

$$a^r = 1 \bmod N$$
$$\Updownarrow$$
$$(a^{r/2})^2 - 1 = 0 \bmod N$$
$$\Updownarrow$$
$$(a^{r/2} - 1)(a^{r/2} + 1) = 0 \bmod N$$

so we might hope that $\gcd(a^{r/2} - 1, N)$ is a nontrivial factor of N.

Miller's reduction

Question

Given $(a^{r/2}-1)(a^{r/2}+1) = 0 \bmod N$, when does $\gcd(a^{r/2}-1, N)$ give a nontrivial factor of N?

Note that $a^{r/2} - 1 \neq 0 \bmod N$ (otherwise the order of a would be $r/2$, or smaller).

So it suffices to ensure that $a^{r/2} + 1 \neq 0 \bmod N$.

Lemma

Suppose $a \in \mathbb{Z}_N^\times$ is chosen uniformly at random, where N is an odd integer with at least two distinct prime factors. Then with probability at least $1/2$, the order r of a is even and $a^{r/2} \neq -1 \bmod N$.

Shor's algorithm

Input: Integer N
Output: A nontrivial factor of N

1. Choose a random $a \in \{2, 3, \ldots, N-1\}$
2. Compute $\gcd(a, N)$; if it is not 1 then it is a nontrivial factor, and otherwise we continue
3. Perform phase estimation with the multiplication-by-a operator U on the state $|1\rangle$ using $n = 2\log_2 N$ bits of precision
4. Compute the continued fraction expansion of the estimated phase, and find the best approximation with denominator less than N; call the result r
5. Compute $\gcd(a^{r/2} - 1, N)$. If it is a nontrivial factor of N, we are done; if not, go back to step 1

Quantum vs. classical factoring algorithms

Best known classical algorithm for factoring N

- Proven running time: $2^{O((\log N)^{1/2}(\log\log N)^{1/2})}$
- With plausible heuristic assumptions: $2^{O((\log N)^{1/3}(\log\log N)^{1/3})}$

Shor's quantum algorithm

- QFT modulo 2^n with $n = O(\log N)$: takes $O(n^2)$ steps
- Modular exponentiation: compute a^x for $x < 2^n$. With repeated squaring, takes $O(n^3)$ steps
- Running time of Shor's algorithm: $O(\log^3 N)$

Part V

Quantum search

Unstructured search

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

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

We are given a black box $f\colon X \to \{0,1\}$ 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"

Unstructured search

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

On the other hand, a quantum algorithm can do much better!

Phase oracle

We assume that we have a unitary operator U satisfying

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

By phase kickback, this can easily be implemented using a reversible black fox for f

Target state

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

This is essentially the hardest case

Goal: prepare the state $|m\rangle$

Initial state

We have no information about which item might be marked

$\Rightarrow$ 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 $|\psi\rangle$

Prepare the target state $|m\rangle$ by implementing a rotation that moves $|\psi\rangle$ toward $|m\rangle$

Realize the rotation as a product of two reflections

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

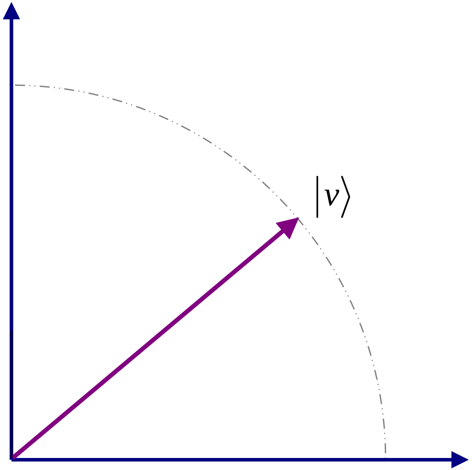

Visualization of a reflection in $\mathbb{R}^2$

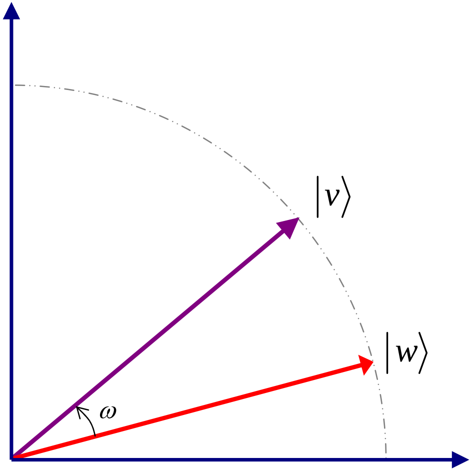

Visualization of a reflection in $\mathbb{R}^2$

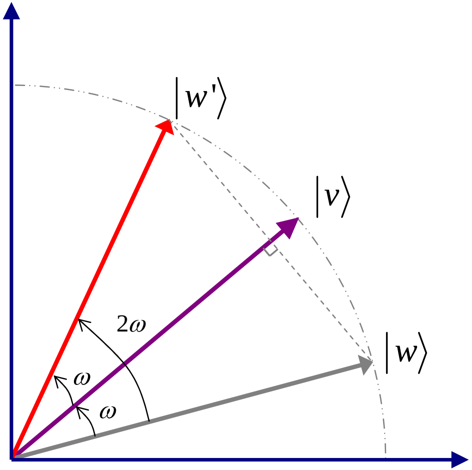

Reflections

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

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

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

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

Structure of Grover

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

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

The Gram-Schmidt process yields

$$|\phi\rangle = \frac{|\psi\rangle - \alpha|m\rangle}{\sqrt{1-\alpha^2}}$$

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

Invariant subspace

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

$$
\begin{aligned}
|\psi\rangle &= \sin\theta|m\rangle + \cos\theta|\phi\rangle \text{ where } \sin\theta = \langle m|\psi\rangle = 1/\sqrt{N} \\
U &= \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \\
V &= I - 2|\psi\rangle\langle\psi| \\
&= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - 2\begin{pmatrix} \sin\theta \\ \cos\theta \end{pmatrix}\begin{pmatrix} \sin\theta & \cos\theta \end{pmatrix} \\
&= \begin{pmatrix} 1 - 2\sin^2\theta & -2\sin\theta\cos\theta \\ -2\sin\theta\cos\theta & 1 - 2\cos^2\theta \end{pmatrix} \\
&= -\begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix}
\end{aligned}
$$

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

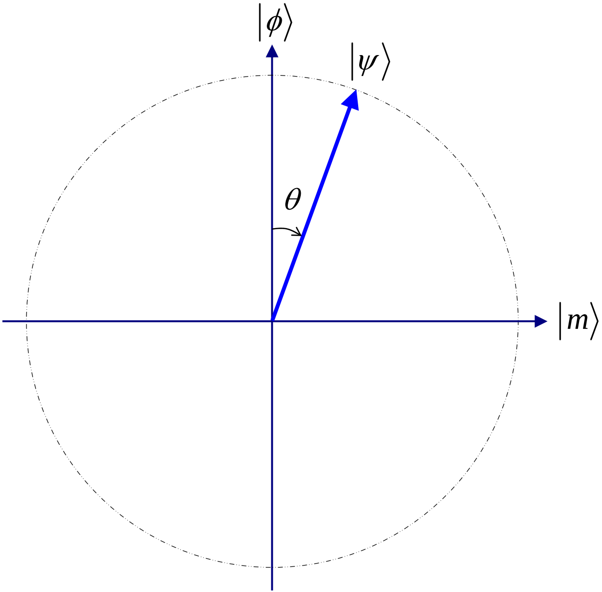

Visualization of first Grover iteration

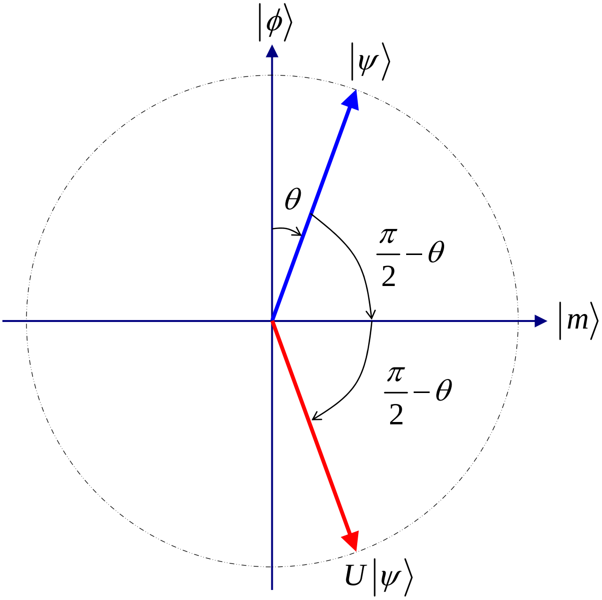

Visualization of first Grover iteration

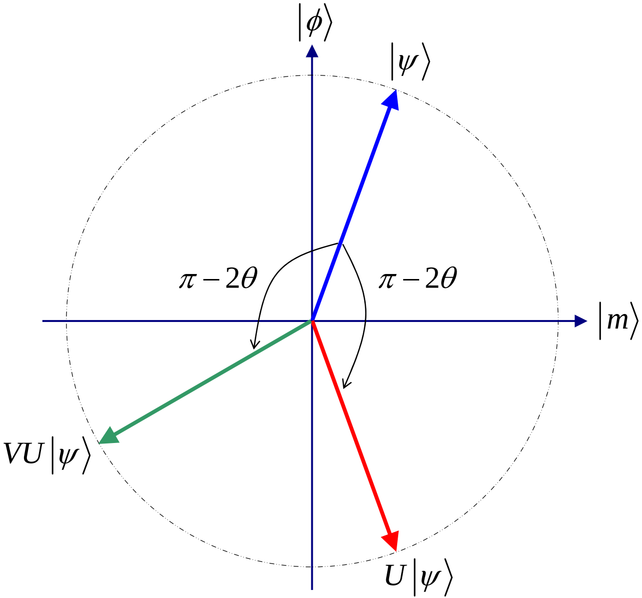

Visualization of first Grover iteration

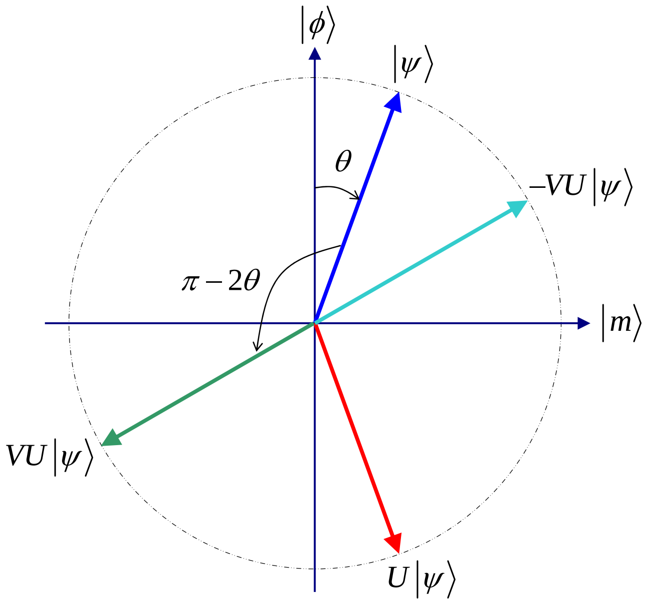

Visualization of first Grover iteration

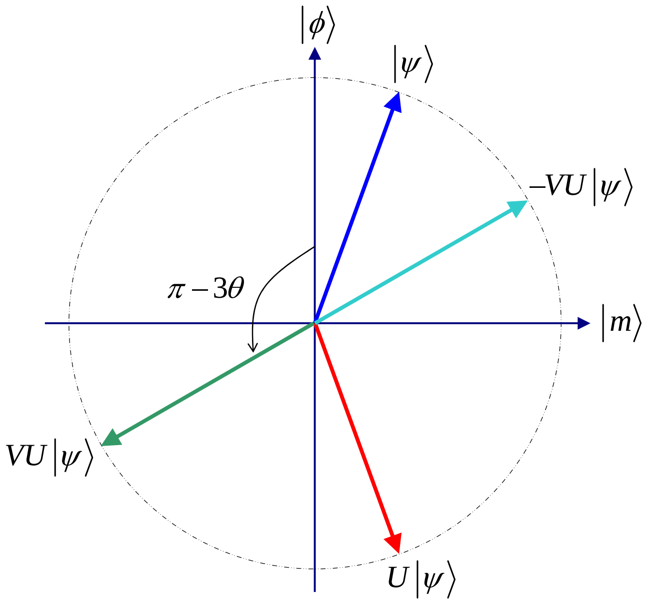

Visualization of first Grover iteration

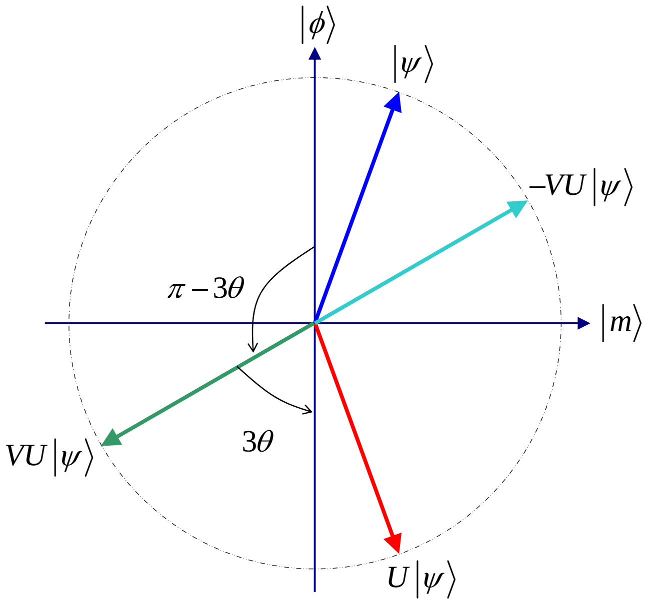

Visualization of first Grover iteration

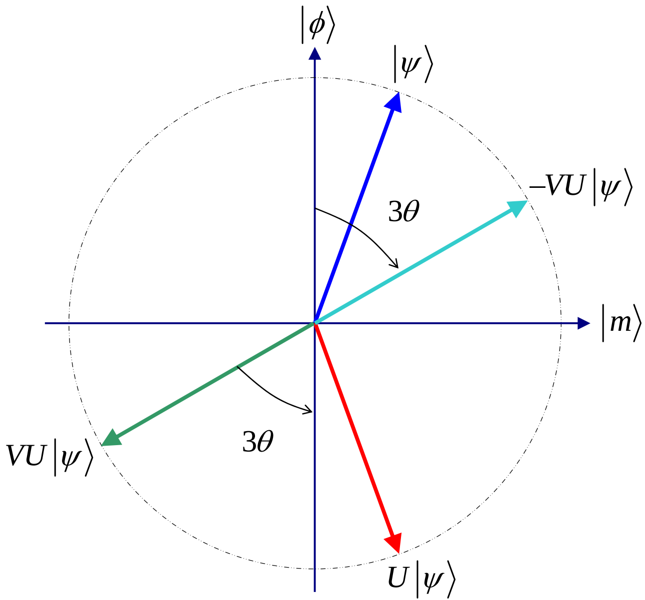

Visualization of first Grover iteration

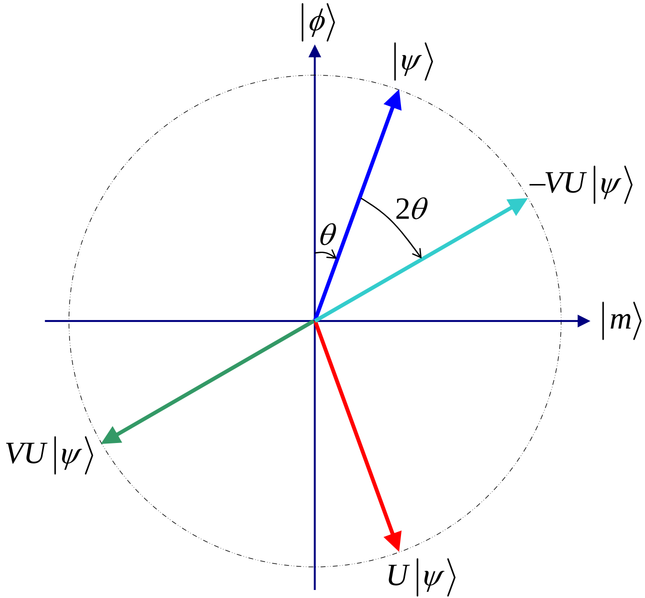

Grover search

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

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

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

$\Rightarrow$ 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

It is easy to calculate that

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

This is the probability that, after k steps of the algorithm, a measurement reveals the marked state

We are solving a completely unstructured search problem with N possible solutions, yet we can find 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

Grover search

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)

Part VI

Beyond Shor and Grover

Beyond factoring

There are many fast quantum algorithms based on ideas related to Shor's factoring algorithm:

- Computing discrete logarithms
- Decomposing abelian/solvable groups
- Estimating Gauss sums
- Counting points on algebraic curves
- Computations in number fields (Pell's equation, etc.)
- Abelian hidden subgroup problem
- Non-abelian hidden subgroup problem?

Beyond unstructured search

- Quantum counting
- Amplitude amplification
- Applications
  - Collision problem
  - Finding the median, minimum, etc.
  - Graph problems (spanning trees, matchings, flows, ...)
  - And many more ...

Quantum walk

Quantum analogs of random walks have led to several new quantum algorithms:

- Exponential speedup for a black-box problem
- General framework for search on graphs
  - Spatial search
  - Element distinctness
  - Subgraph finding (e.g., triangle problem)
  - Checking matrix multiplication
  - Checking group commutativity
- Evaluating Boolean formulas

Other quantum algorithms

- Oracle interrogation
- Simulating Hamiltonian dynamics
- Gradient estimation
- Approximation of #P-hard problems
  - Jones polynomial
  - Tutte polynomial
  - Manifold invariants
  - Tensor networks
- Linear systems, differential equations

Further reading

- P. Kaye, R. Laflamme, and M. Mosca, *An Introduction to Quantum Computing* (Oxford University Press, 2007)
- M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000)
- M. Mosca, *Quantum Algorithms*, in Encyclopedia of Complexity and Systems Science (Springer, 2009), arXiv:0808.0369
- A. M. Childs and W. van Dam, *Quantum algorithms for algebraic problems*, Reviews of Modern Physics 82, 1–52 (2010), arXiv:0812.0380