

ADVANCED QUANTUM INFORMATION THEORY

Lecture notes

Ashley Montanaro, University of Bristol
ashley.montanaro@bristol.ac.uk

Contents

1	Introduction	2
2	Classical and quantum computational complexity	4
3	Grover's algorithm	10
4	The Quantum Fourier Transform and periodicity	17
5	Integer factorisation	22
6	Phase estimation	28
7	Hamiltonian simulation	31
8	Quantum walk	35
9	Noise and the framework of quantum channels	45
10	Quantum error-correction	50
11	Quantum state discrimination and tomography	57

Health warning: There are likely to be changes and corrections to these notes throughout the unit. For updates, see <http://www.maths.bris.ac.uk/~csxam/aqit2016/>.

Version 1.5 (February 26, 2016).

1 Introduction

The field of quantum information theory studies the remarkable ways in which quantum information – and the processing thereof – differs from information stored in classical systems. Nowhere is this difference more pronounced than the dramatic speedups obtained by quantum computation over classical computation. These notes aim to cover (some of) the theoretical topics which any self-respecting quantum information theorist, or experimentalist working in the field of quantum information processing, should know. These include the famous algorithms of Shor and Grover, and the simulation of quantum systems; the more recent topic of quantum walk algorithms; decoherence and quantum error-correction.

1.1 Complementary reading

These lecture notes have benefited significantly from the expositions in the following lecture courses, which may be of interest for background reading:

- *Quantum Computation*, Richard Jozsa, University of Cambridge
<http://www.qi.damtp.cam.ac.uk/node/261>
The material here on the QFT and Shor's algorithm follows this exposition closely.
- *Quantum Algorithms*, Andrew Childs, University of Maryland
<http://cs.umd.edu/~amchilds/qa/>
A combination of lecture notes for three graduate courses. The material here on quantum walk algorithms is based on the exposition in these notes.
- *Theory of Quantum Information*, John Watrous, University of Waterloo
<https://cs.uwaterloo.ca/~watrous/LectureNotes.html>
A particularly useful resource for the theory of quantum channels.

The following books and survey papers may also be useful:

- *Quantum Computation and Quantum Information*, Nielsen and Chuang
Cambridge University Press, 2001
The Bible of quantum computing.
- *Classical and Quantum Computation*, Kitaev, Shen and Vyalı
American Mathematical Society, 2002
A more concise introduction to many important topics in quantum computation.
- *Quantum algorithms for algebraic problems*, Childs and van Dam
Reviews of Modern Physics, 82:1, 2010; <http://arxiv.org/pdf/0812.0380.pdf>
Covers many other quantum algorithms than those discussed here.

1.2 Notation

We write $[n] := \{1, \dots, n\}$ for the integers between 1 and n . $\lceil x \rceil$, $\lfloor x \rfloor$ and $\llbracket x \rrbracket$ denote the smallest integer y such that $y \geq x$, the largest integer z such that $z \leq x$, and the closest integer to x , respectively. We use $\binom{n}{k}$ for the binomial coefficient “ n choose k ”, $n!/(k!(n-k)!)$. Finally, when we say “bounded-error”, we mean with failure probability upper-bounded by some constant below $1/2$.

We use standard “computer science style” notation relating to asymptotic complexity:

- $f(n) = O(g(n))$ if there exist real $c > 0$ and integer $n_0 \geq 0$ such that for all $n \geq n_0$, $f(n) \leq c g(n)$.
- $f(n) = \Omega(g(n))$ if there exist real $c > 0$ and integer $n_0 \geq 0$ such that for all $n \geq n_0$, $f(n) \geq c g(n)$. Clearly, $f(n) = O(g(n))$ if and only if $g(n) = \Omega(f(n))$.
- $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$.

O , Ω and Θ can be viewed as asymptotic, approximate versions of \leq , \geq and $=$.

Box 1: Big-O notation

1.3 Change log

- v1.0: first version covering first part of the 2016 unit.
- v1.1: addition of material on the QFT, periodicity and Shor’s algorithm.
- v1.2: addition of material on phase estimation and Hamiltonian simulation.
- v1.3: addition of material covering the remainder of the unit.
- v1.4: update tomography section to mention recent results.
- v1.5: correct typos in stabilizer quantum error-correction section.

2 Classical and quantum computational complexity

Computational complexity theory aims to classify different problems in terms of their difficulty, or in other words the resources required in order to solve them. Two of the most important types of resources one might study are *time* (the number of computational steps used by an algorithm solving a problem) and *space* (the amount of additional work space used by the algorithm). Classically, the formal model underpinning the notion of an “algorithm” is the Turing machine. We will not go into details about this here, instead taking the informal approach that the number of steps used by an algorithm corresponds to the number of lines of code executed when running the algorithm, and the space usage is the amount of memory used by the algorithm. For much more on the topic, see the book *Computational Complexity* by Papadimitriou, for example.

Rather than looking at the complexity of algorithms for solving one particular instance of a problem, the theory considers asymptotics: given a family of problems, parametrised by an instance size (usually denoted n), we study the resources used by the best possible algorithm for solving that family of problems. The term “problem” is used henceforth as shorthand for “family of problems”. A dividing line between efficient and inefficient algorithms is provided by the notion of polynomial-time computation. An algorithm running in time polynomial in n , i.e. $O(n^c)$ for some fixed c , is considered efficient. For example, consider the following two problems:

- Primality testing: given an integer N expressed as n binary digits, is it a prime number?
- Factorisation: given an integer N expressed as n binary digits, output the prime factors of N .

As the input is of size n , we would like to solve these problems using an algorithm which runs in time $\text{poly}(n)$ (not $\text{poly}(N)!$). No such classical algorithm is known for the factorisation problem; as we will see later, the situation is different for quantum algorithms. However, surprisingly, there is a polynomial-time classical algorithm for the tantalisingly similar problem of primality testing.

An important class of problems is known as decision problems; these are problems that have a yes-no answer. The first of the above problems is a decision problem, while the second is not. But it can be made into a decision problem without changing its underlying complexity significantly:

- Factorisation (decision variant): given integers N and K expressed as n binary digits each, does N have a prime factor smaller than K ?

It is clear that, if we can solve the usual “search” variant of the factorisation problem, solving the decision variant is easy. Further, solving the decision variant allows us to solve the search variant of the problem using binary search. Given an integer N whose prime factors we would like to determine, and an algorithm which solves the decision variant of the factorisation problem, we can use $O(\log N) = O(n)$ evaluations of this algorithm with different values of K to find the smallest prime factor M of N . The other factors can be found by dividing N by M and repeating. This is a simple example of a *reduction*: conversion of one problem into another.

A natural way to compare the complexity of problems is via the notion of complexity classes, where a complexity class is simply a set of problems. Some important classical complexity classes are:

- P: the class of decision problems which can be solved in polynomial time by a classical computer.

- NP: the class of decision problems such that, if the answer is “yes”, there is a proof of this fact which can be verified in polynomial time by a classical computer.
- PSPACE: the class of decision problems which can be solved in polynomial space by a classical computer.

Primality testing is in P, although this was shown for the first time only in 2002. The decision variant of factorisation is in NP, because given a claimed prime factor of N smaller than K , it can be easily checked whether the claim is correct. However, factorisation is not known to be in P. Every problem in P is automatically in NP, because the verifier can simply ignore any claimed proof and solve the problem directly. In addition, any problem in NP is automatically in PSPACE, because one can loop over all polynomial-length proofs in polynomial space in order to determine whether the answer to a problem instance should be “yes” or “no”. Thus we have $P \subseteq NP \subseteq PSPACE$.

A problem is said to be NP-complete if it is in NP, and every other problem in NP reduces to it in polynomial time. So NP-complete problems are, informally, the “hardest” problems in NP. These include many problems of practical importance in areas such as timetabling, resource allocation, and optimisation. One simple example is the Subset Sum problem. An instance of this problem is a sequence of integers x_1, \dots, x_n ; our task, given such a sequence, is to determine whether there is a subset of the integers which sums to 0. Given such a subset, we can easily check that it sums to 0; however, finding such a subset seems to require checking exponentially many subsets.

NP stands for “nondeterministic polynomial-time”, *not* “non-polynomial time”. In fact, it is currently unknown whether every problem in NP can be solved in polynomial time. This is the famous P vs. NP question; resolving it would win you a prize of \$1M from the Clay Mathematics Institute.


2.1 Quantum computational complexity

How are we to measure resource usage by a *quantum* algorithm running on a quantum computer? One framework within which to do this is the quantum circuit model. A quantum computation running for T steps and using space S corresponds to a unitary operation on S qubits (i.e. operating on \mathbb{C}^{2^S}) expressed as a product of T elementary operations picked from some family \mathcal{F} . Each elementary operation is assumed to take time $O(1)$ and act on $O(1)$ qubits. We assume that the initial state of the quantum computer is $|0\rangle^{\otimes S}$ and the computation finishes with a measurement of some of the qubits in the computational basis, which gives the output of the computation. If we prefer, we can allow intermediate measurements during the circuit; this turns out not to change the power of the model.


The set \mathcal{F} of allowed elementary operations will depend on our physical architecture. However, it turns out that most “reasonable” sets of operations on $O(1)$ qubits – called *quantum gates*, by analogy with logic gates in classical circuits – are universal, in the sense that any unitary matrix on S qubits can be decomposed as a product of these basic operations.

A quantum circuit can be drawn as follows. For convenience, in the diagram we have drawn multi-qubit gates as only acting on nearest-neighbour qubits, but this is not an essential restriction

of the model.




Beware that a circuit is read left to right, with the starting input state on the far left, but the corresponding unitary operators act right to left! For example, the circuit




corresponds to the unitary operator $(I \otimes V)(U \otimes I)(H \otimes I \otimes X)$ on 3 qubits. For any gate G , the corresponding “controlled- G ” gate CG uses an extra qubit to control whether the gate is applied or not. That is,

$$CG|0\rangle|\psi\rangle = |0\rangle|\psi\rangle, \quad CG|1\rangle|\psi\rangle = |1\rangle G|\psi\rangle.$$

In a circuit diagram, this is denoted using a filled circle on the control line:



A particularly useful such gate is controlled-NOT (CNOT), denoted . Written as a matrix with respect to the computational basis,

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

For any fixed gate set \mathcal{F} , some large unitary matrices cannot be decomposed efficiently in terms of gates from \mathcal{F} , in the sense that to write them as a product of gates from \mathcal{F} requires exponentially many such gates. For a rough way of seeing this, consider the problem of producing an arbitrary quantum state of n qubits $\sum_{x \in \{0,1\}^n} \alpha_x |x\rangle$, where each coefficient $\alpha_x \in \{\pm 1/2^{n/2}\}$. There are 2^{2^n} such states. Any circuit on n qubits made up of T gates, each acting on k qubits, picked from a gate set of size G can be described by one of

$$\left(G \binom{n}{k} \right)^T = O\left((Gn^k)^T \right) = O\left(2^{T \log(Gn^k)} \right)$$

different sequences of gates, so for $k, G = O(1)$ we need $T \log n = \Omega(2^n)$ to be able to produce $2^{\Omega(2^n)}$ different unitary operators, and hence 2^{2^n} different states. A similar argument still works if we allow approximate computation or continuous gate sets.

In general, just as in the classical world, we look for efficient quantum circuits which use $\text{poly}(n)$ qubits and $\text{poly}(n)$ gates to solve a problem on an input of size n . The class of problems which

can be solved by a quantum computer, in time polynomial in the input size, with probability of failure at most $1/3$, is known as BQP (“bounded-error quantum polynomial-time”). This class encapsulates the notion of efficient quantum computation. Observe that in the quantum circuit picture we can perform multiple operations in parallel, so we in fact have two possible ways to measure “time” complexity: circuit size (number of gates) and circuit depth (number of time steps to execute all the gates). But these can only differ by a factor of $O(S)$, where S is the number of qubits.

2.2 Classical and reversible circuits

Any classical computation which maps a bit-string to another bit-string can be broken down into a sequence of logical operations, each of which acts on a small number of bits (e.g. AND, OR and NOT gates). Such a sequence is called a (classical) circuit. We would like to show that any classical circuit can be implemented as a quantum circuit. But there is a difficulty: in quantum mechanics, if we wish the state of our system to remain pure, the evolution that we apply has to be unitary, and hence reversible. Some classical logical operations (such as AND, written \wedge) are not reversible. However, reversible variants of these can be developed using the following trick. If we wish to compute an arbitrary classical operation $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, we attach an “ancilla” register of m bits, each originally set to 0, and modify f to give a new operation $f' : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^n \times \{0, 1\}^m$ which performs the map

$$f'(x, y) = (x, y \oplus f(x)),$$

where \oplus is bitwise XOR. Then if we input $y = 0^m$, we get $(x, f(x))$, from which we can extract our desired output $f(x)$. If we perform f' twice, we get $(x, y \oplus f(x) \oplus f(x)) = (x, y)$. So f is reversible. And any reversible function that maps bit-strings to bit-strings corresponds to a permutation matrix, which is unitary, so can be implemented as a sequence of quantum gates. If we combine many gates of this form to compute a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, say, we will finish with an output of the form (junk, $x, f(x)$). If we wish to remove the junk, we can simply copy the output $f(x)$ onto a fresh ancilla bit in state 0, and then repeat all the previous gates in reverse. As each is its own inverse, the final state of the computation is $(0, x, f(x))$.

To obtain universal deterministic classical computation, it is sufficient to be able to implement the NOT and AND gates. The NOT gate is immediately reversible. Applying the above construction to AND we get the map $(x_1, x_2, y) \mapsto (x_1, x_2, y \oplus (x_1 \wedge x_2))$. The unitary operator which implements this is then simply the map

$$|x_1\rangle|x_2\rangle|y\rangle \mapsto |x_1\rangle|x_2\rangle|y \oplus (x_1 \wedge x_2)\rangle.$$

Written as a matrix with respect to the computational basis this is

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix},$$

an operation known as the Toffoli gate. In a circuit diagram, the Toffoli gate is written as “controlled-controlled-NOT”, i.e.



It is known that the Toffoli gate, together with the Hadamard gate, are even sufficient for universal *quantum* computation. Another representative universal set of quantum gates is $\{H, X, \text{CNOT}, T\}$, where $T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$.

Randomised classical computation can also be embedded in a quantum circuit. Imagine we have a classical computation which makes use of some random bits, each of which is 0 or 1 with equal probability. We can simulate this by applying a Hadamard gate to $|0\rangle$ to produce the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. Then we can either measure this qubit immediately to obtain a uniformly random bit, or if we prefer, apply classical gates to it and then measure it at the end of the computation; the result is the same.

2.3 Query complexity

While time complexity is a practically important measure of the complexity of algorithms, it suffers from the difficulty that it is very hard to prove lower bounds on it, and that technical details can sometimes obscure the key features of an algorithm. One way to sidestep this is to use a model which is less realistic, but cleaner and more mathematically tractable: the model of query complexity.

In this model, we assume we have access to an *oracle*, or “black box”, to which we can pass *queries*, and which returns answers to our queries. Our goal is to determine some property of the oracle using the minimal number of queries. On a classical computer, we can think of the oracle as a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$. We pass in inputs $x \in \{0, 1\}^n$, and receive outputs $f(x) \in \{0, 1\}^m$. How does this fit into physical reality? We imagine we are given access to the oracle either as a physical device which we cannot open and look inside, or as a circuit which we can see, but for which it might be difficult to compute some property of the circuit. For example, even given a description of a circuit computing some function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, it might be hard to find an input x such that $f(x) = 1$. Sometimes it is more natural to think of an oracle function f as a memory storing n strings of m bits each, where we can retrieve an arbitrary string at the cost of one query.

We can give a quantum computer access to a oracle using the standard reversible computation construction discussed in the previous section. That is, instead of having a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, we produce a unitary operator O_f which performs the map

$$O_f|x\rangle|y\rangle = |x\rangle|y \oplus f(x)\rangle.$$

If $m = 1$, so f returns one bit, it would also make sense to consider an oracle U_f which does not use an ancilla, but instead flips the phase of an input state $|x\rangle$ by applying the map

$$U_f|x\rangle = (-1)^{f(x)}|x\rangle.$$

This variant is thus sometimes known as the *phase oracle*. Given access to a bit oracle, we can simulate a phase oracle by attaching an ancilla qubit in the state $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$:

$$O_f|x\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) = \frac{1}{\sqrt{2}}(|x\rangle|f(x)\rangle - |x\rangle|f(x) \oplus 1\rangle) = (-1)^{f(x)}|x\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

Note that the ancilla qubit is left unchanged by this operation, which is called the phase kickback trick. Also note that the effect of the phase oracle is not observable if we apply it to just one basis state $|x\rangle$, but only if we apply it to a superposition:

$$\sum_{x \in \{0,1\}^n} \alpha_x |x\rangle \mapsto \sum_{x \in \{0,1\}^n} (-1)^{f(x)} \alpha_x |x\rangle.$$

Importantly, note that to implement the oracles O_f and U_f we do not need to understand any more about the inner workings of f than we do classically. That is, if we are given a classical circuit computing f , we can follow a purely mechanical construction to create quantum circuits implementing O_f and U_f .

3 Grover's algorithm

A simple example of a problem that fits into the query complexity model is unstructured search on a set of N elements (Box 2).

In the unstructured search problem, we are given access to a function $f : [N] \rightarrow \{0, 1\}$ with the promise that $f(x_0) = 1$ for a unique element x_0 . Our task is to output x_0 .

Box 2: Unstructured search for a unique marked element

It is intuitively clear that the unstructured search problem should require about N queries to be solved (classically!). We can formalise this as the following proposition:

Proposition 3.1. *Let \mathcal{A} be a classical algorithm which solves the unstructured search problem on a set of N elements with failure probability $\delta < 1/2$. Then \mathcal{A} makes $\Omega(N)$ queries in the worst case.*

Proof sketch. We can think of any classical algorithm \mathcal{A} as choosing in advance, either deterministically or randomly, a sequence S of distinct indices x_1, \dots, x_N to query, and then querying them one by one. Imagine an adversary chooses the marked element x_0 uniformly at random. Then, on average, x_0 will appear at position $N/2$ in the sequence S . As for a random choice of x_0 the algorithm makes $\Omega(N)$ queries on average, the average number of queries made in the worst case must also be $\Omega(N)$. \square

In the quantum setting, we will see that the unstructured search problem can be solved with significantly fewer queries.

Theorem 3.2 (Grover '97). *There is a quantum algorithm which solves the unstructured search problem using $O(\sqrt{N})$ queries.*

For simplicity, assume that $N = 2^n$ for some integer n (this is not an essential restriction). Then Grover's algorithm is described in Box 3.


We are given access to $f : \{0, 1\}^n \rightarrow \{0, 1\}$ with the promise that $f(x_0) = 1$ for a unique element x_0 . We use a quantum circuit on n qubits with initial state $|0\rangle^{\otimes n}$. Let H denote the Hadamard gate, and let U_0 denote the n -qubit operation which inverts the phase of $|0\rangle$: $U_0|0\rangle = -|0\rangle$, $U_0|x\rangle = |x\rangle$ for $x \neq 0$.

1. Apply $H^{\otimes n}$.
2. Repeat the following operations T times, where $T = \lfloor \frac{\pi}{4} \sqrt{N} \rfloor$:
 - (a) Apply U_f .
 - (b) Apply $D := -H^{\otimes n} U_0 H^{\otimes n}$.
3. Measure all the qubits and output the result.

Box 3: Grover's algorithm

The overall unitary operation performed is thus $(-H^{\otimes n} U_0 H^{\otimes n} U_f)^T H^{\otimes n}$, where $T = \lfloor \frac{\pi}{4} \sqrt{N} \rfloor$. (Incidentally, note that the minus sign in front of D can actually be omitted without affecting the

correctness of the algorithm, but it is helpful for the analysis.) In circuit diagram form, Grover's algorithm looks like this:



It may be far from clear initially why this algorithm works, or indeed whether it does work. To describe the algorithm, we introduce unitary operators $I_{|\psi\rangle}$ and $R_{|\psi\rangle}$, where $|\psi\rangle$ is an arbitrary state. These are defined as follows:

$$I_{|\psi\rangle} := I - 2|\psi\rangle\langle\psi|, \quad R_{|\psi\rangle} := -I_{|\psi\rangle} = 2|\psi\rangle\langle\psi| - I,$$

where I is the identity. $I_{|\psi\rangle}$ can be seen as an “inversion about $|\psi\rangle$ ” operation, while $R_{|\psi\rangle}$ can be seen as an “reflection about $|\psi\rangle$ ” operation. An arbitrary state $|\phi\rangle$ can be expanded as

$$|\phi\rangle = \alpha|\psi\rangle + \beta|\psi^\perp\rangle$$

for some α and β , and some state $|\psi^\perp\rangle$ such that $\langle\psi|\psi^\perp\rangle = 0$. Then

$$I_{|\psi\rangle}|\phi\rangle = -\alpha|\psi\rangle + \beta|\psi^\perp\rangle,$$

so $I_{|\psi\rangle}$ has flipped the phase of the component corresponding to $|\psi\rangle$, and left the component orthogonal to $|\psi\rangle$ unchanged. $R_{|\psi\rangle}$ has the opposite effect. Observe that, in the unstructured search problem with marked element x_0 , $U_f = I_{|x_0\rangle}$. Further observe that

$$H^{\otimes n}U_0H^{\otimes n} = H^{\otimes n}(I - 2|0\rangle\langle 0|)H^{\otimes n} = I - 2|+\rangle\langle +| = I_{|+\rangle},$$

where $|+\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle$, so $D = -I_{|+\rangle}$. By moving the minus sign, the algorithm can equally well be thought of as alternating the operations $-I_{|x_0\rangle}$ and $I_{|+\rangle}$, or equivalently $R_{|x_0\rangle}$ and $-R_{|+\rangle}$.

We have the following claims:

1. For any states $|\psi\rangle$, $|\phi\rangle$, and any state $|\xi\rangle$ in the 2d plane spanned by $|\psi\rangle$ and $|\phi\rangle$, the states $R_{|\psi\rangle}|\xi\rangle$ and $R_{|\phi\rangle}|\xi\rangle$ remain in this 2d plane.

This is immediate from geometric arguments, but one can also calculate explicitly:

$$\begin{aligned} R_{|\psi\rangle}(\alpha|\psi\rangle + \beta|\phi\rangle) &= R_{|\psi\rangle}(\alpha|\psi\rangle + \beta(\gamma|\psi\rangle + \delta|\psi^\perp\rangle)) = (\alpha + \beta\gamma)|\psi\rangle - \beta\delta|\psi^\perp\rangle \\ &= (\alpha + 2\beta\gamma)|\psi\rangle - \beta(\gamma|\psi\rangle + \delta|\psi^\perp\rangle) = (\alpha + 2\beta\gamma)|\psi\rangle - \beta|\phi\rangle. \end{aligned}$$

2. Within the 2d plane spanned by orthogonal states $|\psi\rangle$, $|\psi^\perp\rangle$, $I_{|\psi\rangle} = -R_{|\psi\rangle} = R_{|\psi^\perp\rangle}$.

Again, one can calculate explicitly that


$$-R_{|\psi\rangle}(\alpha|\psi\rangle + \beta|\psi^\perp\rangle) = -\alpha|\psi\rangle + \beta|\psi^\perp\rangle = R_{|\psi^\perp\rangle}(\alpha|\psi\rangle + \beta|\psi^\perp\rangle).$$

3. If $|\xi\rangle$ is within the 2d plane spanned by $|\psi\rangle$, $|\psi^\perp\rangle$,


$$R_{|\psi\rangle}|\xi\rangle = \langle\psi|\xi\rangle|\psi\rangle - \langle\psi^\perp|\xi\rangle|\psi^\perp\rangle.$$

This is just a straightforward calculation.

Combining these claims, we see that each step of Grover's algorithm consists of two reflections in the plane spanned by $|+\rangle$, $|x_0\rangle$: a reflection about $|x_0\rangle$ followed by a reflection about $|+\perp\rangle$, a state orthogonal to $|+\rangle$ within this plane. We can illustrate this with the following diagram, demonstrating the effect of these operations on an arbitrary state $|\xi\rangle$ within this 2d plane:



We see that $|\xi\rangle$ has moved closer to $|x_0\rangle$. In fact, geometrically speaking, the composition of two reflections is a rotation! If the angle between $|\xi\rangle$ and $|x_0\rangle$ is θ , and the angle between $|x_0\rangle$ and $|+\perp\rangle$ is γ , composing these two reflections rotates $|\xi\rangle$ in the direction of $|+\perp\rangle$ by an angle of $2\theta - 2(\theta - \gamma) = 2\gamma$. This is proven by picture in the following diagram but could also be shown using the representation of rotations and reflections by 2d matrices.



Repeating the Grover iteration continues to rotate $|\xi\rangle$ within this plane by angle 2γ . We stop when we are as close to $|x_0\rangle$ as possible. We start with $|\xi\rangle = |+\rangle$, so the initial angle between $|\xi\rangle$ and $|x_0\rangle$ is $\pi/2 - \gamma$. We can calculate what γ is by using the formula $\cos \gamma = \langle x_0 | + \perp \rangle$, so $\sin \gamma = \langle x_0 | + \rangle = 1/\sqrt{N}$. As $\sin x \approx x$ for small x , we expect the number of iterations required to move from an angle of $\pi/2 - \gamma$ down to an angle of 0 to be about $(\pi/4)\sqrt{N}$. One can calculate this more precisely: after T iterations, the angle between $|\xi\rangle$ and $|x_0\rangle$ is

$$\gamma_T := \pi/2 - (2T + 1) \arcsin(1/\sqrt{N}),$$

so the probability of obtaining the outcome x_0 when we measure is precisely

$$|\langle \xi | x_0 \rangle|^2 = \cos^2(\gamma_T) = \sin^2((2T + 1) \arcsin(1/\sqrt{N})). \quad (1)$$

Maximising this by taking T to be the integer nearest to

$$\frac{\pi}{4 \arcsin(1/\sqrt{N})} - \frac{1}{2} = \frac{\pi}{4} \sqrt{N} - \frac{1}{2} - O\left(\frac{1}{N}\right),$$

we learn x_0 with probability $1 - O(1/N)$ using $O(\sqrt{N})$ queries. (The above expression uses $\arcsin x = x + O(x^3)$ for small x .) Figure 4 illustrates the success probabilities for $N = 100$.




Figure 4: Success probabilities of Grover’s algorithm for $N = 100$.

We see that, as the number of uses of the Grover iterate increases past $(\pi/4)\sqrt{N}$, the success probability starts to decrease. This is sometimes referred to as the “soufflé” property of Grover’s algorithm: if we open the oven too early, or too late, the soufflé falls.

A particularly nice case, where we can determine an exact solution, is $N = 4$. Here we have $\arcsin(1/2) = \pi/6$, so if we plug in $T = 1$ to Eqn. (1), the probability of getting the outcome x_0 is $\sin^2(\pi/2) = 1$; so we get the right answer with certainty after only one query.

We have calculated the query complexity of Grover’s algorithm; what is the time complexity? As well as the calls to U_f , we need to implement the operation D . But this can be done efficiently: D consists of two layers of n Hadamard gates and an operation which flips the phase if the input is not all 0’s. This operation – which is based on computing the bitwise OR of n bits – can be implemented using $O(\log n)$ layers of classical gates. So the overhead is $O(n)$ gates per iteration, and depth only $O(\log n)$. This is minor compared with the number of iterations, which is $\Theta(2^{n/2})$.

3.1 Multiple marked elements

Grover’s algorithm can also be used when there are $M > 1$ marked elements. In this setting, the operator U_f inverts the phase of input elements $x \in S$, for some unknown subset $S \subseteq [N]$, where $|S| = M$. U_f is still related to an inversion operator, but now an inversion about an M -dimensional subspace:

$$U_f = I - 2\Pi_S,$$

where $\Pi_S = \sum_{x \in S} |x\rangle\langle x|$. If we define the state $|S\rangle := \frac{1}{\sqrt{M}} \sum_{x \in S} |x\rangle$, we see that

$$\begin{aligned} I_{|S\rangle}|+\rangle &= (I - 2|S\rangle\langle S|)|+\rangle = |+\rangle - 2 \left(\frac{1}{M} \sum_{x,y \in S} |x\rangle\langle y| \right) \left(\frac{1}{\sqrt{N}} \sum_{x \in \{0,1\}^n} |x\rangle \right) \\ &= |+\rangle - \frac{2}{\sqrt{N}} \sum_{x \in S} |x\rangle = (I - 2\Pi_S)|+\rangle = U_f|+\rangle \end{aligned}$$

and similarly

$$I_{|S\rangle}|S\rangle = -|S\rangle = (I - 2\Pi_S)|S\rangle = U_f|S\rangle.$$

That is, the U_f operation behaves like an inversion-about- $|S\rangle$ operator for any states in the subspace spanned by $|+\rangle$ and $|S\rangle$. The whole of the previous analysis goes through, except that now the angle γ moved at each step satisfies $\sin \gamma = \langle S|+\rangle = \sqrt{M/N}$. Thus after T iterations we have

$$|\langle \xi|S\rangle|^2 = \cos^2(\gamma_T) = \sin^2((2T + 1) \arcsin(\sqrt{M/N})).$$

By a similar argument to before we can pick $T \approx (\pi/4)\sqrt{N/M}$ to obtain overlap with $|S\rangle$ close to 1. When we measure at the end of the algorithm, we get an element of the subset S (and in fact a uniformly random such element) with probability $|\langle \xi | S \rangle|^2$. In particular, observe that when $M = N/4$, we again measure an element of S with certainty using only one query.

What if we do not know the number of marked elements in advance? The following simple trick can deal with this. First run the algorithm assuming there is 1 marked element; if it fails, try again assuming there are 2 marked elements; then 4, 8, etc. The total number of queries used is roughly

$$\sum_{k=0}^{\log_2 N} \frac{\pi}{4} \sqrt{\frac{N}{2^k}} = \frac{\pi}{4} \sqrt{N} \sum_{k=0}^{\log N} 2^{-k/2} = O(\sqrt{N}).$$

If the number of marked elements is M , at least one of the iterations must choose a value of T which is within a factor of 2 of the optimal value $T' \approx (\pi/4)\sqrt{N/M}$. Then, as $(2T' + 1) \arcsin(\sqrt{M/N}) = \pi/2 + O(M/N)$,

$$\begin{aligned} \sin^2((2T + 1) \arcsin(\sqrt{M/N})) &= \sin^2\left(\frac{2T + 1}{2T' + 1} (2T' + 1) \arcsin(\sqrt{M/N})\right) \\ &= \sin^2\left(\frac{2T + 1}{2T' + 1} (\pi/2 + O(M/N))\right), \end{aligned}$$

which is lower-bounded by a strictly positive constant if M is small with respect to N . Repeating the whole algorithm $O(1)$ times allows us to achieve an arbitrarily high success probability.

This algorithm might still have a high probability of failing in the case where $M = \Omega(N)$. To find a marked element in this case we can just sample $O(1)$ random values of $f(x)$ classically; we will find a marked element with high probability.

3.2 Problems in NP and “database search”

Grover’s algorithm is often presented as a way of searching an unstructured database, or a database which is not structured in a way that is useful to us; for example, trying to search by phone number in a phone book ordered by name. However, the primary use of Grover’s algorithm (at least initially) is likely not to be searching physical databases, but instead searching for solutions to computational problems.

Grover’s algorithm gives a quadratic quantum speedup over classical exhaustive search for any problem in NP. This is because we can choose the oracle operation f to be the classical checking circuit which takes an input a claimed solution, and outputs 1 if the solution is correct, and 0 otherwise. If there are N possible solutions to the problem, Grover’s algorithm lets us find a solution using only $O(\sqrt{N})$ checks. Note that this does not immediately imply that Grover’s algorithm is better than any classical algorithm; in some cases, there could be a more efficient classical algorithm based on using the structure of the problem.

But could we also use Grover’s algorithm to search a real database? This would rely on the use of a “quantum RAM” which allowed elements of the memory to be efficiently queried in superposition. In principle, there do not seem any fundamental reasons why such a memory could not be constructed. However, in practice building a quantum RAM is likely to be challenging.

3.3 Amplitude amplification

The basic idea behind Grover’s algorithm can be generalised remarkably far, to an algorithm for finding solutions to any problem using a heuristic. This algorithm is known as *amplitude amplification*.

Imagine we have N possible solutions, of which a subset S are “good”, and we would like to find a good solution. As well as having access to a “checking” algorithm f as before, where $f(x) = 1$ if and only if x is marked, we now have access to a “guessing” algorithm \mathcal{A} , which has the job of producing potential solutions to the problem. It performs the map

$$\mathcal{A}|0\rangle = \sum_{x=0}^{N-1} \alpha_x |x\rangle$$

for some coefficients $\{\alpha_x\}$. So, if we were to apply \mathcal{A} and then measure, the probability that we would obtain a good solution is

$$p := \sum_{x \in S} |\alpha_x|^2;$$

we think of \mathcal{A} as a heuristic which tries to output a good solution. We can use f to check whether a claimed solution is actually good. If we repeated Algorithm \mathcal{A} until we got a good solution, the expected number of trials we would need is $\Theta(1/p)$.

We now describe the amplitude amplification algorithm.

We are given access to \mathcal{A} and U_f as above.

1. Apply \mathcal{A} to the starting state $|0\rangle$.
2. Repeat the following operations T times, for some T to be determined:
 - (a) Apply U_f .
 - (b) Apply $-\mathcal{A}I_0\mathcal{A}^{-1}$.
3. Measure all the qubits and output the result.

Box 5: Amplitude amplification

Note that this is exactly the same as Grover’s algorithm, except that we have replaced the $H^{\otimes n}$ operations with \mathcal{A} or \mathcal{A}^{-1} . Write

$$|\psi\rangle = \mathcal{A}|0\rangle, \quad |G\rangle = \frac{\Pi_S|\psi\rangle}{\|\Pi_S|\psi\rangle\|},$$

where again $\Pi_S = \sum_{x \in S} |x\rangle\langle x|$. We now repeat the analysis of the previous section, except that we replace $|+\rangle$ with $|\psi\rangle$ and $|S\rangle$ with $|G\rangle$. We observe that everything goes through just as before! The first operation applied is equivalent to $I_{|G\rangle}$, and the second is equivalent to $-I_{|\psi\rangle}$. We start with the state $|\psi\rangle$ and rotate towards $|G\rangle$. The angle γ moved at each step now satisfies

$$\sin \gamma = \langle \psi | G \rangle = \|\Pi_S|\psi\rangle\| = \sqrt{p},$$

so the number of iterations required to move from $|\psi\rangle$ to $|G\rangle$ is $O(1/\sqrt{p})$ – a quadratic improvement.

Finally observe that we can generalise one step further, by replacing the algorithm U_f with inversion about an arbitrary subspace, rather than a subspace defined in terms of computational basis vectors. This allows us to use amplitude amplification to drive amplitude towards an arbitrary subspace, or indeed to create an arbitrary quantum state, given the ability to reflect about that state.

7	3	4	2	9	7	3	4	2	9	7	3	4	2	9	7	3	4	2	9
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Figure 6: A periodic sequence, with period 5, which is one-to-one on each period.

4 The Quantum Fourier Transform and periodicity

We now introduce an important unitary transformation which is used in a number of different contexts in quantum information theory: the quantum Fourier transform (QFT) over \mathbb{Z}_N , the integers modulo N . This can be seen as a generalisation of the familiar Hadamard gate. The QFT is the map

$$Q_N|x\rangle = \frac{1}{\sqrt{N}} \sum_{y \in \mathbb{Z}_N} \omega_N^{xy},$$

where $\omega_N := e^{2\pi i/N}$, and xy is just the product of the two numbers x and y , thought of as integers. We sometimes omit the subscript N where there is no ambiguity. Some examples of the QFT in small dimension:

$$Q_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad Q_3 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 1 & 1 \\ 1 & e^{2\pi i/3} & e^{\pi i/3} \\ 1 & e^{\pi i/3} & e^{2\pi i/3} \end{pmatrix}, \quad Q_4 = \frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{pmatrix}.$$

To see that the QFT is unitary, we calculate the inner product of rows x and z , which equals

$$\frac{1}{N} \sum_{y \in \mathbb{Z}_N} (\omega_N^{xy})^* \omega_N^{zy} = \frac{1}{N} \sum_{y \in \mathbb{Z}_N} \omega_N^{(z-x)y}.$$

To compute this sum, we use the formula for the sum of a geometric series:

$$\sum_{k=0}^{r-1} x^k = \begin{cases} \frac{1-x^r}{1-x} & \text{if } x \neq 1 \\ r & \text{if } x = 1 \end{cases}, \quad (2)$$

implying that the inner product is equal to 1 if $z = x$, and $\frac{1-\omega_N^{(z-x)N}}{1-\omega_N}$ otherwise. But as $\omega_N^N = 1$, the inner product is 0 if $z \neq x$. More generally, for any integer j ,

$$\frac{1}{N} \sum_{y \in \mathbb{Z}_N} \omega_N^{jy} = \begin{cases} 0 & \text{if } j \not\equiv 0 \pmod{N} \\ 1 & \text{if } j \equiv 0 \pmod{N} \end{cases}, \quad (3)$$

a fact which will be useful later.

The QFT is exactly the same transformation as the Discrete Fourier Transform (DFT) used for classical computation and signal processing, up to the nonstandard normalisation of $1/\sqrt{N}$.

4.1 Periodicity determination

One of the most important applications of the QFT is determining the period of a periodic function. Imagine we are given access to an oracle function $f : \mathbb{Z}_N \rightarrow \mathbb{Z}$, such that:

- f is **periodic**: there exists r such that r divides N and $f(x+r) = f(x)$ for all $x \in \mathbb{Z}_N$;
- f is **one-to-one** on each period: for all pairs (x, y) such that $|x - y| < r$, $f(x) \neq f(y)$.

Our task is to determine r .

The periodicity determination algorithm is presented in Box 7.

We are given access to a periodic function f with period r , which is one-to-one on each period. We start with the state $|0\rangle|0\rangle$.

1. Apply Q_N to the first register.
2. Apply O_f to the two registers.
3. Measure the second register.
4. Apply Q_N to the first register.
5. Measure the first register; let the answer be k .
6. Simplify the fraction k/N as far as possible and return the denominator.

Box 7: Periodicity determination

The initial sequence of operations which occur during the algorithm is:

$$|0\rangle|0\rangle \xrightarrow{1} \frac{1}{\sqrt{N}} \sum_{x \in \mathbb{Z}_N} |x\rangle|0\rangle \xrightarrow{2} \frac{1}{\sqrt{N}} \sum_{x \in \mathbb{Z}_N} |x\rangle|f(x)\rangle.$$

When the second register is measured, we receive an answer, say z . By the periodic and one-to-one properties of f , all input values $x \in \mathbb{Z}_N$ for which $f(x) = z$ are of the form $x_0 + jr$ for some x_0 and integer j . The state therefore collapses to something of the form

$$\sqrt{\frac{r}{N}} \sum_{j=0}^{N/r-1} |x_0 + jr\rangle.$$

After we apply the QFT, we get the state

$$\frac{\sqrt{r}}{N} \sum_{j=0}^{N/r-1} \left(\sum_{y \in \mathbb{Z}_N} \omega_N^{y(x_0 + jr)} |y\rangle \right) = \frac{\sqrt{r}}{N} \sum_{y \in \mathbb{Z}_N} \omega_N^{yx_0} \left(\sum_{j=0}^{N/r-1} \omega_N^{jry} \right) |y\rangle.$$

Observe that, as r divides N , $\omega_N^r = e^{2\pi i(r/N)} = \omega_{N/r}$. This state is thus equivalent to

$$\frac{\sqrt{r}}{N} \sum_{y \in \mathbb{Z}_N} \omega_N^{yx_0} \left(\sum_{j=0}^{N/r-1} \omega_{N/r}^{jy} \right) |y\rangle.$$

By Eqn. (3), the sum over j is 0 unless $y \equiv 0 \pmod{N/r}$, or in other words if $y = \ell N/r$ for some integer ℓ . So we can rewrite this state as

$$\frac{1}{\sqrt{r}} \sum_{\ell=0}^{r-1} \omega_N^{\ell x_0 N/r} |\ell N/r\rangle.$$




Figure 8: Periodicity determination as above with $N = 20$, $r = 5$. First diagram illustrates the probabilities of measurement outcomes after step 3 (for one possible measurement result for the second register), second diagram illustrates probabilities after step 5.

When we perform the final measurement, we receive an outcome $k = \ell_0 N/r$, for some ℓ_0 picked uniformly at random from $0, \dots, r - 1$. We know that

$$k = \frac{\ell_0 N}{r}, \quad \text{so } \frac{k}{N} = \frac{\ell_0}{r}.$$

In this equation, we know N and k and would like to determine r . If it happened that ℓ_0 were coprime to r , we could cancel the fraction on the left-hand side and output the denominator. What is the probability that we are lucky in this way?

Fact 4.1. *Fix an positive integer a and pick b uniformly at random from the integers between 0 and a . Then the probability that b is coprime to a is $\Omega(1/\log \log a)$.*

Thus, if we repeat the whole procedure $O(\log \log r) = O(\log \log N)$ times, we are quite likely to find the period r . Why? If we have a probabilistic procedure which succeeds with probability p , the probability that it fails every time over R repetitions is exactly

$$(1 - p)^R \leq e^{-pR},$$

so it suffices to take $R = O(1/p)$ to achieve, say, 99% success probability. Each time the algorithm returns a claimed period, we can check whether it is really a period of the function using two additional queries. Each use of the quantum algorithm therefore makes 3 queries, so the whole algorithm makes $O(\log \log N)$ queries in total. In terms of time complexity, the most complicated classical processing required is the elementary arithmetic in step 6, which can be implemented (via Euclid's algorithm) using $\text{poly}(\log N)$ arithmetic operations. However, we have not yet shown that we can implement the QFT Q_N efficiently.

4.2 Efficient implementation of the QFT

We will show here how to implement Q_N efficiently – i.e. using a circuit of size $O(\text{poly log } N)$ – in the case where N is a power of 2. (In fact, the QFT can also be implemented (approximately) efficiently when N is not a power of 2.) The efficient implementation is based on the same ideas as the classical Fast Fourier Transform (FFT). To begin with, we observe that the output of the QFT, when applied to a computational basis state, has an efficient description as a product state.

Assume that $N = 2^n$ for some integer n , and represent each $y \in \mathbb{Z}_N$ by the n -bit string

$(y_0, y_1, \dots, y_{n-1})$, where $y = y_0 + 2y_1 + 4y_2 + \dots + 2^{n-1}y_{n-1}$. Then

$$\begin{aligned}
Q_N|x\rangle &= \frac{1}{2^{n/2}} \sum_{y \in \mathbb{Z}_{2^n}} \omega_2^{xy} |y\rangle \\
&= \frac{1}{2^{n/2}} \sum_{y \in \{0,1\}^n} \omega_2^{x(\sum_{j=0}^{n-1} 2^j y_j)} |y_{n-1}\rangle |y_{n-2}\rangle \dots |y_0\rangle \\
&= \left(\frac{1}{\sqrt{2}} \sum_{y_{n-1} \in \{0,1\}} \omega_2^{2^{n-1} x y_{n-1}} |y_{n-1}\rangle \right) \left(\frac{1}{\sqrt{2}} \sum_{y_{n-2} \in \{0,1\}} \omega_2^{2^{n-2} x y_{n-2}} |y_{n-2}\rangle \right) \dots \left(\frac{1}{\sqrt{2}} \sum_{y_0 \in \{0,1\}} \omega_2^{x y_0} |y_0\rangle \right) \\
&= \bigotimes_{j=1}^n \left(\frac{1}{\sqrt{2}} \sum_{y_{n-j} \in \{0,1\}} \omega_2^{x y_{n-j}} |y_{n-j}\rangle \right).
\end{aligned}$$

Because $x y_{n-j} \equiv 0 \pmod{2^j}$ when x is an integer multiple of 2^j , we see that the j 'th qubit of the output only depends on the j bits x_0, \dots, x_{j-1} . We can write this another way, as

$$Q_N|x\rangle = \frac{1}{2^{n/2}} \left(|0\rangle + e^{2\pi i(.x_0)} |1\rangle \right) \left(|0\rangle + e^{2\pi i(.x_1 x_0)} |1\rangle \right) \dots \left(|0\rangle + e^{2\pi i(.x_{n-1} \dots x_0)} |1\rangle \right),$$

where the notation $(.x_j \dots x_0)$ is used for the binary fraction

$$\frac{x_j}{2} + \frac{x_{j-1}}{4} + \dots + \frac{x_0}{2^{j-1}}.$$

So we see that the first qubit of the output depends on only the last qubit of the input, the second qubit depends on the last two, etc. We can utilise this structure by building up the output state in reverse order. The last stage of the circuit creates the correct state for the first qubit, which is then not used again; the last but one stage creates the correct state for the second qubit, etc. To produce the correct state for each qubit, we can use the gates H and R_d , where


$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad R_d = \begin{pmatrix} 1 & 0 \\ 0 & e^{\pi i/2^d} \end{pmatrix}.$$

Here we are assuming that we have access to R_d gates for arbitrary d ; as briefly discussed in Section 2, this is not essential as any universal gate set will allow us to approximately implement these gates. Observe that the Hadamard gate can be written as the map

$$H|x\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle + e^{2\pi i(-.x)} |1\rangle \right)$$

for $x \in \{0,1\}$. We can use this to start building up a binary fraction in the phase of the basis state $|1\rangle$. Applying a R_d gate to this state will add $1/2^{d+1}$ to this binary fraction. To apply R_d conditional on the bits of x , we will use controlled- R_d gates.

The easiest way to illustrate this process is with an example. The overall circuit for the QFT on 4 qubits can be depicted as



Observe that the output state is backwards, i.e. the qubits appear in reverse order. They can be returned to the original order, if desired, using swap gates. How many gates in total are used in the circuit? The j 'th stage of the circuit, for $j = 1, \dots, n$, uses one Hadamard gate and $n - j$ R_d gates. Thus the overall number of gates used is $O(n^2)$; $n(n - 1)/2$ R_d gates and n Hadamard gates, then n additional swap gates, if used. This is $O(\log^2 N)$, so we have indeed obtained an efficient circuit.

This complexity can be improved further, to $O(n \log n)$, if we are content with an approximate version of the QFT. The observation which implies this is that many of the operations in the circuit are R_d gates for large values of d , which do not affect the output significantly. Indeed, it turns out that there is a constant C such that if we omit the gates R_d with $d \geq C \log n$, for any input state $|x\rangle$ the output is close to the input up to an error of at most $1/\text{poly}(n)$. The modified circuit uses $O(\log n)$ gates at each stage, so $O(n \log n)$ in total.

5 Integer factorisation

The main application of periodicity determination is Shor’s quantum algorithm for integer factorisation. Given an n -digit integer N as input, this algorithm outputs a non-trivial factor of N (or that N is prime, if this is the case) with success probability $1 - \epsilon$, for any $\epsilon > 0$, in time $O(n^3)$. The best classical algorithm known (the general number field sieve) runs in time $e^{O(n^{1/3} \log^{2/3} n)}$. In fact, this is a heuristic bound and this algorithm’s worst-case runtime has not been rigorously determined; the best proven bound is somewhat higher. Shor’s algorithm thus achieves a super-polynomial improvement. This result might appear only of mathematical interest, were it not for the fact that the widely-used RSA public-key cryptosystem relies on the hardness of integer factorisation. Shor’s efficient factorisation algorithm implies that this cryptosystem is insecure against attack by a large quantum computer.

Unfortunately (?), proving correctness of Shor’s algorithm requires going through a number of technical details. First we need to show that factoring reduces to a periodicity problem – though in this case an *approximate* periodicity problem. This part uses only classical number theory. Then we need to show that periodicity can still be determined even in the setting where the input function is only approximately periodic. This part uses the theory of continued fractions.

5.1 From factoring to periodicity

The basic skeleton of the quantum factorisation algorithm is given in Box 9. It is based on two “magic” subroutines. The first is a classical algorithm for computing the greatest common divisor (gcd) of two integers. This can be achieved efficiently using Euclid’s algorithm. The second ingredient is an algorithm for computing the order of an integer a modulo N , i.e. the smallest integer r such that $a^r \equiv 1 \pmod{N}$; this is where we will use periodicity determination. As long as a and N are coprime, such an integer p exists:

Fact 5.1 (Euler’s theorem). *If a and N are coprime then there exists p such that $a^p \equiv 1 \pmod{N}$.*

Let N denote the integer to be factorised. Assume that N is not even or a power of a prime.

1. Choose $1 < a < N$ uniformly at random.
2. Compute $b = \gcd(a, N)$. If $b > 1$ output b and stop.
3. Determine the order r of a modulo N . If r is odd, the algorithm has failed; terminate.
4. Compute $s = \gcd(a^{r/2} - 1, N)$. If $s = 1$, the algorithm has failed; terminate.
5. Output s .

Box 9: Integer factorisation algorithm (overview)

We start by showing that this algorithm does work, assuming that the subroutines used all work correctly. If a is coprime to N , there exists r such that $a^r \equiv 1 \pmod{N}$. If, further, such an r is even, we can factor

$$a^r - 1 = (a^{r/2} + 1)(a^{r/2} - 1) \equiv 0 \pmod{N}.$$

So N divides the product $(a^{r/2} + 1)(a^{r/2} - 1)$. Because r is the *smallest* integer x such that $a^x \equiv 1 \pmod N$, we know that $a^{r/2} - 1$ is not divisible by N . Then, if in addition $s = \gcd(a^{r/2} - 1, N) \neq 1$, s would be a nontrivial factor of N . This condition would hold if $a^{r/2} + 1$ is not divisible by N . This, and r being even, turn out to occur with quite high probability:

Fact 5.2. *Let N be odd and not a power of a prime. If $1 < a < N$ is chosen uniformly at random with $\gcd(a, N) = 1$, then $\Pr[r \text{ is even and } a^{r/2} \not\equiv -1 \pmod N] \geq 1/2$.*

The algorithm thus succeeds with probability at least $1/2$. If it fails, we simply repeat the whole process. After K repetitions, we achieve a failure probability of at most $1/2^K$. Even if the order-finding procedure has some small probability of error (which will turn out to be the case), we can check whether the algorithm's output s is correct by attempting to divide N by s .

We assumed throughout that N is not even or a power of a prime. If N is even, we simply output 2 as a factor. To deal with the case that $N = p^\ell$ for some prime p and some integer $\ell > 0$, we observe that we can efficiently compute the roots $N^{1/k}$, for $k = 2, \dots, \log_2 N$. If any of these is an integer, we have found a factor of N . Finally, what about if N is itself prime? In this case the algorithm will fail every time. We can therefore output "prime" after a suitable number of failures.

Example 5.3. *Consider $N = 15$. Imagine we choose $a = 7$ at random. Then $\gcd(7, 15) = 1$. The function $f(x) = 7^x \pmod{15}$ takes values $1, 7, 4, 13, 1, \dots$ for $x \geq 0$. So $r = 4$ and we have $(7^2 + 1)(7^2 - 1) \equiv 0 \pmod{15}$. The greatest common divisor of $7^2 - 1 = 48$ and 15 is 3 , which is indeed a factor of 15 .*

It remains to show how to implement step 3. Consider the function $f : \mathbb{Z} \rightarrow \mathbb{Z}_N$ defined by

$$f(x) = a^x \pmod N.$$

We have $f(x + y) = f(x)f(y)$ and, by Euler's theorem, $f(r) = 1$. So $f(x + r) = f(x)f(r) = f(x)$ for all x , i.e. f is periodic with period r . Since r is the smallest integer such that $f(r) = 1$, we also have that f is one-to-one on each period. However, although this function is periodic on the whole domain \mathbb{Z} , we will need to truncate it to a finite size. If we knew what the period was, we could choose this size to make the function periodic again, but of course we do not know this in advance. This will lead to the function becoming no longer exactly periodic, but just approximately periodic.

5.2 Approximate periodicity

We restrict the function $f(x) = a^x \pmod N$ to the set $x \in \{0, \dots, M - 1\}$, where $M = 2^m$ is the smallest power of 2 greater than N^2 (we will see the reasons behind this choice later). Write $M = Br + b$ for $0 \leq b < r$, where $B = \lfloor M/r \rfloor$. That is, the function is periodic up to the last period, which is truncated and contains only b elements, rather than r . We apply the steps of the periodicity-determination algorithm to f as in Section 4.1. That is, we first construct the state

$$\frac{1}{\sqrt{M}} \sum_{x=0}^{M-1} |x\rangle |f(x)\rangle,$$

then measure the second register to receive an answer $z = f(x_0)$ for some x_0 . The state of the first register then becomes

$$|\psi\rangle = \frac{1}{\sqrt{A}} \sum_{j=0}^{A-1} |x_0 + jr\rangle,$$




Figure 10: The probabilities of different measurement outcomes for a function with period 5, with $M = 32$. Note the peaks around multiples of $32/5 = 6.4$.

where $A = B + 1$ if $x_0 < b$, and $A = B$ if $x_0 \geq b$. Write

$$Q_M|\psi\rangle = \sum_{y=0}^{M-1} \alpha_y|y\rangle$$

for the resulting state when we apply the QFT to $|\psi\rangle$. By direct calculation we have

$$\alpha_y = \frac{1}{\sqrt{MA}} \sum_{j=0}^{A-1} \omega_M^{y(x_0+jr)} = \frac{\omega_M^{yx_0}}{\sqrt{MA}} \sum_{j=0}^{A-1} (\omega_M^{yr})^j.$$

Using the formula (2) for the sum of a geometric series, the sum evaluates to

$$\frac{1 - \omega_M^{yrA}}{1 - \omega_M^{yr}}$$

if $yr \not\equiv 0 \pmod{M}$, and evaluates to A otherwise. Previously, in the case of exact periodicity determination, we had $A = B = M/r$, so the numerator was 0 unless $yr \not\equiv 0 \pmod{M}$, or in other words y is a multiple of M/r . Now we aim to show that, when we measure, we get an outcome y which is *close* to a multiple of the non-integer value M/r with high probability. This situation is illustrated in Figure 10.

When we measure, the probability of obtaining outcome y is

$$\Pr[y] = \frac{1}{MA} \left| \frac{1 - \omega_M^{yrA}}{1 - \omega_M^{yr}} \right|^2 = \frac{1}{MA} \left| \frac{1 - e^{2\pi i yr A/M}}{1 - e^{2\pi i yr/M}} \right|^2 = \frac{\sin^2(\pi yr A/M)}{MA \sin^2(\pi yr/M)}.$$

To see the third equality, note that $|1 - e^{i\theta}| = |e^{i\theta/2} - e^{-i\theta/2}| = 2|\sin(\theta/2)|$ for any real θ . Now consider values y of the form $y = \lfloor \ell M/r \rfloor$ for some integer ℓ . We can write any such integer as $y = \ell M/r + \epsilon$ for some small ϵ such that $|\epsilon| \leq 1/2$. Indeed, we have the slightly stronger bound (which we will need) that $|\epsilon| \leq 1/2 - 1/r$. This holds because (a) $\ell M/r$ is an integer divided by r , so the distance from the closest integer is an multiple of $1/r$; (b) $r < N$ and $M > N^2$ is a power of 2, so any factors of 2 in the denominator of the fraction $\ell M/r$ can be cancelled and we cannot have $|\epsilon| = 1/2$.

Then

$$\Pr[y] = \frac{\sin^2(\pi(\ell M/r + \epsilon)rA/M)}{MA \sin^2(\pi(\ell M/r + \epsilon)r/M)} = \frac{\sin^2(\ell A\pi + \epsilon r A\pi/M)}{MA \sin^2(\ell\pi + \epsilon r\pi/M)} = \frac{\sin^2(\epsilon r A\pi/M)}{MA \sin^2(\epsilon r\pi/M)}$$

by periodicity of the function $|\sin \theta|$. We now claim (see (4) in Box 11) that the following pair of inequalities hold for any θ in the range $0 \leq \theta \leq \pi/2$:

$$(2/\pi)\theta \leq \sin \theta \leq \theta.$$

Assuming these inequalities, we have

$$\Pr[y] \geq \frac{(2\epsilon r A/M)^2}{MA(\epsilon r \pi/M)^2} = \frac{4A}{\pi^2 M} \geq \frac{4}{\pi^2 M} \left(\frac{M}{r} - 1 \right) = \frac{4}{\pi^2 r} - O(1/N^2).$$

Note that the lower bound on the numerator is valid because $|\epsilon| r A \pi / M \leq \pi/2$, which was a consequence of $|\epsilon| \leq 1/2 - 1/r$:

$$\frac{|\epsilon| r A \pi}{M} \leq \frac{(\frac{1}{2} - \frac{1}{r}) r (\frac{M}{r} + 1) \pi}{M} = \frac{\pi}{2} \left(1 - \frac{2}{r} \right) \left(1 + \frac{r}{M} \right) \leq \frac{\pi}{2} \left(1 - \frac{2}{r} \right) \left(1 + \frac{2}{r} \right) = \frac{\pi}{2} \left(1 - \frac{4}{r^2} \right) < \frac{\pi}{2},$$

where the second inequality holds because $r^2 \leq N^2 \leq 2M$. Therefore, as there are r “good” integers y of the form $\lfloor \ell M / r \rfloor$, the probability of obtaining at least one of them is at least $4/\pi^2 - O(1/N)$.

The following trigonometric inequalities will be useful for the analysis of various quantum algorithms. We state them without proof, instead referring to the figure below.

$$(2/\pi)\theta \leq \sin \theta \leq \theta \text{ for } 0 \leq \theta \leq \pi/2; \quad (4)$$

$$\cos \theta \geq 1 - \theta^2/2 \text{ for all } \theta. \quad (5)$$



Box 11: Trigonometric inequalities

5.3 Learning r from an approximate period

It remains to extract r from an integer y of the form $y = \lfloor \ell M / r \rfloor$. Divide y by M to obtain a rational number z such that

$$\left| \frac{\ell}{r} - z \right| < \frac{1}{2M} < \frac{1}{2N^2}.$$

We would like to find the fraction ℓ/r from z . We first claim that there is at most one fraction of the form ℓ'/r' with $r' < N$ satisfying the above bound. To prove this, imagine there were two such fractions $\ell'/r', \ell''/r''$. Then

$$\left| \frac{\ell'}{r'} - \frac{\ell''}{r''} \right| = \frac{|\ell' r'' - r' \ell''|}{r' r''} \geq \frac{1}{r' r''} > \frac{1}{N^2}.$$

But, as $\ell'/r', \ell''/r''$ are each within $1/(2N^2)$ of z , they must be at most distance $1/N^2$ apart, so we have a contradiction.

We have seen that it suffices to find any fraction ℓ'/r' such that $r' < N$ to learn ℓ/r . To do this, we use the theory of continued fractions. The continued fraction expansion (CFE) of z is an expression of the form

$$z = \frac{1}{a_1 + \frac{1}{a_2 + \frac{1}{a_3 + \dots}}},$$

where the a_i are positive integers. To find the integers a_i , we start by writing

$$z = \frac{1}{z'},$$

where $z' = a_1 + b$ for some integer a_1 and some $b < 1$; then repeating this process on b . Note that, for any rational z , this expansion must terminate after some number of iterations C . One can show that in fact, for any rational $z = s/t$ where s and t are m -bit integers, $C = O(m)$. Once we have calculated this expansion, if we truncate it after some smaller number of steps, we obtain an approximation to z . These approximations are called convergents of the CFE.

Example 5.4. *The continued fraction expansion of $z = 31/64$ is*

$$\frac{31}{64} = \frac{1}{2 + \frac{1}{15 + \frac{1}{2}}}.$$

The convergents are

$$\frac{1}{2}, \quad \frac{1}{2 + \frac{1}{15}} = \frac{15}{31}.$$

Fact 5.5. *Any fraction p/q with $|p/q - z| < 1/(2q^2)$ will appear as one of the convergents of the CFE of z .*

Therefore, if we carry out the continued fraction expansion of z , we are guaranteed to find ℓ/r , as the unique fraction close enough to z .

Example 5.6. *Imagine we want to factor $N = 21$. We set $M = 512 > 21^2$ and choose $a = 10$ at random. The order of a mod 21 is 6 . So we would expect the measurement outcomes we receive to be close to multiples of $512/6 = 85\frac{1}{3}$. Imagine we receive a measurement result of 427 . This is a “good” result, as the closest integer to $5 \times (512/6) = 426\frac{2}{3}$. The continued fraction expansion of $z = 427/512$ is*

$$\frac{427}{512} = \frac{1}{1 + \frac{1}{5 + \frac{1}{42 + \frac{1}{2}}}}.$$

From this we obtain the sequence of convergents

$$1, \quad \frac{1}{1 + \frac{1}{5}} = \frac{5}{6}, \quad \frac{1}{1 + \frac{1}{5 + \frac{1}{42}}} = \frac{211}{253}.$$

Only the second of these has a denominator smaller than N and is within $1/(2N^2)$ of z . Therefore, we have $\ell/r = 5/6$. We output the denominator, 6 , as our guess for the period $r \dots$ which is correct!

5.4 Complexity analysis

How complex is the final, overall algorithm? Recall that N is n bits in length. We have seen that the QFT on \mathbb{Z}_M can be implemented in time $O(\log^2 M) = O(n^2)$. To implement the modular exponentiation operation $f(x) = a^x \bmod N$ efficiently, we can use repeated squaring to produce $f(2^k)$ for any integer k in k squaring operations. Multiplying the different values $f(2^k)$ together for each k such that the k 'th bit of a is nonzero produces $f(x)$. So we require $O(n)$ multiplications of n -bit integers to compute $f(x)$. Multiplying two n -bit numbers can be achieved classically using standard long multiplication in time $O(n^2)$, so we get an overall complexity of $O(n^3)$.

It turns out (though we will not show it here) that the classical processing of the measurement results based on Euclid's algorithm and the continued fractions algorithm can also be done in time $O(n^3)$. Thus the overall time complexity of the whole algorithm is $O(n^3)$, whereas the best known classical algorithm runs in time exponential in $n^{1/3}$. The asymptotic quantum complexity can in fact be improved a bit further, to $O(n^2 \text{ poly log } n)$, by using more advanced multiplication and division algorithms with runtime $O(n \text{ poly log } n)$. However, these algorithms only become more efficient in practice for very large values of n .

6 Phase estimation

We now discuss an important primitive used in quantum algorithms called *phase estimation*, which provides a different and unifying perspective on the quantum algorithms which you have just seen. Phase estimation is once again based on the QFT over \mathbb{Z}_N , where $N = 2^n$.

Imagine we are given a unitary operator U . U may either be written down as a quantum circuit, or we may be given access to a black box which allows us to apply a controlled- U^j operation for integer values of j . We are also given a state $|\psi\rangle$ which is an eigenvector of U : $U|\psi\rangle = e^{2\pi i\phi}|\psi\rangle$ for some real ϕ such that $0 \leq \phi < 1$. We would like to determine ϕ to n bits of precision, for some arbitrary n .

To do so, we prepend an n qubit register to $|\psi\rangle$, initially in the state $|0\rangle$, and create the state

$$\frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle |\psi\rangle$$

by applying Hadamards to each qubit in the first register. We then apply the unitary operator

$$U' = \sum_{x=0}^{N-1} |x\rangle \langle x| \otimes U^x.$$


This operator can be described as: if the first register contains x , apply U x times to the second register. We are left with the state

$$\frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} e^{2\pi i\phi x} |x\rangle |\psi\rangle;$$

note that the second register is left unchanged by this operation. We now apply the operator Q^{-1} to the first register and then measure it, receiving outcome x (say). We output the binary fraction

$$0.x_1x_2\dots x_n = \frac{x_1}{2} + \frac{x_2}{4} + \dots + \frac{x_n}{2^n}$$

as our guess for ϕ . The following is an explicit circuit for the above algorithm.



Why does this algorithm work? When we perform the final measurement, the probability of getting the outcome x is

$$\frac{1}{N^2} \left| \sum_{y=0}^{N-1} e^{2\pi i\phi y - 2i\pi xy/N} \right|^2 = \frac{1}{N^2} \left| \sum_{y=0}^{N-1} e^{2\pi iy(\phi - x/N)} \right|^2.$$

First imagine that the binary expansion of ϕ is at most n bits long, or in other words $\phi = z/N$ for some $0 \leq z \leq N - 1$. In this case we have

$$\frac{1}{N^2} \left| \sum_{y=0}^{N-1} e^{2\pi iy(\phi-x/N)} \right|^2 = \frac{1}{N^2} \left| \sum_{y=0}^{N-1} e^{2\pi iy(z-x)/N} \right|^2 = \delta_{xz}$$

by the unitarity of the QFT, so the measurement outcome is guaranteed to be z , implying that the algorithm outputs ϕ with certainty. If the binary expansion of ϕ is longer than n bits, we now show that we still get the best possible answer with probability $\Omega(1)$, and indeed are very likely to get an answer close to ϕ . The proof turns out to be very similar to that of correctness of the periodicity determination algorithm in the approximate case.

Theorem 6.1. *The probability that the above algorithm outputs the real number with n binary digits which is closest to ϕ is at least $4/\pi^2$. Further, the probability that the algorithm outputs θ such that $|\theta - \phi| \geq \epsilon$ is at most $O(1/(N\epsilon))$.*

Proof. If the binary expansion of ϕ has n binary digits or fewer, we are done by the argument above. So, assuming it does not, let $\tilde{\phi}$ be the closest approximation to ϕ that has n binary digits, and write $\tilde{\phi} = a/N$ for some integer $0 \leq a \leq N - 1$. For any z , define $\delta(z) := \phi - z/N$ and note that $0 < |\delta(a)| \leq 1/(2N)$. For any ϕ , the probability of getting outcome z from the final measurement is

$$\Pr[z] = \frac{1}{N^2} \left| \sum_{y=0}^{N-1} e^{2\pi iy(\phi-z/N)} \right|^2 = \frac{1}{N^2} \left| \sum_{y=0}^{N-1} e^{2\pi iy\delta(z)} \right|^2 = \frac{1}{N^2} \left| \frac{1 - e^{2\pi i N \delta(z)}}{1 - e^{2\pi i \delta(z)}} \right|^2 = \frac{\sin^2(\pi N \delta(z))}{N^2 \sin^2(\pi \delta(z))}, \quad (6)$$

where we evaluate the sum using the formula for a geometric series. This quantity should be familiar from the proof of correctness of the periodicity determination algorithm.

We first lower bound this expression for $z = a$ to prove the first part of the lemma. As $|\delta(a)| \leq 1/(2N)$, we have $N\pi\delta(a) \leq \pi/2$. Then

$$\Pr[a] = \frac{\sin^2(\pi N \delta(a))}{N^2 \sin^2(\pi \delta(a))} \geq \frac{(2N\delta(a))^2}{N^2 (\pi \delta(a))^2} = \frac{4}{\pi^2}$$

using the trigonometric inequalities (4).

In order to prove the second part of the theorem, we now find an *upper* bound on expression (6). First, it is clear that $\sin^2(\pi N \delta(z)) \leq 1$ always. For the denominator, by the same argument to above we have $\sin(\pi \delta(z)) \geq 2\delta(z)$ and hence, for all z ,

$$\Pr[\text{get outcome } z] \leq \frac{1}{N^2} \left(\frac{1}{2\delta(z)} \right)^2 = \frac{1}{4N^2 \delta(z)^2}.$$

We now sum this expression over all z such that $|\delta(z)| \geq \epsilon$. The sum is symmetric about $\delta(z) = 0$, and as z is an integer, the terms in this sum corresponding to $\delta(z) > 0$ are $\delta_0, \delta_0 + 1/N, \dots$, for some $\delta_0 \geq \epsilon$. The sum will be maximised when $\delta_0 = \epsilon$, when we obtain

$$\begin{aligned} \Pr[\text{get outcome } z \text{ with } |\delta(z)| \geq \epsilon] &\leq \frac{1}{4N^2} \sum_{k=0}^{\infty} \frac{1}{(\epsilon + k/N)^2} \leq \frac{1}{4} \int_0^{\infty} \frac{1}{(N\epsilon + k)^2} dk \\ &= \frac{1}{4} \int_{N\epsilon}^{\infty} \frac{1}{k^2} dk = O\left(\frac{1}{N\epsilon}\right). \end{aligned}$$

□

We observe the following points regarding the behaviour of this algorithm.

- What happens if we do not know an eigenvector of U ? If we input an arbitrary state $|\varphi\rangle$ to the phase estimation algorithm, we can write it as a superposition $|\varphi\rangle = \sum_j \alpha_j |\psi_j\rangle$ over eigenvectors $\{|\psi_j\rangle\}$. Therefore, the algorithm will output an estimate of each corresponding eigenvalue ϕ_j with probability $|\alpha_j|^2$. This may or may not allow us to infer anything useful, depending on what we know about U in advance.
- In order to approximate ϕ to n bits of precision, we needed to apply the operator U^{2^m} , for all $0 \leq m \leq n - 1$. If we are given U as a black box, this may be prohibitively expensive as we need to use the black box exponentially many times in n . However, if we have an explicit circuit for U , we may be able to find a more efficient way of computing U^{2^m} . An example of this is modular exponentiation, where we can efficiently perform repeated squaring.

6.1 Application to quantum counting

An elegant application of phase estimation is to a generalisation of the unstructured (Grover) search problem. Imagine we have an oracle $f : \{0, 1\}^n \rightarrow \{0, 1\}$ which takes the value 1 on k inputs, for some unknown k , and again set $N = 2^n$. We would like to estimate k by querying f .

Classically, a natural way to do this is by sampling. Imagine that we query f on q random inputs and get that f is 1 on ℓ of those inputs. Then as our estimate of k we output $\tilde{k} = \ell N/q$. One can show using properties of the binomial distribution that this achieves

$$|\tilde{k} - k| = O\left(\sqrt{\frac{k(N-k)}{q}}\right)$$

with high probability. We can achieve improved accuracy by using the phase estimation algorithm. Consider the ‘‘Grover iteration’’ $G = -H^{\otimes n}U_0H^{\otimes n}U_f$. As G is a rotation through angle 2θ in a 2-dimensional plane, where θ satisfies $\sin\theta = \sqrt{k/N}$, its eigenvalues are $e^{2i\theta}$ and $e^{-2i\theta}$. In order to estimate k , we can apply the phase estimation algorithm to G to estimate either one of these eigenvalues. As it does not matter which we estimate, we can input any state within this 2-dimensional plane to the phase estimation algorithm as a claimed eigenvector of G . In particular, the state $\frac{1}{\sqrt{N}} \sum_{x \in \{0,1\}^n} |x\rangle$ will work.

By Theorem 6.1, if we apply the phase estimation algorithm to G , we can find the closest m -digit number to θ , for any m , with constant probability of success using $O(2^m)$ queries. For small θ , we have $\theta \approx \sqrt{k/N}$, so we learn $\sqrt{k/N}$ up to additive error $O(1/2^m)$ using $O(2^m)$ queries. Setting $2^m = \sqrt{N}/\delta$ for some real $\delta > 0$, we have learnt \sqrt{k} up to additive error $O(\delta)$ using $O(\sqrt{N}/\delta)$ queries; or in other words have learnt k up to additive error $O(\delta\sqrt{k})$ using $O(\sqrt{N}/\delta)$ queries. In order to achieve a similar level of accuracy classically, we would need $\Omega(N/\delta^2)$ queries for small k .

Another application of phase estimation, to the order finding problem, is discussed in the Exercises.

7 Hamiltonian simulation

One of the earliest – and most important – applications of a quantum computer is likely to be the simulation of quantum mechanical systems. There are quantum systems for which no efficient classical simulation is known, but which we can simulate on a universal quantum computer. What does it mean to “simulate” a physical system? According to the OED, simulation is “the technique of imitating the behaviour of some situation or process (whether economic, military, mechanical, etc.) by means of a suitably analogous situation or apparatus”. What we will take simulation to mean here is approximating the *dynamics* of a physical system. Rather than tailoring our simulator to simulate only one type of physical system (which is sometimes called *analogue* simulation), we seek a general simulation algorithm which can simulate many different types of system (sometimes called *digital* simulation).

According to the laws of quantum mechanics, time evolution of the state $|\psi\rangle$ of a quantum system is governed by Schrödinger’s equation:

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = H(t) |\psi(t)\rangle,$$

where $H(t)$ is the Hamiltonian of the system (for convenience, we will henceforth absorb \hbar into $H(t)$). An important special case on which we will focus is the *time-independent* setting where $H(t) = H$ is constant. In this case the solution of this equation is

$$|\psi(t)\rangle = e^{-iHt} |\psi(0)\rangle.$$

Given a physical system specified by some Hamiltonian H , we would like to simulate the evolution of the system on an arbitrary initial state for a certain amount of time t . In other words, given H , we would like to implement a unitary operator which approximates

$$U(t) = e^{-iHt}.$$

What does it mean to approximate a unitary? The “gold standard” of approximation is approximation in the operator norm (aka spectral norm)

$$\|A\| := \max_{|\psi\rangle \neq 0} \frac{\|A|\psi\rangle\|}{\|\psi\rangle\|},$$

where $\|\psi\rangle\| = \sqrt{\langle\psi|\psi\rangle}$ is the usual Euclidean norm of $|\psi\rangle$. Note that this is indeed a norm, and in particular satisfies the triangle inequality $\|A + B\| \leq \|A\| + \|B\|$. We say that \tilde{U} approximates U to within ϵ if

$$\|\tilde{U} - U\| \leq \epsilon.$$

This is a natural definition of approximation because it implies that, for any state $|\psi\rangle$, $\tilde{U}|\psi\rangle$ and $U|\psi\rangle$ are only distance at most ϵ apart.

7.1 Simulation of k -local Hamiltonians

For simplicity, assume that H is a Hamiltonian of n two-level systems (qubits). In order for our quantum simulation of H to be efficient, we need $U = e^{-iHt}$ to be approximable by a quantum circuit containing $\text{poly}(n)$ gates. A fairly straightforward counting argument shows that not all

Hamiltonians H can be simulated efficiently. However, it turns out that several important physically motivated classes can indeed be simulated. Perhaps the most important of these is k -local Hamiltonians.

A Hamiltonian H of n qubits is said to be k -local if it can be written as a sum

$$H = \sum_{j=1}^m H_j$$

for some m , where each H_j is a Hermitian matrix which acts non-trivially on at most k qubits. That is, H_j is the tensor product of a matrix H'_j on k qubits, and the identity matrix on the remaining $n - k$ qubits. For example, the operator on 3 qubits

$$H = X \otimes I \otimes I - 2I \otimes Z \otimes Y$$

is 2-local. Many interesting physical systems are k -local for small k (say $k \leq 3$), some of which you may have heard of. Simple examples include the two-dimensional Ising model on a $n \times n$ square lattice,

$$H = J \sum_{i,j=1}^n Z^{(i,j)} Z^{(i,j+1)} + Z^{(i,j)} Z^{(i+1,j)}$$

and the Heisenberg model on a line,

$$H = \sum_{i=1}^n J_x X^{(i)} X^{(i+1)} + J_y Y^{(i)} Y^{(i+1)} + J_z Z^{(i)} Z^{(i+1)},$$

both of which are used in the study of magnetism (in the above, $M^{(j)}$ denotes a single qubit operator acting on the j 'th qubit, and J, J_x, J_y, J_z are constants).

Note that, if H is k -local, we can assume that $m \leq \binom{n}{k} = O(n^k)$. We usually assume that k is constant, in which case m is polynomial in n . We first show that each of the individual H_j operators can be simulated efficiently, which will be immediate from the following theorem, which formalises a claim made at the start of the course.

Theorem 7.1 (Solovay-Kitaev theorem). *Let U be a unitary operator which acts non-trivially on $k = O(1)$ qubits, and let S be an arbitrary universal set of quantum gates. Then U can be approximated in the operator norm to within ϵ using $O(\log^c(1/\epsilon))$ gates from S , for some $c < 4$.*

Proof. Sadly beyond the scope of this course. For a readable explanation, see Andrew Childs' lecture notes. \square

As each $e^{-iH_j t}$ acts non-trivially on only at most k qubits, it follows from the Solovay-Kitaev theorem that we can approximate each of these operators individually to within ϵ in time $O(\text{polylog}(1/\epsilon))$. In the special case where all of the H_j operators commute, we have

$$e^{-iHt} = e^{-i(\sum_{j=1}^m H_j)t} = \prod_{j=1}^m e^{-iH_j t}.$$

Thus a natural way to find a unitary operator approximating e^{-iHt} is to take the product of our approximations of $e^{-iH_1 t}, \dots, e^{-iH_m t}$. Although each of these approximates $e^{-iH_j t}$ to within ϵ , this does not imply that their product approximates e^{-iHt} to within ϵ . However, we now show that the approximation error only scales linearly.

Lemma 7.2. *Let $(U_i), (V_i)$ be sequences of m unitary operators satisfying $\|U_i - V_i\| \leq \epsilon$ for all $1 \leq i \leq m$. Then $\|U_m \dots U_1 - V_m \dots V_1\| \leq m\epsilon$.*

Proof. The proof is by induction on m . The claim trivially holds for $m = 1$. Assuming that it holds for a given m , we have

$$\begin{aligned}
& \|U_{m+1}U_m \dots U_1 - V_{m+1}V_m \dots V_1\| \\
&= \|U_{m+1}U_m \dots U_1 - U_{m+1}V_m \dots V_1 + U_{m+1}V_m \dots V_1 - V_{m+1}V_m \dots V_1\| \\
&\leq \|U_{m+1}U_m \dots U_1 - U_{m+1}V_m \dots V_1\| + \|U_{m+1}V_m \dots V_1 - V_{m+1}V_m \dots V_1\| \\
&= \|U_{m+1}(U_m \dots U_1 - V_m \dots V_1)\| + \|(U_{m+1} - V_{m+1})V_m \dots V_1\| \\
&= \|U_m \dots U_1 - V_m \dots V_1\| + \|U_{m+1} - V_{m+1}\| \\
&\leq (m+1)\epsilon.
\end{aligned}$$

□

Thus, in order to approximate $\prod_{j=1}^m e^{-iH_j t}$ to within ϵ , it suffices to approximate each of the H_j to within ϵ/m . We formalise this as the following proposition.

Proposition 7.3. *Let H be a Hamiltonian which can be written as the sum of m commuting terms H_j , each acting non-trivially on $k = O(1)$ qubits. Then, for any t , there exists a quantum circuit which approximates the operator e^{-iHt} to within ϵ in time $O(m \text{ polylog}(m/\epsilon))$.*

7.2 The non-commuting case

Unfortunately, this simulation technique does *not* necessarily work for non-commuting H_j . The reason is that if A and B are non-commuting operators, it need not hold that $e^{-i(A+B)t} = e^{-iAt}e^{-iBt}$. However, we can simulate non-commuting Hamiltonians via an observation known as the Lie-Trotter product formula.

In what follows, the notation $X + O(\epsilon)$, for a matrix X , is used as shorthand for $X + E$, where E is a matrix satisfying $\|E\| \leq C\epsilon$, for some universal constant C (not depending on X or ϵ).

Lemma 7.4 (Lie-Trotter product formula). *Let A and B be Hermitian matrices such that $\|A\| \leq K$ and $\|B\| \leq K$, for some real $K \leq 1$. Then*

$$e^{-iA}e^{-iB} = e^{-i(A+B)} + O(K^2).$$

Proof. From the Taylor series for e^x , for any matrix A such that $\|A\| = K \leq 1$, we have

$$e^{-iA} = I - iA + \sum_{k=2}^{\infty} \frac{(-iA)^k}{k!} = I - iA + (-iA)^2 \sum_{k=0}^{\infty} \frac{(-iA)^k}{(k+2)!} = I - iA + O(K^2).$$

Hence

$$e^{-iA}e^{-iB} = (I - iA + O(K^2))(I - iB + O(K^2)) = I - iA - iB + O(K^2) = e^{-i(A+B)} + O(K^2).$$

□

Applying this formula multiple times, for any Hermitian matrices H_1, \dots, H_m satisfying $\|H_j\| \leq K \leq 1$ for all j ,

$$\begin{aligned}
e^{-iH_1}e^{-iH_2} \dots e^{-iH_m} &= \left(e^{-i(H_1+H_2)} + O(K^2) \right) e^{-iH_3} \dots e^{-iH_m} \\
&= \left(e^{-i(H_1+H_2+H_3)} + O((2K)^2) \right) e^{-iH_4} \dots e^{-iH_m} + O(K^2) \\
&= e^{-i(H_1+\dots+H_m)} + O(K^2) + O((2K)^2) + \dots + O(((m-1)K)^2) \\
&= e^{-i(H_1+\dots+H_m)} + O(m^3K^2).
\end{aligned}$$

Therefore, there is a universal constant C such that if $n \geq Cm^3(Kt)^2/\epsilon$,

$$\left\| e^{-iH_1t/n}e^{-iH_2t/n} \dots e^{-iH_mt/n} - e^{-i(H_1+\dots+H_m)t/n} \right\| \leq \epsilon/n.$$

By Lemma 7.2, for any such n

$$\left\| \left(e^{-iH_1t/n}e^{-iH_2t/n} \dots e^{-iH_mt/n} \right)^n - e^{-i(H_1+\dots+H_m)t} \right\| \leq \epsilon.$$

Given this result, we can simulate a k -local Hamiltonian simply by simulating the evolution of each term for time t/n to high enough accuracy and concatenating the individual simulations. We formalise this as the following theorem.

Theorem 7.5. *Let H be a Hamiltonian which can be written as the sum of m terms H_j , each acting non-trivially on $k = O(1)$ qubits and satisfying $\|H_j\| \leq K$ for some K . Then, for any t , there exists a quantum circuit which approximates the operator e^{-iHt} to within ϵ in time $O(m^3(Kt)^2/\epsilon)$, up to polylogarithmic factors.*

It seems somewhat undesirable that, in order to simulate a Hamiltonian for time t , this algorithm has dependence on t which is $O(t^2)$. In fact, using more complicated simulation techniques, the overall complexity in Theorem 7.5 can be improved to time $O(mKt)$, up to polylogarithmic factors.

8 Quantum walk

We now discuss a quantum generalisation of the classical concept of random walk. Classically, random walks are an important algorithmic tool, and the same has proven to be true for quantum walks. We begin by introducing the most basic variant of quantum walk, walk on the line.

8.1 Quantum walk on the line

Classically, the simple random walk on the line is defined as follows. A particle (“walker”) is placed on an infinite line at position 0. At each time step, the walker flips an unbiased coin. If the result is heads, it moves left by 1; otherwise, it moves right by 1. The probability of being found at position x after t steps is exactly

$$\frac{1}{2^t} \binom{t}{\frac{t+x}{2}},$$

where we define $\binom{t}{r} = 0$ for non-integer r . (This is easily derived as follows: there are 2^t different paths of n steps that could be taken; the paths which end up at x are exactly those with $(t+x)/2$ right-moving steps; and there are $\binom{t}{(t+x)/2}$ such paths.) We can see from this that the probability decays quickly for $|x|$ outside the range $O(\sqrt{t})$, and indeed one can calculate the variance of the position p as

$$\mathbb{E}[p^2] = \mathbb{E} \left[\left(\sum_{i=1}^t Z_i \right)^2 \right] = \sum_{i,j=1}^t \mathbb{E}[Z_i Z_j] = \sum_{i=1}^t \mathbb{E}[Z_i^2] = t,$$

where the $Z_i \in \{\pm 1\}$ are random variables determining whether the i 'th step is to the left or the right. We look for a quantum generalisation of this simple process.

One natural generalisation is as follows. Consider a quantum system with two registers $|x\rangle|c\rangle$, where the first holds an integer position¹ x and the second holds a coin state $c \in \{L, R\}$. As in the classical case, at each step our quantum walk will flip a coin and then decide in which direction to move. These two operations will be unitary: a coin operator C , and a shift operator S . The coin operator acts solely on the coin register, and consists of a Hadamard operation:

$$C|L\rangle = \frac{1}{\sqrt{2}} (|L\rangle + |R\rangle), \quad C|R\rangle = \frac{1}{\sqrt{2}} (|L\rangle - |R\rangle).$$

The shift operator acts on both registers, and simply moves the walker in the direction indicated by the coin state:

$$S|x\rangle|L\rangle = |x-1\rangle|L\rangle, \quad S|x\rangle|R\rangle = |x+1\rangle|R\rangle.$$

Then a quantum walk on the line for t steps consists of applying the unitary operator $(S(I \otimes C))^t$ to some initial state. Remarkably, these simple dynamics can lead to some fairly complicated results. Consider the first few steps of a quantum walk with initial state $|0\rangle|L\rangle$ (position 0, facing left).

¹The reader might object to a register that stores an arbitrary integer. In this case, consider it to store an integer mod m for some large m .

One can calculate that the state evolves as follows.

$$\begin{aligned}
|0\rangle|L\rangle &\mapsto \frac{1}{\sqrt{2}}(|-1\rangle|L\rangle + |1\rangle|R\rangle) \\
&\mapsto \frac{1}{2}(|-2\rangle|L\rangle + |0\rangle|R\rangle + |0\rangle|L\rangle - |2\rangle|R\rangle) \\
&\mapsto \frac{1}{2\sqrt{2}}(|-3\rangle|L\rangle + |-1\rangle|R\rangle + 2|-1\rangle|L\rangle - |1\rangle|L\rangle + |3\rangle|R\rangle) \\
&\mapsto \dots
\end{aligned}$$

Consider the result of measuring the position register after the third step. Positions -3 , 1 and 3 are obtained with probability $1/8$ each, and position -1 is obtained with probability $5/8$. In other words, the most likely position for the particle to be found in is -1 . By contrast, the classical random walk is symmetric about 0 (and in fact is found in position -3 or 3 with probability $1/8$ each, and -1 and 1 with probability $3/8$ each). The bias of the quantum walk is an effect of interference. If the quantum walk is run for more steps before measuring the position, we obtain the pattern illustrated in Figure 12.




Figure 12: The distribution of positions following a quantum walk on the line for 100 steps using the Hadamard coin. Only even positions are shown as the amplitude at odd positions is 0.

Note that, unlike the classical walk, the quantum walk is not symmetric about 0 . Intuitively, this is caused by the Hadamard coin treating the L and R directions differently: only the $|R\rangle$ state gets mapped to a state containing a negative amplitude, leading to destructive interference in this direction. This asymmetry can be removed by changing the initial state of the coin register to $\frac{1}{\sqrt{2}}|0\rangle(|L\rangle + i|R\rangle)$, or alternatively by using a different coin operator. Figure 13 compares the distribution of probabilities obtained for classical and (symmetric) quantum walks.

The quantum walk seems to spread out from the origin faster than the classical walk. Indeed, it can be shown that the variance is $\Omega(t^2)$, noticeably faster than the classical $O(t)$. Unlike the straightforward analysis of the classical random walk, the effect of interference means that it is quite involved to prove this, so we will not do so here. This difference between quantum and classical walks will become more pronounced when we consider more general graphs.




Figure 13: Quantum walk on the line for 100 steps with starting coin state $\frac{1}{\sqrt{2}}|0\rangle(|L\rangle + i|R\rangle)$, compared with classical walk for 100 steps (red dashed line). Only even positions are shown as the amplitude at odd positions is 0.

8.2 Quantum walk on general graphs

There is a natural generalisation of the classical random walk on the line to a random walk on an arbitrary graph G with m vertices. The walker is positioned at a vertex of G , and at each time step, it chooses an adjacent vertex to move to, uniformly at random. Here we will consider only *undirected* graphs where the ability to move from A to B implies the ability to move from B to A. Further, we restrict to *regular* graphs for simplicity, i.e. those whose every vertex has degree d . Labelling vertices by integers between 1 and m , the probability of being at vertex j after t steps, given that the walk started at vertex i , can then be written in compact form as $\langle j|M^t|i\rangle$ for some matrix M , where

$$M_{ij} = \begin{cases} \frac{1}{d} & \text{if } i \text{ is connected to } j \\ 0 & \text{otherwise.} \end{cases}$$

Random walks on graphs have been intensively studied in computer science for their many algorithmic applications, as well as for their intrinsic mathematical interest. For example, one of the most efficient (and simplest) algorithms known for the fundamental problem of boolean satisfiability (SAT) is based on a random walk.

We now consider how a quantum generalisation of this process can be found. We still have position and coin registers, but now the position register is m -dimensional and the coin register is d -dimensional. Label each vertex with a distinct integer between 1 and m , and for each vertex, label its outgoing edges with distinct integers between 1 and d such that, for each i , the edges labelled with i form a cycle. For each vertex $v \in \{1, \dots, m\}$, let $N(v, i)$ denote the i 'th neighbour of v (i.e. the vertex at the other end of the i 'th edge).

Then our quantum walk will once again consist of alternating shift and coin operators S and C , i.e. each step is of the form $(S(I \otimes C))$. The shift operator simply performs the map

$$S|v\rangle|i\rangle = |N(v, i)\rangle|i\rangle,$$

and the coin operator C once again acts only on the coin register. However, as this is now d -dimensional, we have many possible choices for C . One reasonable choice is the so-called *Grover*

coin,

$$C = \begin{pmatrix} \frac{2}{d} - 1 & \frac{2}{d} & \cdots & \frac{2}{d} \\ \frac{2}{d} & \frac{2}{d} - 1 & \cdots & \frac{2}{d} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{2}{d} & \frac{2}{d} & \cdots & \frac{2}{d} - 1 \end{pmatrix}.$$

This is just the iteration used in Grover's algorithm. This operator is an appealing choice because it is permutation-symmetric (i.e. treats all edges equally), and it is far away from the identity matrix (i.e. has a large mixing effect). If $d = 2$, we would get $C = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, so in this case the coins used earlier for the walk on the line lead to more interesting behaviour.

8.3 Exponentially faster hitting on the hypercube

We now focus on one particularly interesting graph: the n -dimensional hypercube (aka the Cayley graph of the group \mathbb{Z}_2^n). This is the graph whose vertices are n -bit strings which are adjacent if they differ in exactly one bit. Each edge $x \leftrightarrow y$ of this graph is naturally labelled by the index of the single bit where x and y differ. We will be interested in the expected time it takes for a random walk on this graph to travel from the “all zeroes” string 0^n to the “all ones” string 1^n , i.e. to traverse the graph from one extremity to the other, which is known as the *hitting time* from 0^n to 1^n .

Classically, this time can be analysed by mapping the walk to a (biased) random walk on the line. Imagine the walker is currently at a vertex with Hamming weight k . The probability of moving to a vertex with Hamming weight $(k - 1)$ is k/n , and the probability of moving to a vertex with Hamming weight $(k + 1)$ is $1 - k/n$. Observe that, as k increases, the probability of a step leading to the Hamming weight increasing decreases, so intuitively the walker becomes “stuck” in the “middle” of the graph (i.e. near Hamming weight $n/2$).

More rigorously, we have the following proposition.

Proposition 8.1. *The hitting time from 0^n to 1^n is at least $2^n - 1$.*

Proof. Let $h(k)$ be the expected number of steps until the walk hits 1^n , given that it starts with Hamming weight k . We have the recurrence

$$h(k) = \frac{k}{n}h(k-1) + \left(1 - \frac{k}{n}\right)h(k+1) + 1,$$

with the boundary case $h(n) = 0$. Writing $\delta(k) = h(k) - h(k+1)$, this can be simplified to

$$\delta(k) = \left(\frac{n}{k+1} - 1\right)\delta(k+1) - \frac{n}{k+1}$$

and the recurrence solved to give

$$\delta(k) = \frac{1}{\binom{n-1}{k}} \sum_{j=0}^k \binom{n}{j}, \text{ and hence } h(0) = \sum_{k=0}^{n-1} \delta(n) = \sum_{k=0}^{n-1} \frac{1}{\binom{n-1}{k}} \sum_{j=0}^k \binom{n}{j}.$$

Thus $h(0)$ can be lower bounded by rearranging the sum to give

$$h(0) = \sum_{j=0}^{n-1} \binom{n}{j} \sum_{k=j}^{n-1} \frac{1}{\binom{n-1}{k}} \geq \sum_{j=0}^{n-1} \binom{n}{j} = 2^n - 1.$$

□

We thus see that the expected time to reach the 1^n vertex is exponential in n . However, for quantum walks the situation is very different, and we have the following result.

Theorem 8.2. *If a quantum walk on the hypercube is performed for $T \approx \frac{\pi}{2}n$ steps starting in position 0^n , and the position register is measured, the outcome 1^n is obtained with probability $1 - O(\text{polylog}(n)/n)$.*

Once again, the proof of this result is too technical to include here. However, we can give a sketch of the first part of one proof strategy, which is analogous to the classical case, and consists of simplifying to a walk on the line. Define a set of $2n$ states $\{|v_k, L\rangle, |v_k, R\rangle\}$ indexed by an integer $k = 0, \dots, n$ as follows:

$$|v_k, L\rangle := \frac{1}{\sqrt{k \binom{n}{k}}} \sum_{x, |x|=k} \sum_{i, x_i=1} |x\rangle |i\rangle, \quad |v_k, R\rangle := \frac{1}{\sqrt{(n-k) \binom{n}{k}}} \sum_{x, |x|=k} \sum_{i, x_i=0} |x\rangle |i\rangle,$$

with the exception of the special cases $|v_0, L\rangle$ and $|v_n, R\rangle$, which will not be used and are undefined. Now observe that the quantum walk on the hypercube preserves the subspace spanned by this set of states:

$$S|v_k, L\rangle = \frac{1}{\sqrt{k \binom{n}{k}}} \sum_{x, |x|=k} \sum_{i, x_i=1} |x \oplus e_i\rangle |i\rangle = \frac{1}{\sqrt{k \binom{n}{k}}} \sum_{x, |x|=k-1} \sum_{i, x_i=0} |x\rangle |i\rangle = |v_{k-1}, R\rangle,$$

and similarly $S|v_k, R\rangle = |v_{k+1}, L\rangle$. In the case of the coin operator, it turns out that

$$\begin{aligned} (I \otimes C)|v_k, L\rangle &= \left(\frac{2k}{n} - 1\right) |v_k, L\rangle + \frac{2\sqrt{k(n-k)}}{n} |v_k, R\rangle \\ \text{and } (I \otimes C)|v_k, R\rangle &= \frac{2\sqrt{k(n-k)}}{n} |v_k, L\rangle + \left(1 - \frac{2k}{n}\right) |v_k, R\rangle. \end{aligned}$$

This behaviour is similar to the quantum walk on the line, with two differences: first, the direction in which the walker is moving flips with each shift, and second, the coin is different at each position (i.e. depends on k). Based on this reduction, it is easy to plot the behaviour of this quantum walk numerically for small n , as shown in Figure 14.

8.4 A general approach for quantising Markov chains

The method described in the previous section, suitably generalised, allows us to define a quantum walk on any undirected graph. However, we now describe a different approach, due to Szegedy, which simultaneously allows us to handle directed graphs, and allows us to analyse the properties of the quantum walk produced.

In the most general setting, a random walk on a directed graph G with n vertices is described by an $N \times N$ matrix P with non-negative entries such that $\sum_{i=1}^N P_{ij} = 1$ for all j (such matrices are called *stochastic*). P_{ij} is the probability of moving to vertex i from vertex j . The quantum walk corresponding to P is a unitary operation W_P on $\mathbb{C}^N \otimes \mathbb{C}^N$. We can think of this as operating on a space spanned by edges of the graph G . To define W_P , we first introduce states

$$|\psi_j\rangle := \sum_{i=1}^N \sqrt{P_{ij}} |j\rangle |i\rangle$$




Figure 14: Quantum walk on a 50-dimensional hypercube for 60 steps starting in state $|\nu_0, R\rangle$, compared with classical walk for 1000 steps (red dashed line). Graph plots probability of being at a point with Hamming weight k . Only even positions are shown as the amplitude at odd positions is 0.

for $j = 1, \dots, N$. As P is stochastic, these are all normalised states. Write

$$\Pi := \sum_{j=1}^N |\psi_j\rangle\langle\psi_j|$$

for the projector onto the subspace spanned by $\{|\psi_j\rangle\}$, and let $S = \sum_{i,j=1}^N |ij\rangle\langle ji|$ be the operator that swaps the two registers. Then a step of the quantum walk is the unitary operator

$$W_P = S(2\Pi - I).$$

$2\Pi - I$ is a reflection about the subspace spanned by $\{|\psi_j\rangle\}$. If P is a random walk on a regular graph with uniform probability of moving to each neighbour, following this procedure gives a quantum walk as described in the previous section using the Grover coin. We now analyse the eigenvalues of W_P .

Theorem 8.3. *For any $N \times N$ stochastic matrix P , let D be the matrix defined by $D_{ij} = \sqrt{P_{ij}P_{ji}}$. Then, for each eigenvalue λ of D , W_P has eigenvalues $\lambda \pm i\sqrt{1 - \lambda^2} = e^{\pm i \arccos \lambda}$. W_P also has eigenvalues ± 1 .*

Proof. Set $T := \sum_{j=1}^N |\psi_j\rangle\langle j|$, and for each eigenvector $|\lambda\rangle$ of D , let $|\tilde{\lambda}\rangle = T|\lambda\rangle$. T is an isometry:

$$T^\dagger T = \sum_{j,k=1}^N |j\rangle\langle\psi_j|\langle\psi_k|\langle k| = \sum_{j,k,\ell,m=1}^N \sqrt{P_{\ell j}P_{mk}} \langle j|k\rangle\langle\ell|m\rangle|j\rangle\langle k| = \sum_{j,\ell=1}^N P_{\ell j}|j\rangle\langle j| = I.$$

On the other hand,

$$TT^\dagger = \sum_{j,k=1}^N |\psi_j\rangle\langle j|k\rangle\langle\psi_k| = \sum_{j=1}^N |\psi_j\rangle\langle\psi_j| = \Pi$$

and finally

$$T^\dagger S T = \sum_{j,k=1}^N |j\rangle\langle\psi_j|S|\psi_k\rangle\langle k| = \sum_{j,k,\ell,m=1}^N \sqrt{P_{\ell j}P_{mk}} \langle j,\ell|S|k,m\rangle|j\rangle\langle k| = \sum_{j,k=1}^N \sqrt{P_{jk}P_{kj}}|j\rangle\langle k| = D.$$

If we apply W_P to $|\tilde{\lambda}\rangle$, we get

$$W_P|\tilde{\lambda}\rangle = S(2\Pi - I)|\tilde{\lambda}\rangle = S(2TT^\dagger - I)T|\lambda\rangle = S(2T - T)|\lambda\rangle = S|\tilde{\lambda}\rangle,$$

and if we apply W_P again we get

$$W_P S|\tilde{\lambda}\rangle = S(2TT^\dagger - I)ST|\lambda\rangle = (2STD - T)|\lambda\rangle = (2\lambda S - I)|\tilde{\lambda}\rangle.$$

So W_P preserves the subspace spanned by $\{|\tilde{\lambda}\rangle, S|\tilde{\lambda}\rangle\}$, implying that there is an eigenvector of W_P within this subspace. Let

$$|\mu\rangle = |\tilde{\lambda}\rangle - \nu S|\tilde{\lambda}\rangle$$

for some $\mu \in \mathbb{C}$ chosen such that $|\mu\rangle$ is indeed an eigenvector of W_P . Then

$$W_P|\mu\rangle = S|\tilde{\lambda}\rangle - \mu(2\lambda S - I)|\tilde{\lambda}\rangle = \mu|\tilde{\lambda}\rangle + (1 - 2\lambda\mu)S|\tilde{\lambda}\rangle,$$

so for $|\mu\rangle$ to be an eigenvector we must have

$$(1 - 2\lambda\mu) = \mu(-\mu), \text{ so } \mu^2 - 2\lambda\mu + 1 = 0$$

and hence $\mu = \lambda \pm i\sqrt{1 - \lambda^2}$. Finally, for any vector orthogonal to all the states $\{|\psi_j\rangle\}$, W_P simply acts as $-S$, which has eigenvalues ± 1 . \square

The matrix D occurring in this theorem is like a ‘‘symmetrised’’ version of P . If P is already symmetric, $D = P$.

8.5 Search by quantum walk

A natural use for random, or quantum walks, on graphs is search. Imagine that a set S of k vertices are marked. A natural approach to find a marked vertex is to start with some arbitrary vertex, then perform a random walk according to P until we find a marked vertex, when we stop. This process can be modelled by modifying the initial matrix P to produce a new matrix P' where

$$P'_{ij} = \begin{cases} 1 & \text{if } i = j \text{ and } j \in S \\ 0 & \text{if } i \neq j \text{ and } j \in S \\ P_{ij} & \text{otherwise.} \end{cases}$$

Assume that P is symmetric and let P_S be the matrix obtained by deleting all the rows and columns of P corresponding to vertices in S . P 's largest eigenvalue is 1; let the second largest, in absolute value, be $1 - \delta$. δ is known as the spectral gap of P .

We first calculate a bound on the largest eigenvalue of P_S .

Lemma 8.4. *If the second largest eigenvalue of P , in absolute value, is at most $1 - \delta$ and $k = \epsilon N$, then $\|P_S\| \leq 1 - \delta\epsilon$.*

Proof. Consider the principal eigenvector $|v\rangle \in \mathbb{R}^{N-k}$ of P_S , and let $|w\rangle \in \mathbb{R}^N$ be the vector produced by adding zeroes at positions corresponding to the marked vertices. We perform the eigendecomposition of P ,

$$P = \sum_{i=1}^N \lambda_i |\lambda_i\rangle \langle \lambda_i|$$

where $\lambda_1 = 1$, and compute

$$\begin{aligned}
\langle v|P_S|v\rangle &= \langle w|P|w\rangle = \sum_{i=1}^N \lambda_i |\langle \lambda_i|w\rangle|^2 = |\langle \lambda_1|w\rangle|^2 + \sum_{i=2}^N \lambda_i |\langle \lambda_i|w\rangle|^2 \\
&\leq |\langle \lambda_1|w\rangle|^2 + (1 - \delta) \sum_{i=2}^N |\langle \lambda_i|w\rangle|^2 \\
&= 1 - \delta \sum_{i=2}^N |\langle \lambda_i|w\rangle|^2 \\
&= 1 - \delta(1 - |\langle \lambda_1|w\rangle|^2).
\end{aligned}$$

Observe that, as P is symmetric, $|\lambda_1\rangle$ must actually be the uniform superposition $|\lambda_1\rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^N |i\rangle$. Then, if we expand $|w\rangle = \sum_{i \notin S} \alpha_i |i\rangle$, we have

$$|\langle \lambda_1|w\rangle|^2 = \frac{1}{N} \left| \sum_{i \notin S} \alpha_i \right|^2 \leq \frac{N - |S|}{N} \left(\sum_{i \notin S} |\alpha_i|^2 \right) = \frac{N - k}{N} = 1 - \epsilon,$$

where the inequality is Cauchy-Schwarz. Hence $\|P_S\| \leq 1 - \delta\epsilon$ as claimed. \square

We now describe an algorithm which uses a quantum walk to solve the more restricted problem of distinguishing between the case where there are no marked vertices, and the case where there are ϵN marked vertices. The algorithm applies phase estimation to the operator $W_{P'}$ using as input the state

$$|\psi\rangle = \frac{1}{\sqrt{N - k}} \sum_{j \notin S} |\psi_j\rangle,$$

the uniform superposition over non-marked vertices. This state can be prepared by starting with the uniform superposition over all vertices, then applying the isometry T , and measuring whether the first register corresponds to a marked vertex. If so, we are done; if not, we have prepared $|\psi\rangle$. Assuming without loss of generality that the last k elements of $[N]$ are marked, the matrix D corresponding to the walk P' is

$$D = \begin{pmatrix} P_S & 0 \\ 0 & I \end{pmatrix},$$

so by Theorem 8.3 the eigenvalues of $W_{P'}$ are ± 1 and $e^{\pm i \arccos \lambda}$ for each eigenvalue λ of P_S . If $S = \emptyset$, then $P_S = P' = P$ and $|\psi\rangle$ is an eigenvector of W_P with eigenvalue 1. So phase estimation will return a phase of 0. On the other hand, if $S \neq \emptyset$, $|\psi\rangle$ is contained entirely within the subspace corresponding to the P_S block, in which $W_{P'}$ has eigenvalues $e^{\pm i \arccos \lambda}$. We have

$$\cos \theta \geq 1 - \theta^2/2$$

for any θ , so $\arccos \lambda \geq \sqrt{2(1 - \lambda)}$. By Lemma 8.4, the phase of each eigenvalue is at least $\sqrt{2\delta\epsilon}$ in absolute value.

So, using phase estimation, it suffices to make $O(1/\sqrt{\delta\epsilon})$ steps of the quantum walk to distinguish between the two cases that (a) there are no marked elements, and (b) there are $k = \epsilon N$ marked elements. One might wonder what the point of all this work was, when Grover's algorithm solves the unstructured search problem with an ϵ fraction of marked elements using $O(1/\sqrt{\epsilon})$ queries, which is faster. The point is that, unlike Grover's algorithm, the quantum walk algorithm

is based around making local moves which respect the structure of the graph. If the data within which we are searching are laid out physically in a particular graph structure, a search algorithm which respects this structure may be more efficient. Even if there is no such physical restriction, bearing a notion of locality in mind can sometimes lead to more efficient algorithms. We will now see an example of this.

8.6 Element distinctness

Given query access to a function $f : \{1, \dots, N\} \rightarrow X$ for some set X , the element distinctness problem is to determine whether there exists a pair (x, y) such that $f(x) = f(y)$. Classically, solving this problem requires $\Omega(N)$ queries to f , as it is at least as hard as unstructured search. Indeed, if we are given the additional promise that any collision is of the form $(1, x)$ for some $x > 1$, then solving element distinctness is equivalent to determining whether $f(x) = f(1)$ for $x \in \{2, \dots, N - 1\}$, which is the unstructured search problem on $N - 1$ elements.

We will now describe a quantum algorithm due to Ambainis which solves the element distinctness problem using $O(N^{2/3})$ queries to f , which is known to be optimal.

The algorithm is based around a quantum walk on the graph whose vertices are subsets of $\{1, \dots, N\}$ of size M , which is called the Johnson graph. Here we will actually consider the graph with N^M vertices whose vertices are M -tuples made up of elements of $\{1, \dots, N\}$ (i.e. ordered subsets of $\{1, \dots, N\}$ where repetitions are allowed), which makes the analysis a bit simpler; this is called the Hamming graph. In this graph there is an edge between any pair of vertices whose corresponding M -tuples differ in exactly one element. For example, with $N = 4$, $M = 3$, there would be an edge between the vertices $(1, 3, 1)$ and $(1, 3, 2)$. A vertex is marked if its corresponding M -tuple contains a pair of duplicate elements.

During the walk, if we are at vertex $v = (v_1, \dots, v_M)$, we will maintain the values $(f(v_1), \dots, f(v_M))$. To set this up initially requires M queries to f . However, once we have created this M -tuple, to move from a vertex v to an adjacent vertex v' , differing at index i , requires only two queries: one to uncompute $f(v_i)$ and one to compute $f(v'_i)$. Checking whether we are at a marked vertex can be done using no additional queries, as we have already computed all the values $f(v_i)$ and just need to check whether there are any collisions among them.

If the elements are not all distinct, the total number of marked vertices is at least $M(M - 1)(N - 2)^{M-2}$, so the fraction of marked vertices is

$$\epsilon \geq \frac{M(M - 1)(N - 2)^{M-2}}{N^M} = \Omega\left(\frac{M^2}{N^2}\right)$$

for $M \ll N$. We now compute the spectral gap of the Hamming graph. The adjacency matrix A can be written down as

$$A = \sum_{i=1}^M C_i,$$

where C is the adjacency matrix of the $N \times N$ complete graph, and the subscript i means that it acts at position i . C has one eigenvalue N , and $N - 1$ eigenvalues -1 , and all the matrices C_i commute. So eigenvalues of A are sums of sequences of M eigenvalues of C . Overall the largest eigenvalue of A is $M(N - 1)$ (the same as the degree of the graph, as expected) and the second largest is $(M - 1)(N - 1) - 1$. After normalising by dividing by the degree, the second largest eigenvalue is

$$\frac{M - 1}{M} - \frac{1}{M(N - 1)} \leq 1 - \frac{1}{M},$$

so the spectral gap $\delta = \Omega(1/M)$. So a marked vertex (i.e. pair of equal elements) can be found using

$$M + O\left(\frac{1}{\sqrt{\delta\epsilon}}\right) = M + O\left(\frac{N}{\sqrt{M}}\right)$$

queries. To minimise this, we set $M = \Theta(N^{2/3})$ to achieve an overall complexity of $O(N^{2/3})$ queries, as claimed.

9 Noise and the framework of quantum channels

Not all processes which occur in quantum mechanics are reversible. As a very simple example, consider the operation of throwing away the input state and replacing it with the state $|0\rangle$:

$$|\psi\rangle \mapsto |0\rangle$$

for all $|\psi\rangle$. This clearly cannot be implemented as a unitary operator. Just as mixed states generalise pure states, we can generalise unitary operations to so-called completely positive trace-preserving (CPTP) maps, also known as quantum channels. These occur throughout quantum information theory and are particularly useful for describing noisy operations and *decoherence*, the bane of quantum computers.

Recall that a mixed state ρ of n qubits, which describes a probabilistic mixture of pure states $|\psi\rangle$, is a Hermitian $2^n \times 2^n$ matrix which is positive semidefinite (all its eigenvalues are non-negative) and has trace 1. What axioms would we like a physically reasonable map (“superoperator”) \mathcal{E} , which takes mixed states to mixed states, to satisfy?

1. \mathcal{E} should be linear: $\mathcal{E}(p\rho + q\sigma) = p\mathcal{E}(\rho) + q\mathcal{E}(\sigma)$ for all real p, q .
2. \mathcal{E} should be trace-preserving: $\text{tr } \mathcal{E}(\rho) = \text{tr } \rho$.
3. \mathcal{E} should preserve positivity: if $\rho \geq 0$, then $\mathcal{E}(\rho) \geq 0$. But there is a further constraint: if we apply \mathcal{E} to part of an entangled quantum state, the whole state should remain positive semidefinite. That is, $(\mathcal{E} \otimes \mathcal{I})(\rho) \geq 0$ for $\rho \geq 0$, where \mathcal{I} is the identity map on an arbitrarily large ancilla system. This constraint is called *complete positivity*.

A completely positive, trace-preserving linear map is called a quantum channel.

9.1 Representations of quantum channels

There are a number of ways in which quantum channels can be represented. We describe two here.

We first consider the **Kraus** (aka “operator-sum”) representation. Here a channel \mathcal{E} with input dimension d_{in} and output dimension d_{out} is described by a sequence of $d_{\text{out}} \times d_{\text{in}}$ matrices E_k such that

$$\sum_k E_k^\dagger E_k = I.$$

The effect of \mathcal{E} on a state ρ is then

$$\mathcal{E}(\rho) = \sum_k E_k \rho E_k^\dagger.$$

A trivial example of a channel in Kraus form is the identity channel $\mathcal{I}(\rho) = \rho$, which has one Kraus operator $E_1 = I$. For a product channel $\mathcal{E} = \mathcal{E}_1 \otimes \mathcal{E}_2$, where

$$\mathcal{E}_1(\rho) = \sum_k E_k^{(1)} \rho (E_k^{(1)})^\dagger, \quad \mathcal{E}_2(\sigma) = \sum_k E_k^{(2)} \sigma (E_k^{(2)})^\dagger$$

we have

$$\begin{aligned} \mathcal{E}(\rho \otimes \sigma) &= \mathcal{E}_1(\rho) \otimes \mathcal{E}_2(\sigma) = \left(\sum_k E_k^{(1)} \rho (E_k^{(1)})^\dagger \right) \otimes \left(\sum_\ell E_\ell^{(2)} \sigma (E_\ell^{(2)})^\dagger \right) \\ &= \sum_{k,\ell} (E_k^{(1)} \otimes E_\ell^{(2)}) (\rho \otimes \sigma) ((E_k^{(1)})^\dagger \otimes (E_\ell^{(2)})^\dagger), \end{aligned}$$

so \mathcal{E} must have Kraus operators which are tensor products of each pair of those of \mathcal{E}_1 and \mathcal{E}_2 .

We claim that any superoperator in Kraus form obeys our three axioms:

1. \mathcal{E} is linear: for any X and Y ,

$$\mathcal{E}(X + Y) = \sum_k E_k(X + Y)E_k^\dagger = \sum_k E_k X E_k^\dagger + \sum_k E_k Y E_k^\dagger = \mathcal{E}(X) + \mathcal{E}(Y).$$

2. \mathcal{E} is trace-preserving:

$$\text{tr } \mathcal{E}(\rho) = \text{tr} \left(\sum_k E_k \rho E_k^\dagger \right) = \sum_k \text{tr} \left(E_k \rho E_k^\dagger \right) = \sum_k \text{tr} \left(\rho E_k^\dagger E_k \right) = \text{tr} \left(\rho \sum_k E_k^\dagger E_k \right) = \text{tr } \rho,$$

where in the second and fourth equalities we use linearity of trace, and in the third we use its invariance under cyclic shifts.

3. \mathcal{E} is completely positive. Here it is sufficient to show that $\mathcal{E} \otimes \mathcal{I}$ maps positive operators to positive operators. Let ρ be an arbitrary positive operator on the extended system (with ancilla). Then

$$(\mathcal{E} \otimes \mathcal{I})(\rho) = \sum_k (E_k \otimes I) \rho (E_k^\dagger \otimes I).$$

For this to be positive semidefinite, it is sufficient that for any positive semidefinite ρ and any matrix M , $M\rho M^\dagger$ is positive semidefinite. This holds because we can expand ρ as a convex combination of pure states $|\psi_i\rangle$, and

$$M|\psi_i\rangle\langle\psi_i|M^\dagger = |\psi'_i\rangle\langle\psi'_i|$$

for some (unnormalised) vector $|\psi'_i\rangle$, which is positive semidefinite.

In fact, it turns out that the equivalence goes the other way too: every quantum channel can be written in Kraus form. We omit the proof (see Nielsen & Chuang, Theorem 8.1).

The Kraus representation is not unique. For example, the channels described by pairs of Kraus operators

$$E_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad E_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad \text{and} \quad F_1 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad F_2 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

have different representations, but are actually the same channel. This can be seen by writing $F_1 = (E_1 + E_2)/\sqrt{2}$, $F_2 = (E_1 - E_2)/\sqrt{2}$, and calculating

$$F_1 \rho F_1^\dagger + F_2 \rho F_2^\dagger = \frac{(E_1 + E_2)\rho(E_1^\dagger + E_2^\dagger) + (E_1 - E_2)\rho(E_1^\dagger - E_2^\dagger)}{2} = E_1 \rho E_1^\dagger + E_2 \rho E_2^\dagger.$$

Second, we describe the **Stinespring** representation. In this representation, a channel \mathcal{E} from system A to system B , $\mathcal{E} : \mathcal{B}(\mathbb{C}^{d_A}) \rightarrow \mathcal{B}(\mathbb{C}^{d_B})$, is described by an isometry from A to the pair of systems B and E , $V : \mathbb{C}^{d_A} \rightarrow \mathbb{C}^{d_B} \otimes \mathbb{C}^{d_E}$, such that

$$\mathcal{E}(\rho) = \text{tr}_E(V\rho V^\dagger).$$

That is, the channel consists of applying an isometry to ρ , mapping it into a larger-dimensional space, then tracing out part of the space. We can think of this picture as modelling the dynamics

of an open quantum system, where ρ interacts with the outside environment E , then to calculate the final state of the system we discard the environment.

Given a Kraus representation of a channel with operators E_k , we can write down a Stinespring representation in a straightforward way:

$$V = \sum_k E_k \otimes |k\rangle. \quad (7)$$

Then

$$V^\dagger V = \sum_{k,\ell} E_k^\dagger E_\ell \langle k|\ell\rangle = \sum_k E_k^\dagger E_k = I,$$

so V is indeed an isometry. Further, for any ρ ,

$$\begin{aligned} \text{tr}_E(V\rho V^\dagger) &= \sum_{k,\ell} \text{tr}_E[(E_k \otimes |k\rangle)\rho(E_\ell^\dagger \otimes \langle\ell|)] = \sum_{k,\ell} \text{tr}_E[(E_k\rho \otimes |k\rangle)(E_\ell^\dagger \otimes \langle\ell|)] \\ &= \sum_{k,\ell} (E_k\rho E_\ell^\dagger)\langle\ell|k\rangle = \sum_k E_k\rho E_k^\dagger, \end{aligned}$$

so this faithfully represents the original channel. Also observe that any isometry V can be decomposed as in Eqn. (7), so the equivalence goes both ways. As we can think of any isometry as a unitary operator on a larger space, this correspondence implies that any quantum channel can be thought of as unitary evolution of a larger system, followed by tracing out a subsystem.

9.2 Examples of quantum channels

- **Unitary evolution** is a quantum channel described by one Kraus operator: $\rho \mapsto U\rho U^\dagger$.
- Discarding ρ and **replacing it** with some state $|\psi\rangle$ is a quantum channel. If ρ is d -dimensional, we have d Kraus operators $|\psi\rangle\langle k|$, $k = 1, \dots, d$. Then, for any state ρ ,

$$\sum_{k=1}^d |\psi\rangle\langle k|\rho|k\rangle\langle\psi| = |\psi\rangle \left(\sum_{k=1}^d \langle k|\rho|k\rangle \right) \langle\psi| = (\text{tr } \rho)|\psi\rangle\langle\psi| = |\psi\rangle\langle\psi|.$$

- The **qubit depolarising channel** replaces the input state ρ with the maximally mixed state $I/2$, with probability p ; with probability $1 - p$, the input state is unchanged. That is,

$$\mathcal{E}_D(\rho) = p \frac{I}{2} + (1 - p)\rho.$$

This represents a simple form of noise: with probability p , our knowledge of the qubit is completely randomised, and with probability $1 - p$ the qubit is unharmed. The qubit depolarising channel has Kraus operators

$$E_1 = \sqrt{1 - 3p/4}I, \quad E_2 = (\sqrt{p}/2)X, \quad E_3 = (\sqrt{p}/2)Y, \quad E_4 = (\sqrt{p}/2)Z.$$

This can be generalised to d dimensions, where we have

$$\mathcal{E}_D(\rho) = p \frac{I}{d} + (1 - p)\rho.$$

- The **amplitude damping channel** models a scenario in quantum optics where a photon may be lost, with some probability. We imagine that we have a basis $\{|0\rangle, |1\rangle\}$ where 0 represents “no photon” and 1 represents “photon” (so this is a very simple example of Fock space). Then the Kraus operators of the amplitude damping channel \mathcal{E}_{AD} are

$$E_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}, \quad E_1 = \begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix}$$

for some γ , which corresponds to the probability of losing a photon. So

$$\mathcal{E}_{\text{AD}}(|0\rangle\langle 0|) = |0\rangle\langle 0|, \quad \mathcal{E}_{\text{AD}}(|1\rangle\langle 1|) = \gamma|0\rangle\langle 0| + (1-\gamma)|1\rangle\langle 1|$$

as expected.

- Every **measurement** is a quantum channel. Imagine we have a projective measurement, i.e. a set of orthogonal projectors $\{P_k\}$ such that $\sum_k P_k = I$. Then the probability that we get outcome k when we perform this measurement on state ρ is $\text{tr} P_k \rho$, and the resulting state if this occurs is

$$\rho'_k := \frac{P_k \rho P_k}{\text{tr}(P_k \rho)}.$$

The Kraus operators of this channel are just the projectors P_k . This is correct because

$$\sum_k P_k \rho P_k = \sum_k \text{tr}(P_k \rho) \frac{P_k \rho P_k}{\text{tr}(P_k \rho)};$$

the output state is a probabilistic mixture of the states ρ'_k with the correct probabilities.

If we think of a qubit as a point in \mathbb{R}^3 using the Bloch ball representation, any channel mapping a qubit to a qubit can be thought of as an affine map on this space. That is, if

$$\rho = \frac{I}{2} + \alpha_x X + \alpha_y Y + \alpha_z Z,$$

a channel corresponds to an affine map acting on the 3-dimensional vector $v = (\alpha_x, \alpha_y, \alpha_z)^T$, i.e. a map of the form

$$v \mapsto Av + b$$

for some 3×3 matrix A and 3-dimensional vector b .

9.3 Master equations

We remark that it is frequently physically reasonable to think of noise as being a continuous process parametrised by a time t . This leads to the concept of Markovian quantum channels (aka “Markov processes”). A Markovian channel is a member of a family $\{\mathcal{E}_t\}$ such that

$$\mathcal{E}_{s+t} = \mathcal{E}_s \circ \mathcal{E}_t$$

for all times s, t . This implies that

$$\mathcal{E}_t = e^{t\mathcal{L}}$$

for some superoperator \mathcal{L} . If we write $\rho(t) = \mathcal{E}_t(\rho(0))$, $\frac{d}{dt}\rho(t) = \mathcal{L}(\rho(t))$. If we demand that \mathcal{E}_t is completely positive for all times t , it can be shown that \mathcal{L} satisfies the *master equation*

$$\mathcal{L}(\rho) = i[\rho, H] + \sum_{\alpha, \beta} G_{\alpha, \beta} (F_{\alpha} \rho F_{\beta}^{\dagger} - \frac{1}{2} \{F_{\beta}^{\dagger} F_{\alpha}, \rho\}_+),$$

where H is a Hamiltonian, G and F_{α} are matrices defining a decoherence process, and $[\cdot, \cdot]$ and $\{\cdot, \cdot\}$ denote the commutator and anticommutator, respectively (with the subscript $+$ denoting the positive part of this operator).

The qubit depolarising and amplitude damping channels, for example, are Markovian. However, not all quantum channels are Markovian.

10 Quantum error-correction

Modern computer hardware is extremely reliable. Indeed, it can usually be assumed to be error-free for all intents and purposes¹. However, early quantum computing hardware is likely to be far from reliable. Even worse, efficient quantum algorithms rely on delicate quantum effects (superposition and entanglement) which must be preserved in the presence of errors. Luckily, it turns out that errors can be fought using the notion of quantum error-correcting codes. To understand these codes, it is helpful to first consider a basic classical notion of error correction.

Imagine we have a single bit x which we would like to store in a noisy storage device. A natural model for this noise is that each bit stored in the device gets flipped with probability p , for some $0 \leq p \leq 1$, and is left the same with probability $1 - p$. So if we store x in the device and then read it back later, we get the wrong answer with probability p . One way to improve this works as follows. Instead of just storing x , store the string xxx , i.e. repeat x three times. Then read out each of the bits of the (potentially corrupted) string to get $y := y_1y_2y_3$, and output 0 if the majority of the bits of y are 0, and 1 otherwise.


What is the probability of failing to output x if this strategy is followed? The wrong answer will be returned if two or more of the bits of y are flipped by noise, which will occur with probability $3p^2(1 - p) + p^3 = p^2(3 - 2p) = O(p^2)$. Thus, if p is small, this strategy has improved our resistance to noise. Indeed, for any p such that $0 < p < 1/2$, we have

$$p^2(3 - 2p) < p,$$

so the probability of error has been reduced. Another way of looking at this situation is that we have stored a bit in such a way that it is impervious to an error affecting a single bit in the storage device. The map $x \mapsto xxx$ is a very simple example of an *error correcting code* known as the binary repetition code of length 3.

10.1 Quantum errors and error-correction

We would like to find a quantum analogue of this notion of error correction. Rather than preserving classical bits x , our quantum error correcting code should preserve a qubit $|\psi\rangle$ under some notion of error. For the time being, we pretend that an error affecting one or more qubits is simply an arbitrary and unknown unitary operator N applied to those qubits. The classical bit-flip error discussed above is an example of this, as it can be seen as simply applying the operator X to a qubit in a computational basis state (recall that $X|0\rangle = |1\rangle$ and $X|1\rangle = |0\rangle$). The process of correcting errors in a qubit state $|\psi\rangle$ can be written diagrammatically as




for some unitary encoding operation E , noise operation N , and decoding operation D . In other words, we encode some qubit state $|\psi\rangle$ as a larger state $|E(\psi)\rangle$ using some ancilla qubits (initially in the state $|0^n\rangle$), some noise is applied, and later we decode the noisy encoded state to produce a state $|\psi'\rangle$. The goal is that after this process $|\psi'\rangle \approx |\psi\rangle$ for some set of correctable noise operations N .

There are two obvious ways in which the classical repetition code could be translated to the quantum regime, both of which unfortunately do not work. First, we could measure $|\psi\rangle$ in the


¹Software, of course, is another matter.

computational basis to obtain a bit 0 or 1, then just encode this with the classical repetition code. This is not suitable for quantum error correction because it does not preserve quantum coherence: if $|\psi\rangle$ is in a superposition of 0 and 1 and will be used as input to a subsequent quantum algorithm, it is necessary to preserve this superposition to see any interesting quantum effects. A second idea is that we could map $|\psi\rangle \mapsto |\psi\rangle|\psi\rangle|\psi\rangle$, by analogy with the classical code. However, this is impossible (for general $|\psi\rangle$) by the no-cloning theorem.

We therefore have to take a different approach, which will be split into two steps. In the first step, we try encoding $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ as $|E(\psi)\rangle = \alpha|000\rangle + \beta|111\rangle$. Note that this is *not* the same as the “cloning” map discussed previously. Indeed, the map $\alpha|0\rangle + \beta|1\rangle \mapsto \alpha|000\rangle + \beta|111\rangle$ can be implemented via the following simple quantum circuit.



Our decoding algorithm for this code will be based on the following quantum circuit.



Call the first three qubits the input qubits and the last two the output qubits. Following this circuit, for any basis state input $|x_1x_2x_3\rangle$, the first of the two output qubits contains $x_1 \oplus x_2$, and the second contains $x_1 \oplus x_3$. Each of these quantities is invariant under the operation of flipping all the bits of x . Thus, for any input superposition of the form $\alpha|x_1x_2x_3\rangle + \beta|x_1x_2x_3 \oplus 111\rangle$, the circuit performs the map

$$(\alpha|x_1x_2x_3\rangle + \beta|x_1x_2x_3 \oplus 111\rangle)|0\rangle|0\rangle \mapsto (\alpha|x_1x_2x_3\rangle + \beta|x_1x_2x_3 \oplus 111\rangle)|x_1 \oplus x_2\rangle|x_1 \oplus x_3\rangle.$$

This implies that, if we measure the two output qubits, we learn both $x_1 \oplus x_2$ and $x_1 \oplus x_3$ without disturbing the input quantum state. Now observe that the encoded state of $|\psi\rangle$ is always of this form, even after arbitrary bit-flip errors are applied to $|E(\psi)\rangle$:

$$\begin{aligned} |E(\psi)\rangle &= \alpha|000\rangle + \beta|111\rangle, \\ (X \otimes I \otimes I)|E(\psi)\rangle &= \alpha|100\rangle + \beta|011\rangle, \\ (X \otimes X \otimes X)|E(\psi)\rangle &= \alpha|111\rangle + \beta|000\rangle, \text{ etc.} \end{aligned}$$

The result of measuring the output qubits is known as the *syndrome*. We now consider the different syndromes we get when different noise operators N are applied to $|E(\psi)\rangle$. First, if $N = I$ (so there has been no error applied to $|E(\psi)\rangle$), we always measure 00. On the other hand, if $N = X \otimes I \otimes I$ (i.e. a bit-flip error on the first qubit) we obtain 11 with certainty. We can write all the possible


outcomes in a table as follows.

N	Syndrome
$I \otimes I \otimes I$	00
$I \otimes I \otimes X$	01
$I \otimes X \otimes I$	10
$I \otimes X \otimes X$	11
$X \otimes I \otimes I$	11
$X \otimes I \otimes X$	10
$X \otimes X \otimes I$	01
$X \otimes X \otimes X$	00


Observe that the syndromes corresponding to no error, and to bit flips on single qubits (i.e. $I \otimes I \otimes I$, $I \otimes I \otimes X$, $I \otimes X \otimes I$ and $X \otimes I \otimes I$) are all distinct. This means that, if one of these four errors occurs, we can detect it. After we detect a bit-flip error on a given qubit, we can simply apply the same bit-flip operation to that qubit to restore the original encoded state $\alpha|000\rangle + \beta|111\rangle$, which can easily then be mapped to $\alpha|0\rangle + \beta|1\rangle$ by reversing the original encoding circuit. On the other hand, if bit-flip errors occur on more than one qubit, we do not detect them (and indeed this “error correction” process can make matters worse!).

While this code is sufficient to protect against single bit-flip errors, there are other, less classical, errors acting on single qubits which it does not protect against. For example, consider the effect of a Z (“phase”) error acting on the first qubit of the encoded state $|E(\psi)\rangle$, which maps it to $\alpha|000\rangle - \beta|111\rangle$ (recall that $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$). It is easy to see that the syndrome measurement still returns 00, so the error correction operation does nothing and the Z error is not corrected.

However, these Z errors can be detected using a different code. Observe that $Z = HXH$, where H is the Hadamard gate. Thus Z acts in the same way as X , up to a change of basis. If we use the same code as before, but perform this change of basis for each qubit, we obtain a code which corrects against Z errors. In other words, we now encode $|\psi\rangle$ as $\alpha|+++ \rangle + \beta|--- \rangle$. Our new encoding circuit is simply



and our decoding circuit is



The analysis for the previous code goes through without change to show that this code protects against Z errors on an individual qubit. However, it is easy to see that the new code no longer protects against X errors! Can we protect against both errors simultaneously? The answer is yes, by *concatenating* these two codes. We first encode $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ using the code protecting against phase flips, and then encode each of the resulting qubits using the code that protects against

bit flips. In other words, we perform the map

$$\begin{aligned} \alpha|0\rangle + \beta|1\rangle &\mapsto \frac{1}{2\sqrt{2}}(\alpha(|0\rangle + |1\rangle)(|0\rangle + |1\rangle)(|0\rangle + |1\rangle) + \beta(|0\rangle - |1\rangle)(|0\rangle - |1\rangle)(|0\rangle - |1\rangle)) \\ &\mapsto \frac{1}{2\sqrt{2}}(\alpha(|000\rangle + |111\rangle)(|000\rangle + |111\rangle)(|000\rangle + |111\rangle) \\ &\quad + \beta(|000\rangle - |111\rangle)(|000\rangle - |111\rangle)(|000\rangle - |111\rangle)). \end{aligned}$$

The single qubit $|\psi\rangle$ is now encoded using 9 qubits. These qubits can naturally be split into three blocks, each of which encodes one qubit of the state $\alpha|+++ \rangle + \beta|--- \rangle$. To decode this encoded state, first the decoding circuit for the bit-flip code is applied to each block. Assuming at most one bit-flip error has occurred in each block, the result will be the state $\alpha|+++ \rangle + \beta|--- \rangle$, perhaps with a Z error applied to one of the qubits. This state can then be mapped back to $\alpha|0\rangle + \beta|1\rangle$ using the decoding algorithm for the phase-flip code.

Example. Imagine a ZX error occurs on the fourth qubit of the encoded state. The input to the decoding algorithm is thus the state

$$\frac{1}{2\sqrt{2}}(\alpha(|000\rangle + |111\rangle)(|100\rangle - |011\rangle)(|000\rangle + |111\rangle) + \beta(|000\rangle - |111\rangle)(|100\rangle + |011\rangle)(|000\rangle - |111\rangle)).$$

We apply the bit-flip decoding algorithm to each of the three blocks of three qubits, and get syndromes of 00, 11, 00 (“no error”, “error on first qubit”, “no error”). So we perform an X operation on the fourth qubit to correct this, and then the map $|000\rangle \mapsto |0\rangle$, $|111\rangle \mapsto |1\rangle$ on each block of three qubits. The result is the state

$$\alpha|+-+\rangle + \beta|-+-\rangle.$$

Applying the phase-flip decoding algorithm to this state gives $\alpha|0\rangle + \beta|1\rangle$ as required.

We now have a code that can protect against X or Z errors acting on an arbitrary qubit. It may seem that this is only the beginning of the gargantuan task of protecting against every one of the infinitely many possible errors that can occur. In fact, it turns out that we have already done this! The key observation is that the matrices $\{I, X, Y, Z\}$, where $Y = ZX$, form a basis for the complex vector space of all 2×2 matrices, so an arbitrary error operation acting on a single qubit can be written as a linear combination of these matrices.

To be more precise, if we have a quantum channel \mathcal{N} representing some noise process, defined by

$$\mathcal{N}(\rho) = \sum_k N_k \rho N_k^\dagger,$$

observe that to protect $|\psi\rangle$ against \mathcal{N} it is sufficient to produce an encoded state $|E(\psi)\rangle$ such that, after the decoding operation, $N_k|E(\psi)\rangle$ is mapped to a vector proportional to $|\psi\rangle$, for all k . If \mathcal{N} represents noise acting on at most one qubit, we can expand the nontrivial part of each N_k as $N_k = \alpha_k I + \beta_k X + \gamma_k Y + \delta_k Z$. If our code protects against (say) an X error on the first qubit, we know that $(X \otimes I \otimes \dots \otimes I)|E(\psi)\rangle$ is decoded to $|\psi\rangle$. The same applies to all other single-qubit errors. By linearity, this implies that $\mathcal{N}(|E(\psi)\rangle\langle E(\psi)|)$ is decoded to $|\psi\rangle\langle\psi|$, so our code in fact can correct an arbitrary error on an individual qubit.

The following general statement about when quantum error-correction is possible is known, but we will not prove it here.

Theorem 10.1 (Quantum error correction criterion). *Assume we have a code subspace with basis $\{|\psi_i\rangle\}$. A necessary and sufficient condition for the set of errors $\{E_a\}$ to be correctable is*

$$\langle \psi_i | E_a^\dagger E_b | \psi_j \rangle = C_{ab} \delta_{ij}$$

for all a, b, i and j .

10.2 The stabilizer formalism

We now describe an elegant way of describing quantum states which is particularly useful in the setting of quantum error-correction, but is also an important tool elsewhere in quantum information theory: the stabilizer formalism. This provides a concise way of describing states of n qubits which can be highly entangled.

The starting point is the Pauli matrices on n qubits: the set of n -qubit matrices of the form

$$M = M_1 \otimes M_2 \otimes \cdots \otimes M_n,$$

where for each i , $M_i \in \{I, X, Y, Z\}$. If we multiply by a global phase from the set $\{1, i, -1, -i\}$, this set becomes a group under multiplication, called the Pauli group. Every pair of Pauli matrices either commutes or anticommutes. Consider a set \mathcal{M} of k Pauli matrices on n qubits, such that all pairs of matrices in the set commute. We can write down such a set as an $k \times n$ matrix whose entries are picked from the set $\{I, X, Y, Z\}$, where the i 'th row specifies an n -qubit Pauli matrix $M^{(i)}$. For example:

$$\begin{pmatrix} X & X \\ Z & Z \end{pmatrix}, \begin{pmatrix} I & X & Z \\ Y & Y & X \\ X & Z & X \end{pmatrix}, \begin{pmatrix} I & X & Z & I \\ X & Y & X & Z \end{pmatrix}.$$

We say that $|\psi\rangle$ is *stabilized* by $M^{(i)}$ if

$$M^{(i)}|\psi\rangle = |\psi\rangle.$$

Observe that if $M^{(i)}$ and $M^{(j)}$ stabilize $|\psi\rangle$, so does $M^{(i)}M^{(j)}$. Assume that all the matrices $M^{(i)}$ are independent; that is, none of them can be expressed as the product of any of the others (even up to a phase), and consider the set S of all n -qubit states stabilized by all of these matrices, i.e. the set of states $|\psi\rangle$ such that

$$M^{(i)}|\psi\rangle = |\psi\rangle$$

for all $i = 1, \dots, k$. S is a subspace of \mathbb{C}^{2^n} . As all the matrices commute, the projector onto S , Π_S , is the product of the projectors onto all the $+1$ eigenspaces of the matrices $M^{(i)}$:

$$\Pi_S = \frac{1}{2^k} \prod_{i=1}^k (I + M^{(i)}).$$

We have

$$\text{tr } \Pi_S = \frac{1}{2^k} \text{tr} \prod_{i=1}^k (I + M^{(i)}) = \frac{1}{2^k} \sum_{T \subseteq [k]} \text{tr} \prod_{i \in T} M^{(i)} = 2^{n-k}.$$

The final equality holds because $\prod_{i \in T} M^{(i)} \neq \pm I$ for all $T \neq \emptyset$, by our assumption that the matrices are all independent. If the product is not proportional to the identity, it must be a Pauli matrix with trace 0.

Therefore, the subspace S has dimension 2^{n-k} . If $k = n$, there is a unique state stabilized by all the matrices; such a state is known as a *stabilizer state*.

Example 10.2. $|1\rangle$ and $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ are both stabilizer states. $|1\rangle$ is stabilized by $\{-Z\}$, and $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ is stabilized by $\{XX, ZZ\}$.

A stabilizer state can thus be described by the pair of an $n \times n$ matrix whose entries are picked from $\{I, X, Y, Z\}$, and n signs ± 1 ; this is significantly more concise than the 2^n complex numbers required to describe a generic quantum state. As well as this static description, we can efficiently describe some dynamics using this picture. Consider a unitary operation U that preserves the Pauli group under conjugation: $UPU^\dagger = P'$ for some Pauli matrix P' . Then, if P stabilizes $|\psi\rangle$, we have

$$P'(U|\psi\rangle) = UPU^\dagger U|\psi\rangle = UP|\psi\rangle = U|\psi\rangle,$$

so P' stabilizes $U|\psi\rangle$. The set of operations which preserve the Pauli group in this way is a group too, called the Clifford group. All of the following gates are in the Clifford group:

$$\mathcal{C} = \{X, \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, H, \text{CNOT}\},$$

and in fact it turns out that the Clifford group is generated by the above set. This is the basis of the following result:

Theorem 10.3 (Gottesman-Knill Theorem). *Any quantum circuit consisting of gates picked from \mathcal{C} and single-qubit measurements can be simulated efficiently classically.*

On the other hand, if we have a set \mathcal{M} of $k < n$ commuting Pauli matrices, we can view the subspace S stabilized by all of these matrices as encoding $n - k$ qubits. One reason for doing this is that such a subspace may have good error-correction properties. Imagine the state $|\psi\rangle$ which we would like to preserve is contained in S and let E be a Pauli error. Then, if $E \in \mathcal{M}$, $E|\psi\rangle = |\psi\rangle$, so $|\psi\rangle$ is unaffected by the error E . On the other hand, if E anticommutes with some element $M \in \mathcal{M}$, we have

$$ME|\psi\rangle = -EM|\psi\rangle = -E|\psi\rangle,$$

so if we measure M and get the answer -1 , we will detect that an error has occurred.

We claim that the quantum error-correction condition (Theorem 10.1) will be satisfied if for each pair of error operations E_a, E_b , either:

1. $E_a^\dagger E_b \in \mathcal{M}$, or
2. There exists $M \in \mathcal{M}$ that anticommutes with $E_a^\dagger E_b$.

In case (1),

$$\langle \psi_i | E_a^\dagger E_b | \psi_j \rangle = \langle \psi_i | \psi_j \rangle = \delta_{ij};$$

in case (2),

$$\langle \psi_i | E_a^\dagger E_b | \psi_j \rangle = \langle \psi_i | E_a^\dagger E_b M | \psi_j \rangle = -\langle \psi_i | M E_a^\dagger E_b | \psi_j \rangle,$$

so we must have $\langle \psi_i | E_a^\dagger E_b | \psi_j \rangle = 0$. In either case, the condition of Theorem 10.1 is satisfied.

A stabilizer code has distance d if each Pauli error matrix $E = E_a^\dagger E_b$ of weight strictly less than d satisfies conditions (1) and (2), but there exists such a matrix E of weight d that fails to satisfy these conditions. A code correcting errors on up to t qubits must have $d \geq 2t + 1$. That is, for all E of weight $< d$,

$$E \in S, \text{ or there exists } M \in \mathcal{M} \text{ that anticommutes with } E.$$

Imagine we have a unitary matrix N such that N commutes with everything in the stabilizer, but is not contained within it. These are precisely the errors which we cannot correct. Looked at another way, such matrices allow us to perform logical operations on data encoded in a stabilizer code, by moving around the code space. In particular, if we can find two such matrices N_1, N_2 such that N_1 and N_2 anticommute, this lets us make logical X and Z operations.

Shor's 9 qubit code is a stabilizer code, with the following stabilizer:

$$\begin{pmatrix} Z & Z & I & I & I & I & I & I & I \\ Z & I & Z & I & I & I & I & I & I \\ I & I & I & Z & Z & I & I & I & I \\ I & I & I & Z & I & Z & I & I & I \\ I & I & I & I & I & I & Z & Z & I \\ I & I & I & I & I & I & Z & I & Z \\ X & X & X & X & X & X & I & I & I \\ X & X & X & I & I & I & X & X & X \end{pmatrix}$$

It can be verified from this representation that the code has distance 3 and encodes one qubit. In addition, the operators $X^{\otimes 9}$ and $Z^{\otimes 9}$ commute with everything in the stabilizer but are not contained within it. These therefore function as our logical X and Z operators. A smaller code (in fact, the smallest that can encode a qubit) is the five qubit code

$$\begin{pmatrix} X & Z & Z & X & I \\ I & X & Z & Z & X \\ X & I & X & Z & Z \\ Z & X & I & X & Z \end{pmatrix}.$$

This code also has distance 3 and encodes one qubit. Another beautiful example of a stabilizer code, but one which is just beyond the scope of this course, is the *toric code* due to Kitaev.

11 Quantum state discrimination and tomography

It is a fundamental feature of quantum mechanics that nonorthogonal states cannot be distinguished with certainty. But just how well can they be distinguished? This question is addressed by the field of quantum state discrimination. A prototypical problem studied here is: given a quantum state $|\psi\rangle$, promised to be picked from a known set of states $\{|\psi_1\rangle, \dots, |\psi_n\rangle\}$ uniformly at random, determine which of the states $|\psi\rangle$ is, with optimal success probability. As well as being a useful subroutine in the design of quantum algorithms, this problem occurs in many other quantum information processing tasks.

To study this problem, we first need to define what a quantum measurement is in full generality. There are measurements that occur in quantum mechanics which go beyond the simple computational basis measurements we have seen so far. The most general measurements we can perform are known as positive operator-valued measures (POVMs), sometimes also referred to as probability operator measures (POMs). A POVM is described by a sequence of positive semidefinite operators μ_k such that $\sum_k \mu_k = I$. The probability of getting outcome k when the measurement is applied to state ρ is $\text{tr}(\mu_k \rho)$. This makes sense because

$$\sum_k \text{tr}(\mu_k \rho) = \text{tr} \left(\left(\sum_k \mu_k \right) \rho \right) = \text{tr} \rho = 1,$$

so the total probability of getting one or other of the measurement outcomes is equal to 1. A computational basis measurement of n qubits, for example, fits into this picture: here the operators are the orthogonal rank 1 projectors $|x\rangle\langle x|$, $x \in \{0, 1\}^n$.

An arbitrary POVM can in fact be expressed as a combination of a unitary operation on a bigger space, followed by a computational basis measurement. This result is known as (a special case of) Neumark's Theorem, and can ultimately be seen as a consequence of the Stinespring dilation discussed previously. For a POVM with measurement operators μ_k , define corresponding operators M_k such that $M_k^\dagger M_k = \mu_k$. Such operators can always be found by non-negativity of the operators μ_k . Then define an isometry

$$V = \sum_k M_k \otimes |k\rangle$$

from the original system into a system-plus-ancilla space. V is indeed an isometry because

$$V^\dagger V = \sum_{k,\ell} M_k^\dagger M_\ell \langle k|\ell\rangle = \sum_k M_k^\dagger M_k = \sum_k \mu_k = I.$$

The desired POVM can be implemented by applying V to the input state ρ , then measuring the ancilla system in the computational basis. The probability of obtaining outcome i is

$$\text{tr}((I \otimes |i\rangle\langle i|)(V\rho V^\dagger)) = \text{tr} \left((I \otimes |i\rangle\langle i|) \sum_{j,k} (M_j \rho M_k^\dagger) \otimes |j\rangle\langle k| \right) = \text{tr} M_i \rho M_i^\dagger = \text{tr} \mu_i \rho$$

as desired, using cyclicity of the trace. Unlike the projective measurements you have seen so far, note that the post-measurement state depends on how the POVM is implemented in terms of operators M_k . Assuming we receive measurement outcome i , the resulting state is

$$\frac{M_i \rho M_i^\dagger}{\text{tr} \mu_i \rho}.$$

11.1 Optimal discrimination of two quantum states

In the case where we have only two states to discriminate, the optimal success probability can be characterised exactly, via a result known as the Holevo-Helstrom theorem. Assume we are given one copy of a state promised to be either ρ_0 or ρ_1 . The probability the state is ρ_0 is p , and the probability the state is ρ_1 is $1 - p$. Then we have the following theorem.

Theorem 11.1. *The optimal probability of success of deciding whether we are given ρ_0 or ρ_1 is precisely*

$$\frac{1}{2} + \frac{1}{2} \|p\rho_0 - (1-p)\rho_1\|_1, \quad (8)$$

where $\|\cdot\|_1$ is the trace norm $\|M\|_1 = \sum_i |\lambda_i(M)|$, with λ_i being the i 'th eigenvalue of M .

Proof sketch. We first show that (8) is an upper bound on the success probability, and then that it can be achieved. Let $\{\mu_0, \mu_1\}$ be an arbitrary POVM with two measurement outcomes. Then the average success probability achieved by this POVM is

$$p \operatorname{tr}(\mu_0 \rho_0) + (1-p) \operatorname{tr}(\mu_1 \rho_1).$$

Using some algebraic manipulations and $\mu_0 + \mu_1 = I$, the following equality holds:

$$p \operatorname{tr}(\mu_0 \rho_0) + (1-p) \operatorname{tr}(\mu_1 \rho_1) = \frac{1}{2} + \frac{1}{2} \operatorname{tr}((\mu_0 - \mu_1)(p\rho_0 - (1-p)\rho_1)).$$

As $\mu_0 - \mu_1 = 2\mu_0 - I$, and the eigenvalues of μ_0 are bounded between 0 and 1, the eigenvalues of $\mu_0 - \mu_1$ are bounded between ± 1 , i.e. $\|\mu_0 - \mu_1\| \leq 1$. By Hölder's inequality for Schatten p -norms (which we will not prove here, but should be plausible...),

$$\operatorname{tr}((\mu_0 - \mu_1)(p\rho_0 - (1-p)\rho_1)) \leq \|\mu_0 - \mu_1\| \|p\rho_0 - (1-p)\rho_1\|_1 \leq \|p\rho_0 - (1-p)\rho_1\|_1.$$

This completes the proof of the upper bound. To see that it can be achieved, consider the POVM with two elements $\{\mu_0, \mu_1\}$, where μ_0 is the projector onto the non-negative part of $p\rho_0 - (1-p)\rho_1$ (the subspace spanned by eigenvectors with non-negative eigenvalues), and μ_1 is the projector onto the negative part of this operator. Then

$$\operatorname{tr}((\mu_0 - \mu_1)(p\rho_0 - (1-p)\rho_1)) = \|p\rho_0 - (1-p)\rho_1\|_1$$

as desired. \square

In the case where ρ_0 and ρ_1 are pure states, the expression (8) can be simplified using the following lemma.

Lemma 11.2. *For states $|\psi\rangle, |\phi\rangle$,*

$$\|p|\psi\rangle\langle\psi| - (1-p)|\phi\rangle\langle\phi|\|_1 = \sqrt{1 - 4p(1-p)|\langle\psi|\phi\rangle|^2}.$$

Proof. We can apply an arbitrary unitary operation to both of $|\psi\rangle, |\phi\rangle$ without changing the left-hand side. So map $\sqrt{p}|\psi\rangle$ to $\sqrt{p}|0\rangle$, and map $\sqrt{1-p}|\phi\rangle$ to $\sqrt{1-p}\alpha|0\rangle + \sqrt{1-p}\sqrt{1-\alpha^2}|1\rangle$ for $\alpha, \beta \in \mathbb{R}$. Then $|\langle\psi|\phi\rangle|^2 = \alpha^2$. We can now explicitly compute

$$\|p|\psi\rangle\langle\psi| - (1-p)|\phi\rangle\langle\phi|\|_1 = \left\| \begin{pmatrix} p - (1-p)\alpha^2 & -(1-p)\alpha\sqrt{1-\alpha^2} \\ -(1-p)\alpha\sqrt{1-\alpha^2} & -(1-p)(1-\alpha^2) \end{pmatrix} \right\|_1 = \sqrt{1 - 4p(1-p)\alpha^2}.$$

\square

In the case where we wish to optimally discriminate between $n > 2$ quantum states, unfortunately no concise expression like (8) is known.

11.2 Quantum state tomography

What if we don't have the prior information that ρ is picked from a known set of states, but instead are asked to determine an input state ρ which is completely unknown? In this setting it is intuitively obvious that we can learn no useful information from only one copy of ρ . However, if we are given n copies of ρ , for some large n , we can attempt to characterise ρ using measurement statistics. This procedure is known as quantum state tomography.

Assume that we are given an n -qubit state ρ that we wish to determine. The simplest quantum state tomography procedure is based around measuring observables corresponding to n -qubit Pauli matrices. This relies on the fact that any n -qubit state can be expressed in terms of Pauli matrices. Indeed, for any state ρ we can write

$$\rho = \sum_{s \in \{I, X, Y, Z\}^n} \hat{\rho}(s) s_1 \otimes s_2 \otimes \cdots \otimes s_n$$

for some real coefficients $\hat{\rho}(s)$, which have the nice characterisation

$$\hat{\rho}(s) = \frac{1}{2^n} \text{tr}(\rho(s_1 \otimes s_2 \otimes \cdots \otimes s_n)).$$

Therefore, if we can learn all the coefficients $\hat{\rho}(s)$, we can learn ρ . To estimate an individual coefficient $\hat{\rho}(s)$, we measure a number of copies of ρ in the eigenbasis of the n -qubit Pauli matrix corresponding to s . Each such matrix has eigenvalues ± 1 corresponding to eigenspaces projected onto by projectors P_+ , P_- , so

$$\text{tr}(\rho(s_1 \otimes s_2 \otimes \cdots \otimes s_n)) = \text{tr}(\rho(P_+ - P_-)) = \Pr[\text{get outcome } +1] - \Pr[\text{get outcome } -1].$$

So, if we can estimate these probabilities, we can estimate ρ . Measuring in the eigenbasis of an n -qubit Pauli matrix $M = M_1 \otimes M_2 \otimes \cdots \otimes M_n$ is easy. For each qubit i such that $M_i = X$ we apply a Hadamard gate; for each qubit i such that $M_i = Y$ we apply $\frac{1}{\sqrt{2}} \begin{pmatrix} i & 1 \\ -i & 1 \end{pmatrix}$; and for each qubit i such that $M_i = Z$, we do nothing. We then measure each qubit in the computational basis, associating 0 with 1, and 1 with -1 ; and finally take the product of the results. If we take the average of many samples, we will get something approximating $\hat{\rho}(s)$.

Given a list of estimates $\widetilde{\hat{\rho}(s)}$, we need to combine these to produce an overall estimate of ρ . This is not a trivial task; for example, given our noisy measurement results we should attempt to output a guess $\tilde{\rho}$ which is actually a quantum state, i.e. positive semidefinite with trace 1. Good strategies for achieving this are currently an active topic of research.

Note that this process is fundamentally inefficient: to characterise an n -qubit state we needed to estimate the probabilities of 4^n different measurement outcomes. This was inevitable as a mixed quantum state of n qubits depends on roughly 4^n parameters. Surprisingly, tight bounds on the complexity of quantum state tomography were only obtained very recently. It was shown by O'Donnell and Wright that an arbitrary d -dimensional state ρ can be learnt up to accuracy ϵ in the trace norm with $O(d^2/\epsilon^2)$ copies of ρ (an independent result of Haah et al. achieves a similar bound, with slightly worse parameters but a shorter proof). These results use a different protocol to that described here.

11.3 Limits on quantum data compression

As we have just discussed, a quantum state of n qubits can be seen as having exponential complexity in terms of n . However, here we finish by describing a way in which an n -qubit state can be seen as

only “really” encoding n bits of information. This is surprising given the many ways, some of which have been discussed in this course, that quantum computation seems to exponentially outperform classical computation.

Imagine we want to encode n bits of information in an m -qubit quantum state, for some $m \leq n$. We do this by associating a state ρ_x with each input bit-string $x \in \{0,1\}^n$. To extract the information, we apply a POVM described by measurement operators $\{\mu_y\}$, and output the bit-string corresponding to the measurement outcome. Assuming that each bit-string x is equally likely, the probability that we output the correct answer is

$$\frac{1}{2^n} \sum_{x \in \{0,1\}^n} \text{tr}(\mu_x \rho_x) \leq \frac{1}{2^n} \sum_{x \in \{0,1\}^n} \|\mu_x\|_1 \|\rho_x\| \leq \frac{1}{2^n} \sum_{x \in \{0,1\}^n} \|\mu_x\|_1 = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} \text{tr} \mu_x = 2^{m-n}.$$

Here the first inequality is Hölder’s inequality for Schatten p -norms; the second inequality is the fact that a density matrix’s eigenvalues are bounded between 0 and 1; the first equality holds because μ_x is positive semidefinite; and the second equality is $\sum_{x \in \{0,1\}^n} \mu_x = I$. Therefore, even for $m = n/2$ (for example), the probability of successfully decoding x goes down exponentially with n . In particular, we cannot learn x with certainty unless $m \geq n$.

This bound, a variant of which is originally due to Nayak but has since been rediscovered multiple times, can be seen as a “toy version” of a famous result called Holevo’s Theorem. Roughly speaking, this result states that the amount of accessible information in an n -qubit quantum state is upper bounded by n bits.