

QUANTUM INFORMATION

This is the preliminary version from 2019.

In the 1920ies quantum mechanics was discovered, by *Heisenberg*, *Born*, *Schrödinger*, and others. It was found to involve probability theory in an intimate way. But classical probability theory was still in a primitive state. *Kolmogorov* gave a proper axiomatization of probability theory in the 1930ies. *Shannon* introduced information theory in the 1940ies. In the following decade the theory of stochastic processes saw great developments. So did operator theory, the work of *Stinespring* being especially important. In the 1970ies the time had come to update the probabilistic structure of quantum mechanics. Thus the subject of quantum information was born, in the hands of *Holevo*, *Kraus*, *Lindblad*, and a few others. A new ingredient was introduced when the idea of quantum computers was put on the table, by *Benioff*, *Deutsch*, and others. The subject grew dramatically in the 1990ies largely because experimentalists caught up with it, enabling theorists like *Bennett* and *Shor* to assume that the theory deals with things you can actually do. At the moment it is one of the main growth points of physics. It will take some time before we know where it leads.

Quantum mechanics has many successes to its credit, such as enabling us to understand the wave-particle nature of light, the structure of atoms, and the stability of matter. Here we just take a lightning tour through quantum information theory, assuming as little as possible in the way of prerequisites. The course falls into four separate parts:

- Lengthy introduction
- Open systems
- Information theory
- Quantum computation

All of this to be crammed into 2×16 hours of teaching. To support you, and keep you company in the evenings, you have the book by *Stenholm* and *Suominen*, and these lecture notes. Since the book starts with a quote, I do so too: “*it would be contrary to ... the perfection of things, if there were no chance events*”. The source is the same as *Stig’s* and *Kalle-Antti’s*.

LENGTHY INTRODUCTION

Throughout the lectures I will stress the connection to classical probability theory rather more than the book does. The book on the other hand stresses the connection to experiments in quantum optics much more than I will do.

The set of classical probabilities

Probability theory is simple to define mathematically if we stick to probability distributions over a finite number n of mutually exclusive events. Then we need n non-negative numbers summing to 1,

$$p_i \geq 0, \quad \sum_{i=1}^n p_i = 1. \quad (1)$$

We can collect these numbers into a vector \vec{p} and refer to any such vector as a probability vector, as a *probability distribution* P , or—in analogy to the quantum states that we will introduce later—as a *classical state*.

It is an embarrassment that the question what it *means* to assign the number $p_1 = 0.561$ (say) is controversial. Is it an objective statement about frequencies? Or a measure of rational belief? Or what? Things are simple only as long as we don't ask.

A bit of extra notation is helpful. A *random variable* is defined as a function that assigns real numbers to events. The random variable A takes values a_i with probabilities

$$p_i = P(A = a_i). \quad (2)$$

If we have several random variables we can define the *joint probability* $p_{i,j} = P(A = a_i, B = b_j)$, which is the probability that A takes the value a_i and B takes the value b_j , and the *conditional probability* $P(A = a_i | B = b_j)$, which is the probability that A takes the value a_i given that B is known to take the value b_j . Joint and conditional probabilities are connected by *Bayes's* formula. In short hand notation

$$P(A, B) = P(A|B)P(B) = P(B|A)P(A). \quad (3)$$

If there are n outcomes for the first and m for the second, the joint probability vector carries a collective index and has nm components. We require

$$p_i = P(A = a_i) = \sum_{j=0}^m p_{i,j} , \quad q_j = P(B = b_j) = \sum_{i=0}^n p_{i,j} . \quad (4)$$

Probability vectors that arise in this way from a joint probability distribution are known as *marginal* distributions.

The random variables may or may not be correlated. If not, they are *independent*, and the probability vector for the joint event can be written as

$$p_{i,j} = P(A = a_i, B = b_j) = P(A = a_i)P(B = b_j) = p_i q_j . \quad (5)$$

I stress that this equation is *not* true in general. Indeed, suppose $n = m = 2$, and order the components of the joint probability vector lexicographically:

$$P_0 = p_{0,0} , \quad P_1 = p_{0,1} , \quad P_2 = p_{1,0} , \quad P_3 = p_{1,1} . \quad (6)$$

Now suppose that the random variables are independent, as in eq. (5). Then the probability vectors that actually occur (the frequencies that will actually be observed, if you look at it that way) are constrained by

$$P_0 P_3 = P_1 P_2 . \quad (7)$$

If the observed frequencies do not obey this condition, you conclude that the two random variables must be correlated.¹

The set of probability vectors form a *convex set*. What this means is that given two probability vectors $\vec{p}_{(1)}$ and $\vec{p}_{(2)}$ and a number $x \in (0, 1)$, the combination

$$\vec{p} = x\vec{p}_{(1)} + (1 - x)\vec{p}_{(2)} \quad (8)$$

is a probability vector too. Conversely, given a probability vector \vec{p} we can ask if it can be formed as a mixture of two distinct probability vectors in this way. If no such decomposition is possible the vector \vec{p} is said to be *pure*. Geometrically, you can think of a convex set as a set of points such that the straight line between any pair of points in the set also belongs to the set. Of course this assumes that we have a rule for adding points, and for drawing

¹Exercise: Prove this. Note that it is an ‘if and only if’ statement.

straight lines between them. But this we have, because we regard our points as vectors in a vector space of a fixed dimension n . Given a set of points you can form the minimal convex set that contains them. This is known as their *convex hull*.

The set of probability vectors is a convex set of a special kind. When $n = 3$ there are three pure vectors, and every probability vector \vec{p} can be written as a convex combination of these pure vectors,

$$\vec{p} = p_1 \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + p_2 \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} + p_3 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad p_i \geq 0, \quad p_1 + p_2 + p_3 = 1. \quad (9)$$

Geometrically this is a triangle, or a *simplex* if we keep the number of outcomes n arbitrary.² A convex combination, also known as a *mixture*, is a quite special case of a linear combination, because the coefficients are required to be non-negative and sum to one. For simplices the decomposition of a point into a convex combination of pure points is unique, but this is a special feature of simplices. A square is also a convex set, but there is no unique way of decomposing an interior point as a mixture of the four pure points at the corners. The same is true for a circular disk, which has infinitely many pure points. This last observation will become relevant when we come to the quantum case.

Dynamics and distinguishability

If you want to introduce dynamics on a probability simplex, the options are rather limited. You can use *stochastic maps*, given by matrices S such that the vector \vec{q} defined by

$$\vec{q} = S\vec{p} \quad (10)$$

is a probability vector for every probability vector \vec{p} .³ An example of a stochastic map is the ‘bit flip’ matrix

²Exercise: When $n = 4$ the simplex is a tetrahedron, and eq. (7) describes a surface inside a tetrahedron. Draw it! You should be able to see that it is formed from straight lines obtained by keeping either \vec{p} or \vec{q} constant.

³Exercise: What conditions on the matrix elements of S ensure that this holds?




Figure 1: We illustrate the stochastic map (11). In information theory stochastic maps are called ‘channels’, and this is the ‘binary symmetric channel’.

$$S = \begin{pmatrix} 1-p & p \\ p & 1-p \end{pmatrix}, \quad (11)$$

which may describe the degrading noise acting on a bit in a computer memory. If we also require the stochastic maps to take pure states into pure states, then the matrix is a permutation matrix. General stochastic maps tend to shrink the simplex towards some fixed point, and can be combined into *Markov chains*.⁴ A *Markov process* is a special kind of stochastic time development in which what happens in a given step depends only on the result of the previous step. To see if a real physical phenomenon—like the Brownian motion studied by *Einstein*—is a Markov process, one needs a careful study of the relevant time scales.

This is a natural point to introduce the notion of *distance* between states, in this case between probability vectors. The first attempt is to say that the distance is defined using the recipe of Euclidean geometry, so that

$$D^2(\vec{p}, \vec{q}) = \sum_{i=1}^n (p_i - q_i)^2. \quad (12)$$

The triangle becomes an ordinary flat triangle. But this turns out to be unsatisfactory from two points of view. From the first point of view one asks how easy it would be to *distinguish* two probability assignments from each other using a finite number N of observations or *samplings*. You may have two competing theories predicting the probability distributions \vec{p} and \vec{q} , respectively. You perform N measurements and obtain a frequency vector \vec{v} , which does not agree exactly with either prediction because of statistical

⁴Exercise: Apply the bit flip map (11) N times, and see what happens to the (one dimensional) probability simplex.

fluctuations. How many measurements do you have to perform before you can sit judgment between the two theories? If that number is large, the predictions are ‘close’, and we want a notion of distance that encodes this idea. The second point of view suggests that the notion of distance should have the property that

$$D(S\vec{p}, S\vec{q}) \leq D(\vec{p}, \vec{q}) . \quad (13)$$

This means that one requires *monotonicity* under arbitrary stochastic maps S . Either way, we insist that a notion of distance must mean something. This is the *operational* point of view, which pervades the whole subject.

Let us look at a special example of a stochastic matrix:

$$S = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{pmatrix} . \quad (14)$$

This is a *coarse graining* of the outcomes. We have decided not to keep track of the distinction between the second and the third outcome, so certainly distinguishability goes down if we perform the map $\vec{p} \rightarrow \vec{q} = S\vec{p}$. But Euclidean distances on a flat simplex can increase under this map, as you can see in Figure 2. One way to handle this is to write

$$x_i = \sqrt{p_i} \geq 0 \quad \Rightarrow \quad \sum_i p_i = \sum_i x_i^2 = 1 . \quad (15)$$

The probability simplex then looks like the positive octant of a sphere. We define the *Fisher–Rao distance* d_{FR} between two probability vectors as the length of the shortest curve connecting them on the sphere that we just defined. It is well known that the shortest curve between two points on the sphere is always an arc of a great circle on the sphere, and you can easily convince yourself that if we connect two points on an octant of the sphere with such an arc, then the entire arc will lie within the octant. The distance d_{FR} between two probability distributions \vec{p} and \vec{q} then becomes simply the angle between the two unit vectors they define, which is given by the formula

$$\cos d_{\text{FR}} = \sum_i \sqrt{p_i} \sqrt{q_i} . \quad (16)$$

This should be a familiar property of the unit sphere.

This notion of distance behaves nicely under stochastic maps. On the round octant of a sphere the transformation (14) is in the direction of an arc of constant latitude, and now distances can only increase. This is true for all stochastic maps. We omit the proof, but Figure 2 shows that it holds for the coarse graining map.




Figure 2: Coarse graining collapses the entire simplex onto one of its edges. Euclidean distances can increase under this operation—two points at distance D from each other will find themselves at a distance $D' > D$. On a round octant this never happens.

The Fisher-Rao metric does have an operational meaning in terms of how well one can distinguish between two probability distributions based on the frequencies observed in a finite but large number of trials. Suppose that there are only two outcomes, and that one of them happens with probability p . The probability to see it happening m times if we do N samplings is

$$P(m) = \binom{N}{m} p^m (1-p)^{N-m} . \quad (17)$$

We get a frequency vector \vec{v} with components $(m/N, (N-m)/N)$. If N is large we expect this to be a fair approximation of the probability vector $(p_1, p_2) = (p, 1-p)$. Indeed, after 20 years of work, *Jakob Bernoulli* proved an important theorem.⁵

Law of Large Numbers. For every $\epsilon > 0$ and $\delta > 0$ the probability to obtain the first outcome m times in N trials obeys

⁵Exercise: Look in a textbook on probability theory for a proof that takes less than 20 years to do.

$$\mathcal{P} \left(\left| \frac{m}{N} - p \right| < \epsilon \right) > 1 - \delta , \quad (18)$$

provided that $N > N_0$ for some sufficiently large N_0 .

Information theory, to which we will arrive later, relies heavily on this.

Here we can make the Law of Large Numbers into a more quantitative statement by bringing in Stirling's formula

$$n! \sim \chi(n) = \sqrt{2\pi n} \left(\frac{n}{e} \right)^n . \quad (19)$$

This formula turns up whenever large numbers are to be counted.⁶ The sign ' \sim ' means that

$$\lim_{n \rightarrow \infty} \frac{n!}{\chi(n)} = 1 , \quad (20)$$

and the approximation is accurate to within a percent already for $n = 10$. Delicate handling of it shows that in a large number N of samplings the probability to see the frequency vector $\vec{\nu}$ is well approximated by

$$\mathcal{P}(\vec{\nu}) = \sqrt{\frac{N}{2\pi p_1 p_2}} e^{-\frac{N}{2} \sum_{i=1}^2 \frac{(\Delta p_i)^2}{p_i}} , \quad \Delta p_i = \nu_i - p_i . \quad (21)$$

The sharpness of the peak varies as you move around the probability simplex. As you can see in Figure 3, two probability assignments close to the edges of the interval will be easy to distinguish with a modest number of samplings, and hence we should regard them as far apart. The opposite holds for two probability assignments close to the centre. Quantifying this leads directly to the Fisher-Rao distance, since

$$\frac{1}{4} \frac{(\Delta p)^2}{p} = \frac{1}{4} \left(\frac{\Delta p}{\sqrt{p}} \right)^2 = (\Delta \sqrt{p})^2 \Rightarrow \frac{1}{4} \sum_{i=1}^2 \frac{(\Delta p_i)^2}{p_i} = D_{\text{FS}}^2(\vec{p}, \vec{p} + \Delta \vec{p}) . \quad (22)$$

A similar statement holds also for more than two outcomes. (Just let the sum go from 1 to n .) When we allow a large number of samplings this provides

⁶Exercise: Derive an approximation for $\ln n!$ by rewriting it as a sum and then approximately as an integral. Compare to Stirling's formula.

a justification of the Fisher–Rao distance, at least if the distances are small. Since we do not expect stochastic maps to increase distinguishability, we do indeed expect that metrics quantifying the latter notion must be monotone.




Figure 3: The broader the Gaussian, the harder it is to decide if the observed frequency matches the probability assignment we are assuming. Here $N = 200$, and we have one component of the probability vector on the horizontal axis.

If we do not know what probability vector to assign to a choice between N outcomes, we expect that the likelihood to find the probability vector in a given region on the probability simplex is proportional to the volume of that region. Here it clearly matters whether the simplex is flat or round. If we use the volume element provided by the Fisher-Rao metric for the purpose, we find that most of the volume is concentrated near the edges and vertices of the simplex. We can take this as a reason why, in real life, most things happen in a fairly predictable way. Random events that are evenly poised are quite rare.⁷




Figure 4: It might seem as if 25 percent of all probability distributions lie in the triangle in the middle. However, using the round metric on the simplex only 12 percent of them do.

⁷Exercise: Verify the calculation that goes into Figure 4.

There are other notions of distance that are operationally meaningful, such as the ‘taxi cab’ l_1 -distance

$$\|\vec{p} - \vec{q}\|_1 = \sum_{i=0}^{n-1} |p_i - q_i| . \quad (23)$$

On a two dimensional probability simplex ‘circles’ at constant distance from a given point will then appear as suitably oriented hexagons.⁸ The taxi cab distance is monotone under stochastic maps, and its operational meaning has to do with how reliably two probability distributions can be distinguished by means of a single sampling.

The quantum generalization

To generalize classical probability theory we first rewrite the definitions in terms of diagonal matrices. We replace all probability vectors with matrices

$$P = \begin{pmatrix} p_1 & 0 & 0 \\ 0 & p_2 & 0 \\ 0 & 0 & p_3 \end{pmatrix} , \quad P \geq 0 , \quad \text{Tr}P = 1 . \quad (24)$$

A random variable A is an otherwise unrestricted diagonal matrix,

$$A = \begin{pmatrix} a_1 & 0 & 0 \\ 0 & a_2 & 0 \\ 0 & 0 & a_3 \end{pmatrix} . \quad (25)$$

The expectation value of a random variable, given a state, is then

$$\langle A \rangle = \text{Tr}PA . \quad (26)$$

There is a fairly obvious generalization of all this. We replace the diagonal matrices with ‘diagonalizable’, and more particularly Hermitian, matrices. This still leaves it open whether they should be real or complex matrices, but it will turn out that complex numbers are the preferred choice, so let us make it right away. A state is now a positive (and therefore Hermitian) matrix,

⁸Exercise: Draw such ‘circles’ around a few selected points in the simplex. Argue that this distance is monotone under the map (14).

$$\rho = \rho^\dagger, \quad \rho \geq 0, \quad \text{Tr}\rho = 1. \quad (27)$$

The notation $\rho \geq 0$ means that all the eigenvalues are greater than or equal to zero. A matrix obeying all these conditions is known as a *density matrix*, or as a *quantum state*. A random variable is an otherwise unrestricted Hermitian matrix,

$$A = A^\dagger. \quad (28)$$

For expectation values, we keep the classical formula (26).

Since positive operators are so important, let us state three equivalent definitions:

$$\begin{aligned} \rho \text{ is Hermitian with non-negative eigenvalues} &\Leftrightarrow \\ \Leftrightarrow \langle v|\rho|v\rangle \geq 0 \text{ for every vector } |v\rangle &\Leftrightarrow \\ \Leftrightarrow \rho = X^\dagger X \text{ for some bounded operator } X. & \end{aligned} \quad (29)$$

It is also good to know that if A and B are positive operators then $\text{Tr}AB \geq 0$, even though AB is not Hermitian unless A and B commute.

The quantum generalization is a significant one. The random variables, and the states, now belong to a non-commutative algebra. If the matrices act on a d dimensional Hilbert space the $d - 1$ dimensional classical state space is turned into a $d^2 - 1 = (d - 1)(d + 1)$ dimensional one.⁹ The set of density matrices is again a convex set. To see this we must prove that if ρ_1 and ρ_2 are density matrices then so is ρ , where ρ is the convex combination

$$\rho = x\rho_1 + (1 - x)\rho_2, \quad 0 \leq x \leq 1. \quad (30)$$

The proof is easy given one of the three equivalent definitions of a positive operators. So the density matrices form a convex set, just as the probability vectors do.

⁹Exercise: Verify that a general matrix subject to (27) depends on $d^2 - 1$ real parameters, where d is the dimension of the complex Hilbert space.

A remark on the Dirac notation

We will be dealing with vectors in Hilbert space, and operators acting on them. The Hilbert space will have a finite dimension d , and an orthonormal basis consisting of d unit vectors $|e_i\rangle$. In fact it has infinitely many different orthonormal bases. Having chosen one, we can write every vector in the form

$$|\psi\rangle = \sum_{i=0}^{d-1} z^i |e_i\rangle , \quad (31)$$

and every operator in the form

$$A = \sum_{i,j=0}^{d-1} |e_i\rangle A^i_j \langle e^j| . \quad (32)$$

Frequently this is on the pedantic side, and we will often regard vectors and operators as ordered arrays of numbers in the usual way (so that an operator becomes a matrix). Having said this, there are situations where things become conceptually more clear if we remember that this is just a shorthand. And it is important to know what properties depend only on the operator itself, and what properties depend also on the chosen basis. Being positive is an example of the first kind, and being diagonal is an example of the second. A Hermitian operator defines its own preferred eigenbasis, but if you are interested in two Hermitian operators that do not commute you will have to make the choice yourself.

The qubit

To understand what we have, set $d = 2$ so that we deal with two-by-two matrices. The most general Hermitian two-by-two matrix with unit trace can be written as

$$\rho = \frac{1}{2} \begin{pmatrix} 1+z & x-iy \\ x+iy & 1-z \end{pmatrix} = \frac{1}{2}(\mathbf{1} + x\sigma_x + y\sigma_y + z\sigma_z) , \quad (33)$$

where Pauli's sigma-matrices were introduced in the last step. Think of the triple (x, y, z) as a real vector. We call it the *Bloch vector*. When the Bloch

vector is zero we have what we call the *maximally mixed* state. Its eigenvalues are clearly positive. We calculate that

$$\det \rho = \frac{1}{4}(1 - x^2 - y^2 - z^2) . \quad (34)$$

So the determinant vanishes if and only if the length of the Bloch vector equals unity. The set of positive matrices is determined by the condition

$$\det \rho \geq 0 \quad \Leftrightarrow \quad x^2 + y^2 + z^2 \leq 1 . \quad (35)$$

We call it the *Bloch ball*. Its surface is the *Bloch sphere*. The pure states sit at the surface. States in the interior are known as *mixed*, because they can be obtained as mixtures of pure states—in a highly non-unique fashion.

To see what goes on at the surface we note that every Hermitian matrix can be decomposed into its eigenvectors. That is, we can write

$$\rho = \lambda_0 |\psi_0\rangle\langle\psi_0| + \lambda_1 |\psi_1\rangle\langle\psi_1| , \quad \langle\psi_i|\psi_j\rangle = \delta_{ij} , \quad \lambda_0 + \lambda_1 = 1 . \quad (36)$$

The eigenvectors are orthonormal, and the eigenvalues sum to 1 because $\text{Tr}\rho = 1$. At the surface of the Bloch ball $\det\rho = 0$, and one of the eigenvalues vanish. Hence we have

$$\rho = |\psi\rangle\langle\psi| , \quad (37)$$

for some unit vector $|\psi\rangle$. Conversely, for every unit vector $|\psi\rangle$ we will obtain a density matrix lying on the Bloch sphere in this way. Let us consider the most general case. Since an overall phase factor of the vector can be chosen freely we can make things completely definite by insisting that the first non-vanishing component is real and positive, and then we get

$$|\psi\rangle = \cos\frac{\theta}{2}|e_1\rangle + \sin\frac{\theta}{2}e^{i\phi}|e_2\rangle , \quad 0 \leq \theta \leq \pi , \quad 0 \leq \phi < 2\pi . \quad (38)$$

We have chosen to parametrize the state vector with angles θ and ϕ with the idea that these parameters should, in the end, serve as coordinates on a sphere.

Putting eqs. (37) and (38) together we find

$$\rho = \begin{pmatrix} \cos^2 \frac{\theta}{2} & \cos \frac{\theta}{2} \sin \frac{\theta}{2} e^{-i\phi} \\ \sin \frac{\theta}{2} \cos \frac{\theta}{2} e^{i\phi} & \sin^2 \frac{\theta}{2} \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 + \cos \theta & \sin \theta e^{-i\phi} \\ \sin \theta e^{i\phi} & 1 - \cos \theta \end{pmatrix}. \quad (39)$$

We can now read off the Bloch vector that describes an arbitrary pure state $|\psi\rangle$, namely

$$(x, y, z) = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta). \quad (40)$$

And we can draw a picture of the Bloch sphere, with some states of interest placed where they should be.

Let us think about the basis that was implicit in the discussion. By agreement the *computational basis* consists of the vectors

$$|e_0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \leftrightarrow \theta = 0, \quad |e_1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \leftrightarrow \theta = \pi. \quad (41)$$

(No pedantry here.) We place them at the north and south poles of our sphere. Another pair of orthogonal states are

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \leftrightarrow (\theta, \phi) = \left(\frac{\pi}{2}, 0\right), \quad \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \leftrightarrow (\theta, \phi) = \left(\frac{\pi}{2}, \pi\right). \quad (42)$$

The map of the Bloch sphere is ready, and we just have to fill in some physics.

Notice that orthogonal pure states always sit on antipodal points on the sphere, as far from each other as they can be. Also that every state in the interior can be created as a mixture of pure states in infinitely many ways.

Polarization states of a photon

We have just described the *qubit*, the smallest information carrying unit in quantum information theory. Many physical systems are as simple as this. A well known example is given by the polarization states of a photon. The word ‘photon’ is a dangerous one. *Lamb* suggested that it should not be used by anyone who does not have a special license to do so. Some people *define* a photon as ‘a click in a detector’. So let us instead think of a monochromatic electromagnetic plane wave travelling in the z -direction. While this is a

rather different thing, the possible polarization states of the two things agree, and this is what we want to parametrize. A polarization state of a plane wave is an ellipse swept out by an electric field vector orthogonal to that direction. The magnetic field sweeps out a similar ellipse. The ellipse degenerates to a line for a linearly polarized state. Why is the set of such polarization states equal to a sphere? The answer was given by *Stokes*. Going through the details gives useful experience of how to handle qubits.

We begin with a general vector in \mathbf{C}^2 . We are interested in vectors only up to an overall normalization and an arbitrary phase factor. It is quite convenient to remove the ambiguity by insisting that its first component equals 1. This works for all vectors except for those whose first component vanish. In \mathbf{C}^2 we miss only one state in this way, and this can be added in later. Taking the sign ‘ \sim ’ to mean ‘equal up to an overall complex factor’ we write

$$|\psi\rangle \sim \cos \frac{\theta}{2} |e_0\rangle + \sin \frac{\theta}{2} e^{i\phi} |e_1\rangle \sim |e_0\rangle + \tan \frac{\theta}{2} e^{i\phi} |e_1\rangle = |e_0\rangle + z |e_1\rangle . \quad (43)$$

The equality defines an arbitrary complex number z . The one-to-one correspondence between pure qubit states and the surface of a sphere will turn into the well known correspondence between the complex plane and the surface of a sphere. The ‘extra’ number ∞ will have to be added to the complex plane in order to ensure that the correspondence becomes one-to-one.

We write down the corresponding density matrix, normalizing as we go, and read off the Bloch vector:

$$\frac{1}{1+|z|^2} \begin{pmatrix} 1 \\ z \end{pmatrix} \begin{pmatrix} 1 & \bar{z} \end{pmatrix} = \frac{1}{1+|z|^2} \begin{pmatrix} 1 & \bar{z} \\ z & |z|^2 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1+s_3 & s_1-is_2 \\ s_1+is_2 & 1-s_3 \end{pmatrix} . \quad (44)$$

The components of the Bloch vector are denoted s_i , as is conventional for the Stokes parameters we are aiming at. Clearly

$$s_1 + is_2 = \frac{2z}{1+|z|^2} \quad s_3 = \frac{1-|z|^2}{1+|z|^2} . \quad (45)$$

The condition $s_1^2 + s_2^2 + s_3^2 = 1$ is built in, because we started from a pure state. This is the famous *stereographic projection* from the sphere to the complex plane.

So far there are no polarization ellipses in evidence. To get them in, we perform an operation that may seem strange at first. We take the square root of z . With the exception of the north and south poles every point on the sphere splits in two:

$$z = re^{i\alpha}, \quad w = \sqrt{z} = \pm\sqrt{r}e^{i\frac{\alpha}{2}}. \quad (46)$$

We read off new Bloch vectors from

$$\frac{1}{1+|w|^2} \begin{pmatrix} 1 & \bar{w} \\ w & |w|^2 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1+x_3 & x_1-ix_2 \\ x_1+ix_2 & 1-x_3 \end{pmatrix}. \quad (47)$$

We find

$$\mathbf{x}_{\pm} = \frac{1}{1+r} \begin{pmatrix} \pm 2\sqrt{r} \cos \frac{\alpha}{2} \\ \pm 2\sqrt{r} \sin \frac{\alpha}{2} \\ 1-r \end{pmatrix}. \quad (48)$$

The idea now is to regard the space in which these Bloch vectors sit as a copy of ordinary three dimensional space, think of a monochromatic plane wave (or a photon) propagating in the z -direction, consider the intersection between the Bloch sphere and the plane orthogonal to any of the pair of Bloch vectors we obtained, and finally to project that circle onto the $x-y$ -plane.¹⁰ Actually we consider a sphere of arbitrary radius R , so the equations are

$$\pm 2\sqrt{r} \cos \frac{\alpha}{2} x \pm 2\sqrt{r} \sin \frac{\alpha}{2} y + (1-r)z = 0, \quad x^2 + y^2 + z^2 = R^2. \quad (49)$$

Solve the first equation for z and insert the result into the second. Then the sign ambiguity goes away. After some manipulations, and reintroducing the Stokes parameters, we find that the projection onto the $x-y$ -plane is given by

$$\frac{x^2}{1-s_1} + \frac{y^2}{1+s_1} + \frac{4s_2}{1-s_1^2}xy = \frac{R^2(1-r)^2(1+s_3)}{2(1-s_1^2)}. \quad (50)$$

This is indeed an ellipse.

Let us now write (with a slightly unusual sign convention)

¹⁰Exercise: Draw the picture.

$$(s_0, s_1, s_2, s_3) = (a_1^2 + a_2^2, a_2^2 - a_1^2, -2a_1a_2 \cos \delta, -2a_1a_2 \sin \delta) . \quad (51)$$

Assume $a_1^2 + a_2^2 = 1$ and choose R suitably. Then the equation for our ellipse becomes

$$\left(\frac{x}{a_1}\right)^2 + \left(\frac{y}{a_2}\right)^2 - \frac{2 \cos \delta}{a_1 a_2} xy = \sin^2 \delta . \quad (52)$$

Now consider a general monochromatic plane wave of fixed amplitude propagating in the z -direction,

$$E_x = a_1 \cos(\tau + \delta_1), \quad E_y = a_2 \cos(\tau + \delta_2), \quad E_z = 0, \quad a_1^2 + a_2^2 = 1. \quad (53)$$

Set $\delta = \delta_2 - \delta_1$. A trigonometric exercise then verifies that the electric field vector sweeps out precisely the ellipse described by eq. (52).

This exercise may strike you as a trick. Actually there is a deeper reason why it works. What is really happening is that we are embedding the spin 1/2 Hilbert space \mathbf{C}^2 in the spin 1 Hilbert space \mathbf{C}^3 . If the plane wave describes a neutrino no square root would be necessary in eq. (46). If it describes a spin 2 graviton we would need to take a fourth root. But explaining this point would take us too far afield, so we leave neutrinos and gravitons to their fates.

The important point is that the geography of the Bloch sphere is now completely fixed, in any problem concerning polarisation measurements on photons. Points at the poles of the Bloch sphere correspond to circular polarisation, right or left. Linear polarisation occurs on its equator, with an angle shifted 90° for antipodal points. Notice though that this specific one-to-one correspondence was written by Stokes, not by God. We can (and people sometimes do) change coordinates so that linear polarization states lie on the Greenwich meridian, where they are described by real state vectors. What is always true is that the linear polarization states lie on a great circle of the sphere, with pairs of physically orthogonal states (for which the relative angle of linear polarization is 90°) sitting at antipodal points.

[Afterthought: It would have been better to talk about linear polarization states, a source producing 10^6 photons per second, a very expensive detector with 90 % efficiency, and two Glan-Thompson prisms, 1000 Swedish crowns

apiece, enabling us to *prepare* and *measure* a circle's worth of linear polarization states. The Stokes parameters are good for serious optics, but not really needed in this course.]

Dynamics and distances on the qubit

In the classical theory dynamics is given by stochastic maps, degenerating to permutation matrices if pure states always evolve into other pure states. In quantum theory the general case is difficult enough to merit a chapter of its own, 'Open systems'. But we can understand those transformations that take pure states to pure states. They are simply rotations of the Bloch sphere. We obtain them by acting with a *unitary* matrix on the pure state vectors. To see this it is enough to observe that we can obtain scalar products between Bloch vectors by turning the set of operators acting on \mathbf{C}^d into a Hilbert space in its own right. The scalar product between operators is defined by

$$A \cdot B = \text{Tr} A^\dagger B . \quad (54)$$

Notice the dagger, guaranteeing that

$$\|A\|^2 = \text{Tr} A^\dagger A \geq 0 . \quad (55)$$

Clearly this scalar product is left invariant by

$$A \rightarrow UAU^\dagger , \quad B \rightarrow UBU^\dagger . \quad (56)$$

The complex dimension of the Hilbert space of operators acting on a complex Hilbert space of dimension d is d^2 . This is an important idea that we will use again.

So far things do not really depend on the Hilbert space dimension d . For $d = 2$ we check that if

$$A = \frac{1}{2} \begin{pmatrix} x_3 & x_1 - ix_2 \\ x_1 + ix_2 & -x_3 \end{pmatrix} , \quad B = \frac{1}{2} \begin{pmatrix} y_3 & y_1 - iy_2 \\ y_1 + iy_2 & -y_3 \end{pmatrix} , \quad (57)$$

then

$$\text{Tr}A^\dagger B = \frac{1}{2}\vec{x} \cdot \vec{y}. \quad (58)$$

It follows that, in \mathbf{C}^2 , unitary transformations can be thought of as rotations of the Bloch ball. Moreover every rotation of the Bloch ball can be obtained in this way, but this is a special feature of the qubit.

We have defined a notion of distance between quantum states that reduces to the ordinary Euclidean distance on the Bloch ball, and on the probability simplex spanned by a set of diagonal density matrices. There is no disputing the usefulness of this, but we have already argued that if a distance is to capture the notion of statistical distinguishability it must be defined differently. If we restrict ourselves to pure states and qubits, another notion of distance between quantum states suggests itself, namely the length of an arc of a great circle connecting the two points on the Bloch sphere. This is known as the *Fubini–Study distance*, and denoted D_{FS} . It is an exercise to show that it is given by the formula

$$\cos^2 D_{\text{FS}} = \frac{|\langle \psi | \phi \rangle|^2}{\langle \psi | \psi \rangle \langle \phi | \phi \rangle}, \quad (59)$$

where the vectors are not necessarily normalized.¹¹ The formula turns out to be useful, in the same sense as the Fisher–Rao distance is useful, and it can be used regardless of the dimension of Hilbert space.

Higher dimensions

Many of the formulas for the qubit generalize immediately to Hilbert spaces of dimensions higher than two. Nevertheless the qubit is a misleadingly simple example in many ways. The set of quantum states, that is positive matrices with trace one, has $d^2 - 1$ dimensions. A general density matrix can be obtained from a diagonal one by means of a unitary transformation

$$\rho \rightarrow U\rho U^\dagger. \quad (60)$$

¹¹Exercise: Do the exercise referred to in the text. You may have to adjust the size of the sphere.

Because an overall phase does not matter it is the $SU(d)$ subgroup that acts effectively on the states. You can still think of this as a rotation in $d^2 - 1$ dimensions, but now it is a quite special rotation. This is so because the set of density matrices is no longer a sphere. In fact the surface of a sphere in $d^2 - 1$ dimensions has, in itself, $d^2 - 2$ dimensions. But the pure states can be described by vectors in \mathbf{C}^d , with one of the d complex numbers being irrelevant. So the set of pure states has only $2d - 2$ dimensions, far less than the dimension of the surface of a sphere. The set of all states is the convex hull of the pure states, and is therefore only a subset of the ball. It is well-nigh impossible to visualize the set of density matrices for $d > 2$.¹²

When $d = 2$ we expanded an arbitrary state in terms of the Pauli matrices. A Pauli matrix has eigenvalues ± 1 , which means that it is both Hermitian and unitary. There are two ways to generalize the Pauli matrices to arbitrary dimensions, because you can regard them either as a basis for the set of all Hermitian matrices, or as a basis for the set of all unitary matrices in the d^2 dimensional Hilbert space of d by d matrices equipped with the trace inner product (61). Saving the second way for later, we observe that for any d one can find a set of $d^2 - 1$ Hermitian traceless matrices such that

$$\text{Tr} \lambda_i \lambda_j = d \delta_{ij} . \quad (61)$$

Any $d \times d$ density matrix can then be written as¹³

$$\rho = \frac{1}{d} \left(\mathbf{1} + \sum_{i=1}^{d^2-1} x_i \lambda_i \right) . \quad (62)$$

This defines a generalized Bloch vector with components x_i , although the conditions one has to impose on this vector to ensure that ρ is positive are distinctly unpleasant.

A couple of useful facts, that you can prove by first diagonalizing ρ : A density matrix lies at the boundary of the set if and only if it has a vanishing eigenvalue. A density matrix is pure if and only if $\text{Tr} \rho^2 = 1$. This quantity obeys

¹²Exercise: Let ρ be a Hermitian matrix obeying $\text{Tr} \rho^3 = \text{Tr} \rho^2 = 1$. Prove that ρ is a pure state.

¹³Exercise: Find such a set of Hermitian matrices for $d = 3$.

$$1 \geq \text{Tr}\rho^2 \geq \frac{1}{d}. \quad (63)$$

It is sufficiently important to have a name, *purity*.¹⁴

From now on we will be a bit vague about the dimension of the Hilbert space in this ‘Introduction’. The dimension is arbitrary but finite when this causes no extra trouble, but you may think in terms of qubits if you so prefer.

A preliminary theory of measurements

In classical theory it is fairly obvious what a measurement is, and what it measures. Not so in quantum theory, and most of the discussion has to wait until we come to ‘Open systems’. But we can make a start. Choose any quantum random variable, that is to say a Hermitian matrix. It has a complete set of orthonormal eigenvectors $\{|e_i\rangle\}_{i=0}^{d-1}$, and they correspond to the d mutually exclusive outcomes of an experiment. Given that the state of the system is defined by a particular density matrix ρ , the probability that a particular outcome occurs is declared to be

$$p_i = \text{Tr}|e_i\rangle\langle e_i|\rho = \langle e_i|\rho|e_i\rangle. \quad (64)$$

These numbers are non-negative because $\rho \geq 0$ is a positive matrix, and they sum to unity because

$$\sum_i p_i = \sum_i \text{Tr}|e_i\rangle\langle e_i|\rho = \text{Tr} \sum_i |e_i\rangle\langle e_i|\rho = \text{Tr}\rho = 1. \quad (65)$$

Note carefully the steps in this calculation, which made use first of the linearity of the trace, then of the completeness of the basis, and finally of the normalization of the density matrix. The density matrix is not a probability distribution, but it stands prepared to give one for every Hermitian observable that you choose to introduce.

It is important to notice that the result of a quantum measurement cannot tell you what the state of a qubit was before the measurement, not even if you are assured that it was in some pure state. The measurement is associated to

¹⁴Exercise: Prove these inequalities. Also write down a formula connecting $\text{Tr}M$, $\text{Tr}M^2$, and $\det M$ for an arbitrary two-by-two matrix M .

an eigenbasis unrelated to the state. If you observe the outcome associated to $|e_1\rangle$, all you can say with certainty is that the state was *not* one of the those orthogonal to $|e_1\rangle$. Every other pure state would yield that outcome with some probability.

The situation improves if you are given a large number N of qubits, and an assurance that their states are identical. Then you will obtain a probability distribution \vec{p} over the eigenstates. A look at the Bloch sphere tells you that this outcome is consistent with any state lying on a circle at the surface of the Bloch sphere. If you have no assurance that the state ρ you are trying to determine is pure, you can only conclude that ρ lies on a disk in the Bloch ball. In order to pin down the state exactly you will need not only a supply of many identically prepared qubits, you will also need to perform quantum measurements associated to more than one eigenbasis.¹⁵

There remains the question about the state of the system after the measurement. The answer depends. A photon detected as a click in a detector does not exist after the measurement. In a *projective* or *von Neumann measurement* corresponding to some Hermitian matrix, the state after the measurement is one of the eigenstates $|e_i\rangle$ of that matrix. The system ‘collapses’ to one of the eigenstates, with probabilities given by (64). Some more detail will be added when we come to ‘Open systems’, but for now it is enough to know that the notion of von Neumann measurements is a useful one in the lab.

Nice error bases

Let us return to the Pauli matrices, and look at them as unitary operators. We even give them new names:

$$X = \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad Z = \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (66)$$

There is one more, but we can regard it as a derived quantity because

¹⁵Exercise: If your Hilbert space has dimension d , how many different measurements would you need to determine an arbitrary state ρ , given an infinite supply of identically prepared systems so that each measurement may be repeated to give a frequency?

$$Y = \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = iXZ . \quad (67)$$

The action on the computational basis is

$$X|0\rangle = |1\rangle , \quad X|1\rangle = |0\rangle , \quad Z|0\rangle = |0\rangle , \quad Z|1\rangle = -|1\rangle . \quad (68)$$

X is sometimes called the *bit flip*, and Z is called the *phase flip*. Another interesting operator in this connection is

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} . \quad (69)$$

It interchanges X and Z ,

$$H^2 = \mathbf{1} , \quad HXH^\dagger = Z , \quad HZH^\dagger = X . \quad (70)$$

It is called the *Hadamard gate*, where ‘gate’ means ‘unitary operator’ in discussions of quantum computers.¹⁶

An important property of the set $\{\mathbf{1}, X, Z, XZ\}$ is that it forms a *unitary operator basis*. This is also known as a *nice error basis* in quantum computing, because such bases play a role in quantum error correction. What it means is that there exists a *group* with d^2 elements g_i , $g_0 = e$ being the identity element, and a set of unitary operators U_g acting on \mathbf{C}^d , such that

$$U_e = \mathbf{1} , \quad g_i \neq e \quad \Rightarrow \quad \text{Tr}U_{g_i} = 0 , \quad U_{g_i}U_{g_j} \sim U_{g_i g_j} . \quad (71)$$

Here the sign ‘ \sim ’ means ‘equal up to a phase factor’. It follows that $U_{g_i}^\dagger = U_{g_i^{-1}}$. But then it also follows that

$$\text{Tr}U_{g_i}^\dagger U_{g_i} = d , \quad g_i \neq g_j \quad \Rightarrow \quad \text{Tr}U_{g_i}^\dagger U_{g_j} = 0 . \quad (72)$$

But this means that this set of unitary operators forms an orthonormal basis in the d^2 -dimensional Hilbert space of all operators acting on \mathbf{C}^d , equipped with the natural inner product (54). Hence the name. Every operator, whether Hermitian or unitary or none of those, can be expanded as

$$A = \frac{1}{d} \sum_g U_g \text{Tr}U_g^\dagger A , \quad (73)$$

¹⁶Exercise: How do X , Z , and H rotate the Bloch sphere?

where the sum runs over the elements in the operator basis. In the qubit example, there are four terms.

Nice error bases exist in all dimensions. One of many possibilities goes as follows: Introduce a basis $\{|i\rangle\}_{i=0}^{d-1}$ in Hilbert space, and treat the labels as integers modulo d (that is to say that $i + d = i$). Let

$$\omega = e^{\frac{2\pi i}{d}} . \quad (74)$$

Define two unitary operators Z and X by

$$Z|i\rangle = \omega^i|i\rangle , \quad X|i\rangle = |i+1\rangle . \quad (75)$$

One can now prove that

$$X^d = Z^d = \mathbf{1} , \quad ZX = \omega XZ . \quad (76)$$

The resulting group is called the *Weyl–Heisenberg group*. Up to phase factors it has d^2 group elements $X^i Z^j$, and it can be shown that these group elements form a nice error basis in dimension d .¹⁷

Composite systems

If, in classical probability theory, we look for the combined outcomes of n events that can be either true or false, then there are two outcomes per event but 2^n possible outcomes altogether. If n is large this gives a probability simplex of a high dimension. Quantum mechanics shows a similar exponential growth in the dimension of Hilbert space if you put several qubits together. The idea is that if you have a system composed of two parts, so that you can choose to do observations either on only one of the parts or on the whole, then the Hilbert space is the *tensor product* of the Hilbert spaces of the parts. One way of constructing the tensor product is to introduce orthonormal bases in each *factor*, say $\{|e_i\rangle\}_{i=0}^{d_1-1}$ and $\{|f_j\rangle\}_{j=0}^{d_2-1}$, and use the $d_1 d_2$ vectors $|e_i\rangle \otimes |f_j\rangle$ as an orthonormal basis for the combined system. This is a somewhat clumsy way of expressing things, but it will do. It does not imply that every vector

¹⁷Exercise: Write the operators X and Z as matrices for $d = 2, 3, 4$. Write out all the operators $X^i Z^j$ for $d = 3$, and check that they form a nice error basis.

in the larger space is a product vector, because we can have superpositions such as

$$|\psi\rangle = \sum_{i,j} c_{ij} |e_i\rangle \otimes |f_j\rangle . \quad (77)$$

This is not a product vector in general.¹⁸ You can put n qubits together by iterating the idea, and you will end up with a Hilbert space of dimension 2^n .

We can define operators of a special form, called *local* operators and denoted $A \otimes B$. They act on the basis states according to

$$A \otimes B |e_i\rangle \otimes |f_j\rangle = A |e_i\rangle \otimes B |f_j\rangle . \quad (78)$$

They obey obvious rules such as

$$A \otimes (B + C) = A \otimes B + A \otimes C , \quad (A \otimes B)(C \otimes D) = AC \otimes BD . \quad (79)$$

Operators of the form $A \otimes \mathbf{1}$ and $\mathbf{1} \otimes B$ commute.

When working with components in a product basis we adopt a lexicographical ordering of the basis vectors. Thus

$$|\psi_1\rangle = \begin{pmatrix} a_0 \\ a_1 \end{pmatrix} , \quad |\psi_2\rangle = \begin{pmatrix} b_0 \\ b_1 \end{pmatrix} \quad \Rightarrow \quad |\psi_1\rangle \otimes |\psi_2\rangle = \begin{pmatrix} a_0 b_0 \\ a_0 b_1 \\ a_1 b_0 \\ a_1 b_1 \end{pmatrix} , \quad (80)$$

$$A = \begin{pmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{pmatrix} , \quad B = \begin{pmatrix} b_{00} & b_{01} \\ b_{10} & b_{11} \end{pmatrix} \quad \Rightarrow \quad (81)$$

$$A \otimes B = \begin{pmatrix} a_{00}b_{00} & a_{00}b_{01} & a_{01}b_{00} & a_{01}b_{01} \\ a_{00}b_{10} & a_{00}b_{11} & a_{01}b_{10} & a_{01}b_{11} \\ a_{10}b_{00} & a_{10}b_{01} & a_{11}b_{00} & a_{11}b_{01} \\ a_{10}b_{10} & a_{10}b_{11} & a_{11}b_{10} & a_{11}b_{11} \end{pmatrix} .$$

¹⁸Exercise: For two qubits, write down the condition that this be a product vector and compare to (7).

We frequently abbreviate $|\psi_1\rangle \otimes |\psi_2\rangle$ to $|\psi_1\rangle|\psi_2\rangle$. More drastic abbreviations are sometimes used, especially for qubits where one often abbreviates $|0\rangle \otimes |1\rangle$ and $X \otimes Z$ (say) to $|01\rangle$ and XZ .

By inspection we see that it may be useful to index the components of the density matrix, and other matrices, by a pair of indices. Thus we write

$$\rho = \sum_{i,j=0}^{d_1-1} \sum_{\alpha,\beta=0}^{d_2-1} |e_i\rangle|f_\alpha\rangle \rho^{i\alpha}_{j\beta} \langle e^j|\langle f^\beta|, \quad (82)$$

where we used both Latin and Greek indices to emphasize that the factor Hilbert spaces may have different dimensions (although the dimensions would be the same for the two-qubit case). We do the same for every operator. For local operators the expressions take a quite special form. In particular

$$(A \otimes \mathbf{1})^{i\alpha}_{j\beta} = A^i_j \delta_\beta^\alpha, \quad (\mathbf{1} \otimes B)^{i\alpha}_{j\beta} = \delta_j^i B^\alpha_\beta. \quad (83)$$

We have constructed the Hilbert space $\mathbf{C}^{d_1 d_2} = \mathbf{C}^{d_1} \otimes \mathbf{C}^{d_2}$. Sometimes we will write this as $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$. Starting from the other end, a Hilbert space with a dimension that is a composite number can always be decomposed into a tensor product, but there are many different ways of doing so. A particular way may be preferred by the physics. Suppose that Hilbert space is four dimensional, and that we are able to identify two sets of mutually commuting operators, each of them as rich as the set of operators on a qubit. Then we declare them to be local operators of the special form that guarantee that they commute, relative to a tensor product structure defined especially for the purpose. Thus operators of the form $A \otimes \mathbf{1}$ and $\mathbf{1} \otimes B$ may be applied at opposite ends of an optical table, while non-local operators may have been applied to the system at an earlier stage. For instance, the system may be a pair of photons created in a special two-photon state by a non-linear crystal. In some experiments pairs of photons have been distributed by satellite to locations that are very far apart, and it seems obvious that whatever is done at one of the locations ‘commutes’ with what is done at the other location. This explains why we refer to ‘local’ operators, because then we really mean ‘local in space’.

We have a general formula for the expectation value of any Hermitian operator A given that the state is ρ , namely

$$\langle A \rangle = \text{Tr}(A\rho). \quad (84)$$

But there are many situations, many more than those hinted at in the preceding paragraph, where we are only concerned with expectation values of local operators of the form $A \otimes \mathbf{1}$. Inspection of the formulas (83) suggests that there may be many complexities of ρ that do not come into play when we calculate such expectation values. We are therefore looking for a *reduced state* density matrix ρ_1 , acting only on the first factor in the tensor product, such that

$$A = A_1 \otimes \mathbf{1}_2 \quad \Rightarrow \quad \text{Tr}(A\rho) = \text{Tr}(A_1\rho_1) , \quad (85)$$

where the matrices and the trace at the end are those relevant for the factor Hilbert space \mathcal{H}_1 . This is achieved by taking the *partial trace* of ρ , written

$$\rho_1 = \text{Tr}_2\rho , \quad \rho_2 = \text{Tr}_1\rho . \quad (86)$$

The definition of the partial trace is conveniently given in the matrix representation (82), as

$$(\text{Tr}_2\rho)^i_j = \sum_{\alpha=0}^{d_2-1} \rho^{i\alpha}_{j\alpha} , \quad (\text{Tr}_1\rho)^\alpha_\beta = \sum_{i=0}^{d_1-1} \rho^{i\alpha}_{i\beta} . \quad (87)$$

Clearly $\text{Tr}\rho = \text{Tr}_1\text{Tr}_2\rho$. One can show that ρ_1 , so defined, is the unique linear operator that meets the requirement (85). The importance of the reduced states ρ_1 and ρ_2 is that they encode all the information about the state ρ that can be extracted by means of operations acting only on one of the factors of the Hilbert space. Formulated in this way it sounds abstract, but it is in fact what may be going on in the lab, and elsewhere.¹⁹

In calculations it is often preferable to use the Dirac notation for the trace. Using the product basis we write

$$\text{Tr}\rho = \sum_{i,j} \langle i_1 | \langle j_2 | \rho | i_1 \rangle | j_2 \rangle , \quad \text{Tr}_2\rho = \sum_j \langle j_2 | \rho | j_2 \rangle , \quad \text{Tr}_1\rho = \sum_i \langle i_1 | \rho | i_1 \rangle . \quad (88)$$

Note carefully that $\text{Tr}\rho$ is a number, while $\text{Tr}_2\rho$ is an operator acting on \mathcal{H}_1 .

The story does not stop at two. We can construct tensor products of an arbitrary number of Hilbert spaces, such as $\mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \dots \otimes \mathcal{H}_n$, and trace out whatever factors we are not controlling, to obtain density matrices like

¹⁹Exercise: Compute the two partial traces of the matrix $A \otimes B$ given in (81).

$$\rho_{12} = \text{Tr}_{34}\rho , \quad (89)$$

where ρ acts on the four partite Hilbert space $\mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \mathcal{H}_3 \otimes \mathcal{H}_4$ in this particular example.

Entanglement

The term ‘entanglement’ was coined by *Schrödinger*, in a far-seeing series of papers in the 1930ies. We begin by considering a pure state in a bipartite Hilbert space $\mathcal{H}_1 \otimes \mathcal{H}_2$, and ask what one can do with it using local operators only. For simplicity we assume that both the factors have the same dimension d . A general pure state can be written as

$$|\Gamma\rangle = \frac{1}{\sqrt{d}} \sum_{i,j=0}^{d-1} \Gamma_{ij} |e_i\rangle |e_j\rangle . \quad (90)$$

This is a product vector if there exists vectors $\vec{\alpha}$ and $\vec{\beta}$ such that

$$\Gamma_{ij} = \alpha_i \beta_j . \quad (91)$$

The state is then said to be *separable*. If it is not separable it is said to be *entangled*, but entanglement is (as we will see) a question of degree. Now let us act on the state with a local unitary $U \otimes V$. The result is²⁰

$$U \otimes V |\Gamma\rangle = \frac{1}{\sqrt{d}} \sum_{i,j} U^i_k V^j_l \Gamma^{kl} |e_i\rangle |e_j\rangle = |U\Gamma V^T\rangle . \quad (92)$$

This is alarming at first sight. Let us choose $\Gamma^{ij} = \delta^{ij}$. This gives the *Jamiolkowski state* $|\delta\rangle$,

$$|\delta\rangle = \frac{1}{\sqrt{d}} \sum_{i=0}^{d-1} |e_i\rangle |e_i\rangle . \quad (93)$$

We see that

$$U \otimes \mathbf{1} |\delta\rangle = |U\rangle , \quad \mathbf{1} \otimes V |\delta\rangle = |V^T\rangle . \quad (94)$$

²⁰Exercise: Perform the calculation, using Dirac notation very carefully.

Everything that Alice can do acting with $U \otimes \mathbf{1}$ can be done by Bob as well, if he acts with $\mathbf{1} \otimes U^T$. Does this mean that information can be transmitted by the state in a non-local way?

Recall that all the information that Alice can extract from the state is given by the partial trace of the density matrix, and similarly for Bob. So we calculate

$$\rho_1 = \text{Tr}_2 |\delta\rangle\langle\delta| = \frac{1}{d} \mathbf{1}, \quad \rho_2 = \text{Tr}_1 |\delta\rangle\langle\delta| = \frac{1}{d} \mathbf{1}. \quad (95)$$

In both cases this is the maximally mixed state for dimension d .²¹ It means that neither Alice nor Bob can make any useful prediction about the experiments that they carry out on their own, using operators of the form $A \otimes \mathbf{1}$ and $\mathbf{1} \otimes B$ respectively. This is true for Bob regardless of what Alice does to the state, and conversely. Alice and Bob will have to cooperate, and use non-local operations, if they want something more specific to come out of the Jamiołkowski state. No information was transmitted in (94).

Let us now consider the reduced states of an arbitrary pure state in $\mathbf{C}^{d_1} \otimes \mathbf{C}^{d_2}$, where $d_1 \leq d_2$. We may be interested in the situation where the first factor is, say, a qubit while the second factor describes, say, most of the rest of the world, so there is no reason for the dimensions to be the same. Let us first write the state as

$$|\psi\rangle = \sum_{i=0}^{d_1-1} \sum_{j=0}^{d_2-1} \Gamma_{ij} |\tilde{e}_i\rangle |\tilde{f}_j\rangle = \sum_{i=0}^{d_1-1} |\tilde{e}_i\rangle |\tilde{\phi}_i\rangle. \quad (96)$$

The second step is always possible, but there is absolutely no reason why the d_1 states $|\tilde{\phi}_i\rangle$ in \mathbf{C}^{d_2} should be orthonormal. We can still calculate the partial trace. The centre piece of the calculation makes use of (88), and the observation that

$$\sum_{k=0}^{d_2-1} \langle k_2 | \tilde{\phi}_i \rangle \langle \tilde{\phi}_j | k_2 \rangle = \sum_{k=0}^{d_2-1} \langle \tilde{\phi}_j | k_2 \rangle \langle k_2 | \tilde{\phi}_i \rangle = \langle \tilde{\phi}_j | \tilde{\phi}_i \rangle. \quad (97)$$

We obtain

$$\rho_1 = \text{Tr}_2 |\psi\rangle\langle\psi| = \sum_{i=0}^{d_1-1} \sum_{j=0}^{d_1-1} \langle \tilde{\phi}_j | \tilde{\phi}_i \rangle |\tilde{e}_i\rangle \langle \tilde{e}_j|. \quad (98)$$

²¹Exercise: Perform the calculation. Also try the two qubit state $|\psi\rangle = z_0|00\rangle + z_1|11\rangle$.

We now have a Hermitian $d_1 \times d_1$ matrix on our hands. We diagonalize it by introducing a new basis $|e_i\rangle$. The eigenvalues will be real and positive, so there exist positive numbers p_i , summing to one, such that

$$\rho_1 = \sum_{i=0}^{d_1-1} p_i |e_i\rangle \langle e_i|. \quad (99)$$

The eigenbasis of ρ_1 is a preferred basis for \mathbf{C}^{d_1} , for any problem in which ρ_1 plays a leading role. So we go back to (96) and repeat the calculation using it, dropping the tildes as we go. This time the result (98) takes a simple form:

$$|\psi\rangle = \sum_{i=0}^{d_1-1} |e_i\rangle |\phi_i\rangle \quad \Rightarrow \quad \rho_1 = \sum_{i=0}^{d_1-1} \sum_{j=0}^{d_1-1} \langle \phi_j | \phi_i \rangle |e_i\rangle \langle e_j|. \quad (100)$$

But this time we are using a basis in which ρ_1 is known to take the form (99). We can therefore conclude that

$$\langle \phi_j | \phi_i \rangle = p_i \delta_{ij}. \quad (101)$$

Rescaling the vectors $|\phi_i\rangle$ with a factor $\sqrt{p_i}$ we have the first d_1 vectors in an orthonormal basis $\{|f_j\rangle\}_{j=0}^{d_2-1}$ for \mathbf{C}^{d_2} . We have therefore proved a theorem saying that by choosing bases in the factors suitably any given pure state in $\mathbf{C}^{d_1} \otimes \mathbf{C}^{d_2}$ can be written on the form

$$|\psi\rangle = \sum_{i=0}^{d_1-1} \sqrt{p_i} |e_i\rangle |f_i\rangle. \quad (102)$$

This standard form of $|\psi\rangle$ is called its *Schmidt decomposition*. The reduced density matrices ρ_1 and ρ_2 have the same spectrum. Their eigenvalues p_i are *entanglement invariants* in the sense that they cannot be changed by local unitary operations. A pure state is separable if the reduced states are pure. A state is *maximally entangled* if the reduced states are maximally mixed.

[It would have been better to do this by means of a *singular value decomposition* of the matrix C in (96). Given a matrix, Mathematica does it for you in no time. Then I could also have explained the polar form $C = RU$.]

Entanglement theory

There is much more to the story. It is usually told from the perspective of *local operations and classical communication*. Then the operations are unitary transformations and measurements associated to operators like $A \otimes \mathbf{1}$ and $\mathbf{1} \otimes B$. The parties are allowed to communicate with each other, so that Alice can suggest to Bob which particular operation he should apply. The theory treats local operations as if they were for free, while entangling operations are expensive. Separable states can be transformed into each other for free, and pure states with isospectral reduced density matrices can be transformed into each other for free, but it is expensive to turn a separable state into an entangled one. We can define the *entanglement cost* of a state ρ as the minimum number m of maximally entangled states needed to create n copies of ρ , or more precisely as the quotient m/n in the limit of large n . We can define the *distillable entanglement* of a state ρ similarly, in terms of the number of maximally entangled states one can create from a number of copies of ρ . (Curiously, these are not equal.)

And we can generalize bipartite entanglement to tripartite or multipartite entanglement. Let me just say that this is not easy. One reason is that the Schmidt decomposition works only in the bipartite case.

The no-cloning property

Here is a simple observation one should know about: It is impossible to construct a machine that inputs an arbitrary quantum state and outputs two copies of the same state. Let us prove this for pure states. Our machine effects some unitary transformation. We need a unitary transformation in a two-particle Hilbert space $\mathcal{H} \otimes \mathcal{H}$, so that

$$U(|\psi\rangle \otimes |0\rangle) = |\psi\rangle \otimes |\psi\rangle , \quad (103)$$

for every state $|\psi\rangle$. Let $|e_1\rangle$ and $|e_2\rangle$ be two orthogonal states in \mathcal{H} . It must be the case that

$$U(|e_1\rangle \otimes |0\rangle) = |e_1\rangle \otimes |e_1\rangle , \quad U(|e_2\rangle \otimes |0\rangle) = |e_2\rangle \otimes |e_2\rangle . \quad (104)$$

Now consider the state $|\psi\rangle = z_1|e_1\rangle + z_2|e_2\rangle$. By linearity

$$\begin{aligned} U(|\psi\rangle \otimes |0\rangle) &= U(z_1|e_1\rangle \otimes |0\rangle + z_2|e_2\rangle \otimes |0\rangle) = \\ &= z_1U(|e_1\rangle \otimes |0\rangle) + z_2U(|e_2\rangle \otimes |0\rangle) = z_1|e_1\rangle \otimes |e_1\rangle + z_2|e_2\rangle \otimes |e_2\rangle \neq \quad (105) \\ &\neq (z_1|e_1\rangle + z_2|e_2\rangle) \otimes (z_1|e_1\rangle + z_2|e_2\rangle) = |\psi\rangle \otimes |\psi\rangle . \end{aligned}$$

End of proof.

Note that were it possible to clone an unknown state, we could perform a large number of measurements on a large number of copies, and in this way determine the state. But this way is blocked.

The no-cloning theorem is not the end of the story. We can still ask for a unitary operator that minimizes the maximal copying error over all possible states. The answer is known, but we go on to other matters.

Quantum teleportation

It is time to introduce Alice and Bob, the pair that does all the quantum communication so far discussed in the literature. Alice is in possession of a system in an unknown quantum state. She wants to transmit this state to Bob without moving her actual system. The idea is that she simply reaches for a classical telephone and sends instructions enabling Bob to put a system of his in the same state. Given that unknown quantum states cannot be copied, and given that Alice does not know what state $|\psi\rangle$ her system is in—indeed, by the quantum theory of measurement, she cannot find it out—this sounds like a tall order.

I will first present the solution in complete generality, and afterwards specialize to qubits to make sure that we understand what is going on. The starting point is a nice error basis. Suppose that $\{U_I\}_{I=0}^{d^2-1}$ is a nice error basis (or, more generally, a unitary operator basis—since we are not doing quantum error correction yet, it does not have to form a group). Using one of these unitaries to replace the matrix Γ in (90) gives a maximally entangled state, as you can check by taking the partial trace of $|U_I\rangle\langle U_I|$. To be explicit about it,

$$|U\rangle\langle U| = \frac{1}{d} \left(\sum_{i,j} U^{ij} |i\rangle\langle j| \right) \left(\sum_{k,l} \langle k|\langle l| U_{kl}^* \right) = \frac{1}{d} \sum_{i,j,k,l} |i\rangle\langle j| U^{ij} U_{kl}^* \langle k|\langle l|. \quad (106)$$

Now take the partial traces following the recipe (87), remembering that U is a unitary matrix. But it is also easy to see that

$$\langle U_I | U_J \rangle = \frac{1}{d} \text{Tr} U_I^\dagger U_J = \delta_{IJ}, \quad (107)$$

because we assumed that the matrices form a unitary operator basis. To sum up, from a unitary operator basis we have created an orthonormal basis consisting solely of maximally entangled states.²² Alice must be sophisticated enough so that she can perform a von Neumann measurement using this basis.

Next Alice and Bob create a Jamiołkowski state, and share it out between them. The total Hilbert space is a tensor product of three factors (of equal dimension). Alice controls the first two factors, Bob (situated elsewhere) controls the third. The shared entangled state sits in the tensor product of the last two factors. The state to be teleported, $|\psi\rangle$, is in the first factor.

When the teleportation is about to start, the state is

$$|\psi\rangle_1 |\delta\rangle_{23} = \mathbf{1}_{12} \otimes \mathbf{1}_3 |\psi\rangle_1 |\delta\rangle_{23} = \left(\sum_{I=0}^{d^2-1} |U_I\rangle_{12} \langle U_I| \otimes \mathbf{1}_3 \right) |\psi\rangle_1 |\delta\rangle_{23}. \quad (108)$$

The equality here is just an odd way to rewrite the expression for the state, using a special basis for the \mathcal{H}_{12} factor. It is an exercise to show that this can be continued to²³

$$|\psi\rangle_1 |\delta\rangle_{23} = \sum_{I=0}^{d^2-1} |U_I\rangle_{12} |U_I^\dagger \psi\rangle_3. \quad (109)$$

So far nothing has happened. But now Alice performs her von Neumann measurement. After looking at the outcome, she finds that the state has

²²Exercise: Write down the basis obtained from the nice error basis provided by the Pauli matrices, and for the one provided by the Weyl–Heisenberg group for $d = 3$, in fully explicit form.

²³Exercise: Do this exercise.

collapsed to an eigenstate in \mathcal{H}_{12} , namely, one of the states in the unitary operator basis. At this point Alice reaches for the classical phone, and tells Bob to apply one out of the d^2 unitary operators to the state he controls. As a result of this two-step process

$$|\psi\rangle_1|\delta\rangle_{23} \rightarrow |U_I\rangle_{12}U_I^\dagger|\psi\rangle_3 \rightarrow |U_I\rangle_{12}|\psi\rangle_3. \quad (110)$$

The transfer of the unknown qubit state $|\psi\rangle$ from Alice to Bob is complete. Alice ends up with all the entanglement. Neither of them have gained any information about what the state $|\psi\rangle$ is.

Now let us specialize to qubits. We start with the state

$$\begin{aligned} |\psi\rangle|\delta\rangle &= \frac{1}{\sqrt{2}}(z_0|0\rangle + z_1|1\rangle)(|00\rangle + |11\rangle) = \\ &= \frac{z_0}{\sqrt{2}}|00\rangle|0\rangle + \frac{z_0}{\sqrt{2}}|01\rangle|1\rangle + \frac{z_1}{\sqrt{2}}|10\rangle|0\rangle + \frac{z_1}{\sqrt{2}}|11\rangle|1\rangle. \end{aligned} \quad (111)$$

We introduce the maximally entangled *Bell basis*

$$\begin{aligned} |U_0\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), & |U_1\rangle &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle), \\ |U_2\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle), & |U_3\rangle &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle). \end{aligned} \quad (112)$$

A quick calculation shows that

$$\begin{aligned} \langle U_0|00\rangle &= \frac{1}{\sqrt{2}} = \langle U_0|11\rangle, & \langle U_1|00\rangle &= \frac{1}{\sqrt{2}} = -\langle U_1|11\rangle, \\ \langle U_2|01\rangle &= \frac{1}{\sqrt{2}} = \langle U_2|10\rangle, & \langle U_3|01\rangle &= \frac{1}{\sqrt{2}} = -\langle U_3|10\rangle. \end{aligned}$$

The remaining scalar products vanish. We can now take the step from (108) to (109), and obtain

$$\begin{aligned}
|\psi\rangle|\delta\rangle = & \frac{1}{2}|U_0\rangle(z_0|0\rangle + z_1|1\rangle) + \frac{1}{2}|U_1\rangle(z_0|0\rangle - z_1|1\rangle) + \\
& \frac{1}{2}|U_2\rangle(z_0|1\rangle + z_1|0\rangle) + \frac{1}{2}|U_3\rangle(z_0|1\rangle - z_1|0\rangle) .
\end{aligned}
\tag{113}$$

A Bell measurement by Alice will not tell her what the amplitudes z_0 and z_1 are, but it does tell her what action Bob must take to transform his state into a copy of the original. The choice is between four possibilities, so the message she sends contains exactly two classical bits, in the language we will introduce later.

The BB84 protocol

This introduction would be incomplete if we did not give an example of a quantum communication protocol that one can actually buy on the market. (I do not know if the companies that sell the hardware have made profits, but that is another matter.) It concerns cryptography. Readers of *Poe* and *Conan Doyle* will know that encrypted messages can be cracked by statistical analysis. There is a way of encrypting messages that will never yield to such methods, however. Suppose Alice wishes to encrypt a message to be sent in binary digits, such as 000111000 (or preferably something longer, to which statistical analysis can be applied). Suppose that Alice and Bob share a sequence of random digits of the same length, such as 101101100. (I constructed this sequence by flipping a coin. This is not a very good method. In fact, since random sequences are valuable, constructing them by quantum mechanical means may be a useful application of quantum mechanics.) Alice now adds the two sequences digit by digit, and sends the sequence 101010100 to Bob. Bob subtracts the random sequence from the sequence he received, and obtains the sequence that Alice wanted to convey. This way of encrypting a message is known as the *Vignère* cipher, and it is in principle unbreakable because the sequence being sent shares complete randomness with the key. The catch with the idea is that it assumes that Alice and Bob do share copies of the same random sequence. The BB84 protocol offers a way of sending such a random sequence over an open channel, with a built-in guarantee that Alice

and Bob can detect whether an eavesdropper has been listening in.

Alice starts with a secret sequence of 0s and 1s, rather longer in fact than the one needed to encrypt the message. She selects, at random and in secret, either X or Z , and encodes a binary digit by preparing a qubit in an eigenstate of the gate she selected. She sends the qubit to Bob, who selects either X or Z and measures the outcome. If they made the same choices, the binary digit has been successfully transferred (because we are ignoring all practical difficulties with noise in the communication channel). If not, Bob's digit is not at all determined by Alice's. After doing this many times, Alice announces, over a telephone line that anyone can listen in to, what choices she made in the preparation. Bob replies with a list of those measurements he made where his choice agrees with Alice's. They both delete those elements of the sequence for which their choices disagreed. Having done so, they share identical random sequences, selected at random from the longer sequence that Alice started out with.

The claim is that this is an absolutely safe way to transmit a random sequence. To see why this is so, suppose an eavesdropper (called Eve) tries to listen in. Since no information about the actual digits has been sent over the telephone line, she must inform herself about the state of the qubits that were sent. This means that she must intercept and measure a qubit. But should she measure in the X or in the Z basis? She has to make a choice. Sometimes her choice will not be the same as Alice's, and then Alice's preparation is undone. When she sends the qubit on to Bob (to hide the fact that she is listening in) she destroys the perfect correlation between Alice and Bob that should have been there whenever they made the same choices between X and Z .

This means that Alice and Bob can test whether Eve has been listening. Alice sends a part of the shared random sequence to Bob over the public telephone line. Bob examines it and if it agrees with what he has, he knows that no-one has tampered with the qubit transmission. He informs Alice accordingly, and the secretive pair can use the remainder of the shared random sequence to encipher their messages using the unbreakable Vignère cipher.

Bell inequalities

Bell inequalities were introduced in the 1960ies by *Bell* to show that quantum

mechanics gives predictions that cannot be obtained from any ‘hidden variable’ theory. We do not discuss hidden variables here, but it is to be noted that the inequalities arise in classical probability theory. The quantum part of the story is precisely that they do not hold, in situations where you might have expected them to hold.

We begin by defining a few *correlation polytopes*. A polytope is the convex hull of a finite set of points, and is bounded by a finite set of faces (called *facets* if their dimension is larger than two). A simple correlation polytope is obtained by first choosing three events a_i , each of which may occur in some experiment. We can make a truth table for the events, as follows

$$\begin{array}{c|cccccccc} a_1 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 1 \\ a_2 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 1 \\ a_3 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 \end{array} \quad (114)$$

The first column covers the case when none of the events happen, the second when only a_1 happens, and so on. Each column can be regarded as a vector in a three-dimensional space, and together these vectors form the corners of a cube—which is a polytope. If the experiment is repeated N times each event will occur n_i times, and is hence associated with a frequency $\nu_i = n_i/N$. The argument to follow applies directly to the observed frequencies as well as to the probabilities that emerge in the limit.

The possible outcomes of the N experiments, taken together, can be labelled by three numbers p_i , which by construction obey the inequalities

$$0 \leq p_1 \leq 1, \quad 0 \leq p_2 \leq 1, \quad 0 \leq p_3 \leq 1. \quad (115)$$

This is related to a dual description of the cube in terms of these six inequalities, each of which says that the cube lies on one side of some two-dimensional plane. Every compact convex set admits two dual descriptions along these lines, one of which describes the body as the convex hull of a set of pure points, and one of which describes it as the intersection of a set of half-spaces. In both cases we are looking for a minimal description of its kind.

Things get more interesting if there are logical connections between the events. For instance, we can assume that the three events are mutually exclusive, so that exactly one of the events happens in each run of the experiment. Only three columns of the truth table survives, there are only three pure points left, and the observed frequencies necessarily obey the condition

$$p_1 + p_2 + p_3 = 1 . \tag{116}$$

This is the case we started out with.

Another interesting restriction is to impose the logical condition that $a_3 = a_1 \& a_2$, that is to say that the event a_3 is that both a_1 and a_2 happen. This gives the truth table

$$\begin{array}{c|cccc} a_1 & 0 & 1 & 0 & 1 \\ a_2 & 0 & 0 & 1 & 1 \\ a_1 \& a_2 & 0 & 0 & 0 & 1 \end{array} \tag{117}$$

The convex hull of the four column vectors is a polytope inscribed in the cube. Its dual description in terms of inequalities describing the faces of the polytope is

$$p_3 \geq 0 , \quad p_1 \geq p_3 , \quad p_2 \geq p_3 , \tag{118}$$

$$p_1 + p_2 - p_3 \leq 1 . \tag{119}$$

The last of these may be a little hard to see, but—like the other three faces—it defines a face passing through exactly three of the vertices, namely in this case the last three of the columns of the truth table. See Fig. 5. Slightly more involved examples of this last kind of inequality are known as *Bell inequalities*. We will see that they may fail to hold in quantum theory, and must try to understand why.




Figure 5: A correlation polytope for three logically connected events.

To get an interesting Bell inequality, consider eight events $a_1, a_2, a_3, a_4, a_1 \& a_3, a_1 \& a_4, a_2 \& a_3, a_2 \& a_4$. The truth table will produce vectors with eight components, in fact 16 vectors altogether, and we do not write it out here. The inequalities of interest to us are

$$-1 \leq p_{1,3} + p_{1,4} + p_{2,4} - p_{2,3} - p_1 - p_4 \leq 0 \quad (120)$$

$$-1 \leq p_{2,3} + p_{2,4} + p_{1,4} - p_{1,3} - p_2 - p_4 \leq 0 \quad (121)$$

$$-1 \leq p_{1,4} + p_{1,3} + p_{2,3} - p_{2,4} - p_1 - p_3 \leq 0 \quad (122)$$

$$-1 \leq p_{2,4} + p_{2,3} + p_{1,3} - p_{1,4} - p_2 - p_3 \leq 0 . \quad (123)$$

You can check that they hold by making all possible truth assignments to the four events a_1, a_2, a_3, a_4 .²⁴

To each event a_i there will correspond an observable α_i taking values $+1$ if the event happens, 0 if it does not. There will be corresponding expectation values

$$p_i = \langle \alpha_i \rangle , \quad p_{i,j} = \langle \alpha_i \alpha_j \rangle . \quad (124)$$

It is customary to introduce observables $A_i = 2\alpha_i - 1$, and to rewrite the inequalities in terms of

$$\langle A_i A_j \rangle = \langle 4\alpha_i \alpha_j - \alpha_i - \alpha_j + 1 \rangle = 4p_{i,j} - 2p_i - 2p_j + 1 . \quad (125)$$

Finally we rename A_3 as B_1 and A_4 as B_2 . Inequality (120) then takes the form

$$-2 \leq \langle A_1 B_1 + A_1 B_2 + A_2 B_2 - A_2 B_1 \rangle \leq 2 . \quad (126)$$

Written in this form it is known as the *Clauser–Horne–Shimony–Holt* inequality. It has to hold for the individual outcomes and for the observed frequencies in any experiment measuring these random variables.

²⁴Exercise: Do it. Also rewrite (120) in the form (126).

Now let A_1, A_2 be operators corresponding to measurements that can be made by Alice, while B_1, B_2 can be made by Bob. The possible outcomes are ± 1 in all cases. In a single run of the experiment only one of $A_1 \otimes B_1, A_1 \otimes B_2, A_2 \otimes B_1, A_2 \otimes B_2$, can be measured. But many measurements of each kind will be made, and it seems reasonable to assume that, once the statistics is collected, it will be true that

$$\langle A_1 B_1 + A_1 B_2 + A_2 B_2 - A_2 B_1 \rangle = \langle A_1 B_1 \rangle + \langle A_1 B_2 \rangle + \langle A_2 B_2 \rangle - \langle A_2 B_1 \rangle . \quad (127)$$

In quantum mechanics the left and right hand side are necessarily equal to each other, and both can be computed once we know what state the system is in. The argument for the CHSH inequality applies to the left hand side. The right hand side will be measured. Its first term is evaluated by collecting the statistics from those instances of the experiment in which the settings were made so that $A_1 \otimes B_1$ was measured, and so on. The system, in most experiments, consists of a pair of photons in a carefully prepared entangled state. In the best experiments the choice of settings is done when the pair of photons are well on their way, and every effort is made to ensure that the choice is made in a random fashion on both sides. But now we do have a problem, namely Problem 1.

Problem 1: To test the CHSH inequality Alice and Bob use

$$\begin{aligned} A_1 &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \otimes \mathbf{1} , & A_2 &= \begin{pmatrix} 0 & e^{-i\alpha} \\ e^{i\alpha} & 0 \end{pmatrix} \otimes \mathbf{1} , \\ B_1 &= \mathbf{1} \otimes \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} , & B_2 &= \mathbf{1} \otimes \begin{pmatrix} 0 & e^{-i\alpha} \\ e^{i\alpha} & 0 \end{pmatrix} . \end{aligned} \quad (128)$$

They can vary α . What do these measurements correspond to if they are measuring photon polarization? What is the largest value they can obtain for the quantity in the CHSH inequality? What quantum state gives this value?

Problem 2: Consider two quantum states, anywhere on the Bloch sphere. Choose a Hermitian matrix A , and expand the two states in its eigenbasis. Calculate the

Fisher–Rao distance between the two states, in terms of the expansion coefficients, for the two-outcome measurement corresponding to A . Prove that by varying A this distance can be made equal to the Fubini–Study distance, but that it cannot be made larger.

OPEN SYSTEMS

The theory of open systems deals with situations where the system of interest is not fully isolated from its environment. It regards the world as divided into a ‘system’ and a ‘reservoir’, where the latter is not controlled by the experimentalist. Sometimes the reservoir is useful when one wants to manipulate the system, and then it may be called an ‘ancilla’. If the system becomes entangled with the reservoir the split of the whole into its parts becomes a subtle matter. The brief account here is mostly concerned with the theory of quantum measurements.

The chicken and the egg

Let us summarize the message of the Schmidt decomposition (102):

- Let ρ_{12} be a pure state on $\mathcal{H}_{12} = \mathcal{H}_1 \otimes \mathcal{H}_2$. Then the reduced states $\rho_1 = \text{Tr}_2 \rho_{12}$ and $\rho_2 = \text{Tr}_1 \rho_{12}$ have the same non-zero eigenvalues.
- Given a state ρ_1 on \mathcal{H}_1 there exists a Hilbert space \mathcal{H}_2 and a pure state $|\psi\rangle$ in \mathcal{H}_{12} such that $\rho_1 = \text{Tr}_2 |\psi\rangle\langle\psi|$.

In these notes density matrices were introduced first, and Hilbert space vectors were added almost as an afterthought. Many accounts start at the other end, density matrices appearing only when the state of the system has not been specified as completely as it could be. What is the correct starting point? The Schmidt decomposition turns this into a chicken-and-egg question because any mixed state can be *purified*, and regarded as a pure state in a larger Hilbert space partly outside our control.

Either way we have a concrete question concerning time evolution to discuss. Were the system isolated it would evolve *unitarily* according to

$$\rho \rightarrow U\rho U^\dagger . \quad (129)$$

If the unitary is of the form $U = e^{-iHt}$ this becomes the differential equation

$$\frac{d\rho}{dt} = i[\rho, H] . \quad (130)$$

We will want to know how the density matrix on \mathcal{H}_1 evolves if we extend it to a state on the larger Hilbert space $\mathcal{H}_{12} = \mathcal{H}_1 \otimes \mathcal{H}_2$, evolve that state with a unitary, and take the partial trace down to \mathcal{H}_1 at the end.

CP maps

It is convenient to assume that there are no correlations between the state and its environment to begin with. In fact we assume that the initial state is

$$\rho_S \otimes \rho_R^0, \quad (131)$$

where the initial state ρ_R^0 of the reservoir is pure

$$\rho_R^0 = |R_0\rangle\langle R_0|. \quad (132)$$

The second assumption is harmless because we can consider a Hilbert space large enough so that it holds. The assumption that we start out with a product state is not harmless, but it is an interesting one to make. It holds if the environment quickly forgets about its past interactions with the system. To see whether this is so requires a detailed analysis of the physics. It seems plausible that it holds if the reservoir consists of photons that quickly disappear to very large distances.

We now evolve the state (131) with a unitary U acting on the full Hilbert space, and then take the partial trace using an orthonormal basis for \mathcal{H}_R . We find

$$\rho_S \rightarrow \text{Tr}_R U (\rho_S \otimes \rho_R^0) U^\dagger = \sum_i \langle i_R | U | R_0 \rangle \rho_S \langle R_0 | U^\dagger | i_R \rangle. \quad (133)$$

The sum will have to include all non-zero terms. Whatever their number is, the answer has an interesting form.

We define the *Kraus operators* A_i by

$$A_i = \langle i_R | U | R_0 \rangle. \quad (134)$$

If you do not lose track of the fact that U operates on the Hilbert space $\mathcal{H}_S \otimes \mathcal{H}_R$, it will be obvious that this is an operator on the Hilbert space of

the system.²⁵ By construction

$$\sum_i A_i^\dagger A_i = \sum_i \langle R_0 | U^\dagger | i_R \rangle \langle i_R | U | R_0 \rangle = \langle R_0 | U^\dagger U | R_0 \rangle = \mathbf{1}_S . \quad (135)$$

Conversely, it can be shown that every set of operators A_i obeying the last equality can be obtained in this way from some unitary U .

We have arrived at the notion of a *completely positive map*, which is the quantum version of the linear maps (10) that take probability vectors to probability vectors. A map Φ taking density matrices to density matrices is completely positive if and only if there exists a set of Kraus operators A_i such that

$$\rho \rightarrow \Phi(\rho) = \sum_i A_i \rho A_i^\dagger \quad \text{where} \quad \sum_i A_i^\dagger A_i = \mathbf{1} . \quad (136)$$

This is clearly a linear map,

$$a_1 \rho_1 + a_2 \rho_2 \rightarrow a_1 \sum_i A_i \rho_1 A_i^\dagger + a_2 \sum_i A_i \rho_2 A_i^\dagger . \quad (137)$$

If you rearrange the density matrix so that it forms a vector with d^2 components, the CP map Φ becomes a $d^2 \times d^2$ matrix. If you like, it is a *superoperator* acting on operators. The trace and the positivity of ρ is preserved by the map, so it takes density matrices to density matrices.²⁶ But positivity is not enough. The extra ingredient leading to the Kraus form comes about in a rather strange way, as we will see.

Positive maps and entangled states

A linear map from matrices to matrices is said to be *positive* if it takes positive matrices to positive matrices. It is very difficult to describe positive maps in general, but a simple example of a positive map is $\rho \rightarrow \rho^T$. The transposed matrix has the same spectrum as the original, so this is clearly a

²⁵Exercise: Expand U in a product basis for the composite Hilbert space (letting $|R_0\rangle$ be one of the basis vectors in one of the factors). Then calculate the Kraus operators. Note that many different U s correspond to the same set of A_i .

²⁶Exercise: Prove this.

trace preserving positive map. But suppose our Hilbert space is $\mathbf{C}^2 \otimes \mathbf{C}^2$. As usual, Alice controls only the first factor, so she performs a *partial transpose* of the density matrix. That is, using the notation (82), she performs the map

$$\rho^{i\alpha}_{j\beta} \rightarrow \rho^{j\alpha}_{i\beta} . \quad (138)$$

A little more abstractly, this is

$$\rho \rightarrow \rho^{\text{TA}} . \quad (139)$$

(It is understood that Alice ‘performs’ the map using pen and paper only. It does not correspond to a physical transformation, as we will see.) Let us take it that the density matrix is pure,

$$\rho = |\psi\rangle\langle\psi| , \quad |\psi\rangle = \sqrt{p_0}|0\rangle|0\rangle + \sqrt{p_1}e^{i\nu}|1\rangle|1\rangle , \quad p_0 + p_1 = 1 . \quad (140)$$

The problem is that when the partial transposition is performed the spectrum of the density matrix changes according to

$$(1, 0, 0, 0) \rightarrow (p_0, p_1, \sqrt{p_0 p_1}, -\sqrt{p_0 p_1}) . \quad (141)$$

Unless the state vector is a product vector negative eigenvalues appear. The matrix is not a density matrix any more.²⁷

The definition of completely positive maps avoids this difficulty. Let us represent positive maps as $d^2 \times d^2$ matrices Φ . The density matrices on which they act to act in their turn on a Hilbert space \mathcal{H}_S . This Hilbert space can be enlarged to a Hilbert space $\mathcal{H}_S \otimes \mathcal{H}_R$, but let us agree that the second factor is irrelevant, so that we perform only positive maps of the form $\Phi \otimes \mathbf{1}$. The partial transposition is of this form. By definition the map Φ is said to be completely positive if $\Phi \otimes \mathbf{1}$ is positive for all possible extensions $\mathcal{H}_S \rightarrow \mathcal{H}_S \otimes \mathcal{H}_R$.

Stinespring’s dilation theorem: a map is completely positive and trace preserving if and only if it can be written in the Kraus form (136).

This is the extra ingredient that led us to the Kraus form.

²⁷Exercise: Do this calculation. Do the same for the case when Bob takes the partial transpose.

It took some twenty years for physicists to understand Stinespring's paper, and Kraus was one of the first to do so. Once you know the result, it is easy to remember and to use. Moreover every completely positive map can be performed in a hypothetical lab equipped to perform every unitary transformation, in Hilbert spaces whose dimension is at most the square of that of the Hilbert space of the physical system being studied.

There is another interesting aspect of this story. We can ask: when is a quantum state entangled? If the state is pure the answer is simple: if its partial trace is a mixed state. But what if the state itself is mixed? There can be classical correlations present, and in fact there will be whenever the state is *not* of the form $\rho^A \otimes \rho^B$. A suitable definition of a separable quantum state, whether mixed or pure, should allow for classical correlation between the two subsystems, but no more than that. Mathematically this means that the separable state lies in the convex hull of uncorrelated states:

Definition: a state ρ is separable if and only if it can be written as

$$\rho = \sum_i p_i \rho_i^A \otimes \rho_i^B, \quad (142)$$

for some density matrices ρ_i^A, ρ_i^B acting on the factors.

This definition is due to *Werner*.

A direct check whether a given state ρ obeys this condition is prohibitively difficult, since we have to try all possible decompositions of this type. It is therefore helpful to know that a state is entangled if and only if there exists a positive but not completely positive map Φ such that $\Phi \otimes \mathbf{1}$ turns the state ρ into a matrix with at least one non-negative eigenvalue. This is still a hard condition to check when the dimensions involved are large. In fact it is an **NP** complete problem, in the language of complexity theory that we will introduce later. But for two qubits it settles things.

*Theorem: A state ρ is entangled if ρ^{TA} has a negative eigenvalue. For two qubits, and for one qubit and one qutrit, this is an if-and-only-if statement.*²⁸

²⁸Exercise: Consider the two qubit state $\rho = \frac{p}{2}(|00\rangle + |11\rangle)(\langle 00| + \langle 11|) + \frac{1-p}{4}\mathbf{1}$. For what values of p is this an entangled state?

Measurements

Let us agree that an isolated quantum system always evolves unitarily, while an open system—perhaps the system is coupled to a measurement apparatus—evolves with completely positive trace preserving maps. In the first step (136) applies, for some choice of Kraus operators A_i . But in a measurement the wave function must collapse to an outcome. Rushing in where angels fear to tread, we resolve the measurement problem with a postulate:

Let there be n possible measurement outcomes. Then there are n Kraus operators A_i . In a non-selective measurement the system changes its state according to the CP map generated by the A_i . In a selective quantum measurement the system changes its state from ρ to one of the n states

$$\rho_i = \frac{A_i \rho A_i^\dagger}{\text{Tr}(A_i \rho A_i^\dagger)} . \quad (143)$$

The transition $\rho \rightarrow \rho_i$ happens with probability

$$p_i = \text{Tr} A_i \rho A_i^\dagger . \quad (144)$$

All the ρ_i are legitimate density matrices and the Kraus operators A_i obey a condition ensuring that the probabilities sum to unity.

We now have two kinds of time evolution, the linear time evolutions effected by CP maps, and the non-linear time evolution $\rho \rightarrow \rho_i$ described by the measurement postulate. An important special case is that of a *projective* or *von Neumann* measurement. Then we choose the measurement operators to be mutually orthogonal projection operators,

$$A_i = P_i = A_i^\dagger , \quad P_i P_j = \delta_{ij} P_i . \quad (145)$$

The completeness property of the measurement operators implies that

$$\sum_i P_i = \mathbf{1} . \quad (146)$$

Sets of projection operators like this are obtained by choosing a Hermitian operator, now called an ‘observable’, and performing a *spectral decomposition*

$$A = \sum_i \lambda_i P_i . \quad (147)$$

For simplicity, suppose that all eigenvalues are non-degenerate. In a non-selective von Neumann measurement of the ‘observable’ A the state changes according to

$$\rho \rightarrow \rho' = \sum_{i=1}^d P_i \rho P_i . \quad (148)$$

The state has been forced to commute with A . This is in fact a CP map effected by the d Kraus operators P_i . In a selective von Neumann measurement the state collapses, and the outcome labelled λ_i occurs with probability p_i ,

$$\rho \rightarrow \rho_i = \frac{P_i \rho P_i}{\text{Tr}(P_i \rho P_i)} , \quad p_i = \text{Tr}(P_i \rho P_i) = \text{Tr}(\rho P_i) . \quad (149)$$

This measurement is repeatable.²⁹ It is also highly idealized. Still the von Neumann measurement is much beloved by people who have the task of actually measuring things in the lab. The picture to have in mind is that of a photon encountering a Glan–Thompson prism that lets through linearly polarized photons only.

If a probability vector counts as a classical state, it also can collapse. A classical probability distribution $P(A)$ can collapse to a conditional probability distribution $P(A = a_i | B = b_j)$ depending on the outcome observed for the random variable B . But the non-selective measurement does not have an analogue. In a non-selective classical measurement nothing happens since

$$\sum_j P(A = a_i | B = b_j) P(B = b_j) = P(A = a_i) . \quad (150)$$

The analogy breaks down because the quantum state ρ' obtained in (148) may be significantly different from the state ρ . Very special cases excepted a quantum measurement always disturbs the state, and the classical notion of ‘conditional probability’ is no longer with us.

²⁹Exercise: Show that the expectation value $\langle A \rangle = \text{Tr}(A\rho)$. Also show that if you repeat the same measurement you get the same result.

This is an important conceptual point, so let us say it in a different way. Suppose a pure quantum state evolves under a unitary transformation,

$$|\psi(t)\rangle = z_0(t)|0\rangle + z_1(t)|1\rangle = U(t)|\psi(0)\rangle . \quad (151)$$

If, at time t , we perform a measurement in this basis we find the outcome $|0\rangle$ with probability $|z_0(t)|^2$. But what is the probability to get this outcome at time $t = 1$ given the outcome at $t = 0$? We cannot talk about the outcome at $t = 0$ unless a measurement was performed at that time. But if so the state at $t = 0$ changed according to

$$|\psi(0)\rangle\langle\psi(0)| = \begin{pmatrix} |z_0|^2 & z_0\bar{z}_1 \\ z_1\bar{z}_0 & |z_1|^2 \end{pmatrix} \rightarrow \rho(0) = \begin{pmatrix} |z_0|^2 & 0 \\ 0 & |z_1|^2 \end{pmatrix} . \quad (152)$$

The off-diagonal elements have disappeared, and

$$U(1)\rho(0)U^{-1}(1) \neq |\psi(1)\rangle\langle\psi(1)| . \quad (153)$$

So we get into a muddle if we use equation (151) to talk about the probability that the system ‘was’ in one of the states $|0\rangle$ and $|1\rangle$ at two different times.

We leave these questions open, and return to general measurements. The idea to use Kraus operators to guarantee that there is some larger Hilbert space behind us has led us to claim that a given outcome occurs with probability

$$p_i = \text{Tr}(\rho A_i^\dagger A_i) , \quad \sum_i A_i^\dagger A_i = \mathbf{1} . \quad (154)$$

The Kraus operators describe how the state of the system changes in a measurement. But maybe we are only interested in the result of the measurement, and therefore we define a general measurement by saying that there exists a number of positive operators E_i , called *effects*, such that

$$E_i \geq 0 , \quad \sum_{i=1}^n E_i = \mathbf{1} . \quad (155)$$

The collection of effects is known as a *POVM*, which can be spelt out as ‘positive operator valued measure’. The POVM allows us to extract a probability distribution from any density matrix,³⁰

³⁰Exercise: Show that this defines a probability distribution.

$$p_i = \text{Tr}(\rho E_i) . \quad (156)$$

There is no requirement that $\text{Tr} E_i E_j = 0$, which means that this is not a probability distribution over mutually exclusive events, nor is there any guarantee that we can choose a state for which p_1 (say) equals 1.³¹ Moreover, unless the effects are one-dimensional projectors they can be written in terms of Kraus operators in many different ways.

Coming back to the chicken-and-egg question, following *Naimark* we can regard any POVM as a von Neumann measurement by adding an ancilla Hilbert space to the Hilbert space of the system. Let us consider the special case of a *rank one* POVM, in which each of the $n > d$ effects is of the form $E_i = |x_i\rangle\langle x_i|$ for some vector $|x_i\rangle$ obeying $\langle x_i|x_i\rangle \leq 1$, with equality only if the effect actually is a projector. Then form the *generator matrix*

$$X = \left(\begin{array}{cccc} |x_1\rangle & |x_2\rangle & \dots & |x_n\rangle \end{array} \right)_{d \times n} = \left(\begin{array}{c} \langle u_1| \\ \langle u_2| \\ \vdots \\ \langle u_d| \end{array} \right)_{d \times n} . \quad (157)$$

We have arranged things so that we can think of X either as n columns or as d rows. You can check that

$$XX^\dagger = \sum_{i=1}^n |x_i\rangle\langle x_i| = \sum_{i=1}^n E_i = \mathbf{1}_{d \times d} . \quad (158)$$

But when you think in terms of the rows, this means that you have a set of n -dimensional vectors obeying

$$\langle u_i|u_j\rangle = \delta_{ij} . \quad (159)$$

Here $1 \leq i, j \leq d$. You can extend this set of d orthonormal row vectors to n orthonormal row vectors by adding more rows to X , so that it becomes a unitary matrix. We interpret it as a set of n column vectors,

³¹Exercise: For the qubit, choose projectors that project onto states forming a regular triangle on the equator. Rescale them so that they form a POVM. What probability distributions over three events can you obtain in this way?

$$U = \begin{pmatrix} |x_1\rangle & |x_2\rangle & \cdots & |x_n\rangle \\ |y_1\rangle & |y_2\rangle & \cdots & |y_n\rangle \end{pmatrix}_{n \times n} . \quad (160)$$

Because the matrix is unitary, the column vectors are mutually orthogonal unit vectors $\{|z_i\rangle\}_{i=1}^n$, yielding n projection operators $\{P_i\}_{i=1}^n$, where

$$|z_i\rangle = \begin{pmatrix} |x_i\rangle \\ |y_i\rangle \end{pmatrix} , \quad P_i = |z_i\rangle\langle z_i| . \quad (161)$$

This defines a von Neumann measurement in the large Hilbert space, related by a simple projection to the POVM we started out with.³²

POVMs provide an interesting perspective on quantum states, but we hasten on to the next topic.

The Lindblad equation

A system may fail—and will fail, more or less—to be isolated because it interacts with some ‘environment’ or ‘thermal reservoir’. In this situation unitary time evolution is at best an approximation to what is observed. We round off this brief introduction to the theory of open systems by giving the generalization to quantum mechanics of the classical Markov process. If the evolution of the system is a continuous unfolding of completely positive maps it is described by the *Lindblad equation*

$$\frac{d\rho}{dt} = \mathcal{L}(\rho) = i[\rho, H] + \sum_i \left(L_i \rho L_i^\dagger - \frac{1}{2}(L_i^\dagger L_i \rho + \rho L_i^\dagger L_i) \right) . \quad (162)$$

The L_i are the Lindblad operators, and \mathcal{L} stands for Liouville. The Hamiltonian H typically has contributions also from the terms that couple the system to the reservoir in the evolution of the composite system. The equation applies when the system dynamics is slow compared to the correlation time scale of the reservoir. In typical quantum optical systems, where correlations in the radiation field leak away quickly, it has fared brilliantly.

For a careful derivation you have to look elsewhere, but to see where the various terms come from consider a CP map

³²Exercise: Do this construction explicitly for the POVM in the preceding footnote.

$$\mathcal{E}(\rho) = A_0 \rho A_0^\dagger + \sum_i A_i \rho A_i^\dagger = \rho + \delta\rho , \quad (163)$$

$$A_0 = \mathbf{1} + (L_0 - iH)\delta t , \quad A_i = L_i \sqrt{\delta t} . \quad (164)$$

One of the Kraus operators is close to the identity, while the others are close to zero. Then

$$\delta\rho = (-i[H, \rho] + L_0 \rho + \rho L_0 + \sum_i L_i \rho L_i^\dagger) \delta t . \quad (165)$$

To ensure that the trace is preserved we must have $\text{Tr} \delta\rho = 0$. This forces

$$L_0 = -\frac{1}{2} \sum_i L_i^\dagger L_i . \quad (166)$$

Dividing through by δt we arrive at the Lindblad equation. It looms large in quantum optics, but we do not go into it any further here.

Problem 3: Consider the unitary operator

$$U = \mathbf{1} \otimes |0\rangle\langle 0| + X \otimes |1\rangle\langle 1| \quad (167)$$

acting on a two qubit Hilbert space (where $X = \sigma_x$). Choose the initial state to be $\rho \otimes |R_0\rangle\langle R_0|$, where

$$|R_0\rangle = \cos \frac{\theta}{2} |0\rangle + \sin \frac{\theta}{2} |1\rangle . \quad (168)$$

Construct the Kraus operators acting on ρ . What happens to the Bloch vector of ρ when we perform the resulting CP map? To what surface in the Bloch ball of the first qubit does the Bloch sphere of the first qubit go?

INFORMATION THEORY

A Mathematical Theory of Communication introduced a quantity called ‘entropy’. Following *Boltzmann* it was denoted by the letter H . The papers were later reissued as a book, now called *The Mathematical Theory of Communication*. Generalizing this theory to the quantum case is a subtle thing.

Information

Shannon worked at the Bell Telephone Laboratories, and his interest in communication was very practical. A message is being sent, using some alphabet of n letters. The theory (as presented here) assumes that the letters are *i.i.d.*, spelt out as ‘independent and identically distributed’. That is, the probability that the i th letter will be sent is p_i , every time. We want to compress the message as much as possible before sending it, and the first goal is to determine how much data compression that can be achieved.




Figure 6: The problem considered by Shannon. The message is compressed at an encoder before it is sent through a possibly noisy communication channel.

Shannon began by asking for a measure of how much ‘choice’ is involved in the selection of a letter, or how uncertain we are of the outcome of the event. Equivalently, this is the amount of information produced when the event happens. The measure he arrived at is a function of the probability distribution associated to some random variable. From now on then *information* is quantified by the *Shannon entropy*

$$H(\vec{p}) = - \sum_{i=1}^n p_i \log p_i . \quad (169)$$

It is understood that $0 \log 0 = 0$. The logarithm is usually taken to be with base 2, and then the information is measured in *bits*. We will slip back to natural logarithms at some stage.³³

We can argue for this measure as follows. We regard $-\log p_i$ as an additive measure of the ‘surprise’ we feel in receiving the i th letter. Then $H = \langle -\log p \rangle$ is the information received per letter, when averaged over a long message. The interpretation makes more sense if we remember the psychophysical law that says that human response is proportional to the logarithm of the stimulus. (This is why the Greek measured the luminosity of stars using logarithmic magnitudes.) But the real justification for the definition lies in the theorems that Shannon proved. Let us place Shannon’s *noiseless coding theorem* on the table right away. Then you will see that these lecture notes carry a certain amount of information, and that Shannon’s definition of information is the relevant one if you want to convert them to JPEG format.

To introduce the theorem we first note that when we use the logarithm with base 2 we assume that the length of the message is measured in terms of the number of binary digits you need to encode it. This length will clearly depend on the coding. Recall that the Morse alphabet uses dashes and dots, with the number of dashes and dots used to encode the letter being lower if that particular letter is in frequent use. The theorem is concerned with the length of a message that has been encoded in an optimal way. Ignoring some fine print, it says that *if a message contains N letters chosen with probabilities p_i from an alphabet consisting of n symbols, then it can be transmitted in the form of a string of bits of length $NH(\vec{p})$, but it cannot be compressed further.*

To see how the theorem comes about, consider the kind of sequences that can arise. In particular, consider the number of sequences that contains N_1 instances of the first letter in the alphabet, N_2 of the second, and so on. This number is

³³Exercise: For $p_1 = p, p_2 = 1 - p$, plot $H(p)$. Then maximize the function for arbitrary n . Also check what happens if you change the base of the logarithm to e .

$$\frac{N!}{N_1!N_2!\dots N_n!} . \quad (170)$$

Using Stirling's formula we can approximate this number, or more conveniently its logarithm, by

$$\begin{aligned} \log \frac{N!}{N_1!N_2!\dots N_n!} &\approx \\ &\approx N \log N - N - \sum_{i=1}^n (N_i \log N_i - N_i) = - \sum_{i=1}^n N_i \log \frac{N_i}{N} . \end{aligned} \quad (171)$$

We assume that all the individual distinguishable sequences are equally likely to occur. It then follows from the Law of Large Numbers that, with overwhelming probability, the sequences to be encoded have

$$\frac{N_i}{N} = p_i . \quad (172)$$

Inserting this in (171) we conclude that the Law of Large Numbers allows us say that, with overwhelming probability, the sequence to be encoded can be regarded as having been chosen from a set of

$$2^{NH(\vec{p})} \quad (173)$$

typical sequences. For large N , and unless all letters are equally likely, this is a small fraction of the number of all possible sequences (equal to n^N).

We see that the compression of the message is possible because at the encoder we can assume that we are dealing with a typical message. We do not have to encode all the 2^N messages. It is enough if the signal carries information about which out of all the typical messages is being sent. This is how the noiseless coding theorem arises. We also see that some fine print must be added to the statement of the theorem, because there is a small probability that the message is, in fact, not typical. If so we must declare an error. However, the Law of Large Numbers guarantees that the probability that this happens can be made as small as we please by making the message long enough.

Some limitations are apparent here. Let us raise some of them. To begin with, the i.i.d. assumption may be too strong. If the message is in English, or in any other natural language, there will in fact be strong correlations between the various letters in the message. We could compress the message by, say, removing all the vowels, and chances are that it would still be decodable. Entire books have been written in this way. (Incidentally, a theory taking the redundancies in the English language into account shows that the amount of information present in my notes would increase if its letters were reordered in a random way.) A second limitation of the theorem is that it is non-constructive. It tells us that compression is possible, but it does not provide a recipe for how to do it. Finally we notice that the encoding cannot even start until the encoder has received a string of letters long enough to ensure that the string is typical. Still the noiseless coding theorem provided information theory with a very good start.

Some properties of the Shannon entropy

The Shannon entropy is a continuous and nowhere negative function of a probability distribution \vec{p} , taking the value zero if and only if the probability vector is pure. It also obeys the recursion property that we illustrate in Figure 7. The idea is that a choice can be broken down in successive choices, so that

$$\begin{aligned}
 H\left(\frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}\right) &= \\
 &= H\left(\frac{3}{8}, \frac{3}{8}, \frac{1}{4}\right) + \frac{3}{8}H\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right) + \frac{3}{8}H\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right) + \frac{1}{4}H\left(\frac{1}{2}, \frac{1}{2}\right) .
 \end{aligned}
 \tag{174}$$

Supplemented by some mild extra conditions the recursion property actually defines the function H uniquely .

A key property is concavity. Let \vec{p} and \vec{q} be two probability vectors, let $x \in [0, 1]$, and consider the mixture $x\vec{p} + (1 - x)\vec{q}$. Then

$$H(x\vec{p} + (1 - x)\vec{q}) > xH(\vec{p}) + (1 - x)H(\vec{q}) .
 \tag{175}$$




Figure 7: The recursion property illustrated: to the right it is used to determine $H(\frac{3}{8}, \frac{3}{8}, \frac{1}{4})$ in terms of the Shannon entropy for uniform distributions.

In words, the Shannon entropy is a *concave* function of its arguments, which means that a straight line between two points on its graph always lies below the graph.³⁴

Concave functions, or *convex* functions for which the direction of the inequality is reversed, are important for many reasons. (To remember which is which, memorize that a convex function has a convex *epigraph*, and think of the epicentre of an earth quake.) Concave functions are easy to optimize. A concave function has at most a single maximum inside its domain.

Conditional entropy, joint entropy, and mutual information

To continue, it is convenient to associate the Shannon entropy with some particular random variable A , so that

$$H = H(A) = - \sum_i P(A = a_i) \log P(A = a_i) . \quad (176)$$

If we have two random variables we have the two probability distributions $P(A)$ and $P(B)$ to play with, as well as the joint and conditional probability distributions. (At this point, please recall eq. (3), known as *Bayes' formula*).

³⁴Exercise: Calculate the matrix of second derivatives of $H(p_1, p_2, \dots, p_n)$. Show that it is negative definite, and that this is enough to prove concavity. Where in these notes have you seen this matrix before?

Thus we have the *conditional entropy*

$$H(A|B) = - \sum_{i,j} P(B = b_j) P(A = a_i | B = b_j) \log P(A = a_i | B = b_j) \quad (177)$$

and the *joint entropy*

$$H(A, B) = - \sum_{i,j} P(A = a_i, B = b_j) \log P(A = a_i, B = b_j) . \quad (178)$$

The joint entropy measures the uncertainty of a joint event, or equivalently the information received when both events are found to happen. The conditional entropy measures the uncertainty of the event A given that we know the outcome of the event B , weighted over all the possible outcomes of B .

In the calculations to follow we will set

$$P(A = a_i | B = b_j) = p_{i|j} , \quad P(A = a_i, B = b_j) = p_{i,j} , \quad (179)$$

$$P(A = a_i) = p_i = \sum_j p_{i,j} , \quad P(B = b_i) = q_i = \sum_j p_{j,i} . \quad (180)$$

In this notation Bayes' formula (3) takes the form

$$p_{i,j} = p_{i|j} q_j = p_{j|i} p_i . \quad (181)$$

For calculational convenience we will also replace the log with the natural logarithm in all the definitions.

There are relations between the joint and conditional entropies. Using Bayes' formula it is easy to see that

$$H(A, B) = H(B) + H(A|B) = H(A) + H(B|A) . \quad (182)$$

The proof consists in writing out the definition of the conditional entropy,

$$\begin{aligned} H(A|B) &= - \sum_{i,j} p_{i,j} \ln \frac{p_{i,j}}{q_j} = - \sum_{i,j} p_{i,j} (\ln p_{i,j} - \ln q_j) = \\ &= - \sum_{i,j} p_{i,j} \ln p_{i,j} + \sum_j q_j \ln q_j = H(A, B) - H(B) . \end{aligned} \quad (183)$$

If the two random variables are uncorrelated it immediately follows that $H(A|B) = H(A)$, so that the joint entropy is just the sum of the two entropies for the individual random variables. We define one more useful quantity, the *mutual information*

$$I(A : B) = H(A) + H(B) - H(A, B) . \quad (184)$$

It equals zero if the events are uncorrelated.³⁵ An important property of the mutual information is that it is concave function. If one of its arguments is a convex mixture $\vec{p} = x\vec{r} + (1 - x)\vec{s}$ then

$$I(\vec{p}, \vec{q}) \geq xI(\vec{r}, \vec{q}) + (1 - x)I(\vec{s}, \vec{q}) . \quad (185)$$

We omit the proof, but keep it in mind. It will be useful when we have to maximize the mutual information in one of its arguments.

There are important inequalities obeyed by the various quantities that we have introduced. The most obvious one is

$$H \geq 0 \quad (186)$$

with equality only for a pure state. Two more are suggested by the interpretation. The information received when a joint event happens must be greater than that received from only one of the events, so we must have

$$H(A, B) \geq H(A) . \quad (187)$$

Indeed this is true. Similarly, it must be the case that the conditional entropy obeys

$$H(A) \geq H(A|B) \quad (188)$$

with equality only for independent random variables A and B . To prove this, let the random variables have the probability distributions \vec{p} and \vec{q} , respectively. Using the definition, and then *Bayes's* formula (3) to rewrite the conditional probabilities, we observe that

³⁵Exercise: Prove this. Also calculate its maximum value.

$$\begin{aligned}
H(A|B) - H(A) &= - \sum_{i,j} q_j \frac{p_{i,j}}{q_j} \ln \frac{p_{i,j}}{q_j} + \sum_{i,j} p_{i,j} \ln p_i = \\
&= \sum_{i=1}^n \sum_{j=1}^m p_{i,j} \ln \frac{p_i q_j}{p_{i,j}} \leq \sum_{i=1}^n \sum_{j=1}^m p_{i,j} \left(\frac{p_i q_j}{p_{i,j}} - 1 \right) = 1 - 1 = 0 .
\end{aligned}
\tag{189}$$

We relied on the inequality

$$\ln x \leq x - 1 , \tag{190}$$

which holds for all $x > 0$.³⁶

Using the connection between joint and conditional entropy the property of *subadditivity* follows immediately,

$$H(A) + H(B) \geq H(A, B) , \tag{191}$$

again with equality if and only if the random variables are independent. Another way of saying this is that mutual information is always positive, or equal to zero for independent events.

Without giving the story away, let me say that neither the conditional entropy nor the conditional probability distribution have any immediate analogues in the quantum case for which we are aiming. We will have to revisit the entropy inequalities when we come to quantum information theory.

Relative entropy

Boltzmann's H-theorem says that entropy cannot decrease as time passes. If the passage of time is governed by stochastic maps this is not true, in general, for the Shannon entropy. The bit-flip map (11) does increase entropy, but the coarse-graining map (14) can increase it. The general statement is that a stochastic map will increase the Shannon entropy, $H(S\vec{p}) \geq H(\vec{p})$ for all \vec{p} , if and only if it has the maximally mixed probability distribution as a fixed point.

³⁶Exercise: Prove the inequality (190). Why is the entropy inequality (188) necessary for the interpretation of $H(A|B)$ to make sense?

This is one reason to introduce the *relative entropy*

$$H(\vec{p}|\vec{q}) = \sum_{i=1}^n p_i \ln \frac{p_i}{q_i} . \quad (192)$$

This does have the desirable property of monotonicity under stochastic maps, namely that

$$H(S\vec{p}|S\vec{q}) \leq H(\vec{p}|\vec{q}) \quad (193)$$

for every stochastic map S . We omit the proof, but we (or you) will prove that relative entropy enjoys other interesting properties as well.

First of all it is non-negative. More than that, it obeys

$$H(\vec{p}|\vec{q}) \geq \frac{1}{2} \sum_{i=1}^n (p_i - q_i)^2 . \quad (194)$$

To prove this, note that any smooth function f obeys

$$f(x) = f(y) + (x - y)f'(y) + \frac{1}{2}(x - y)^2 f''(\xi) , \quad \xi \in (x, y) . \quad (195)$$

Apply this to the function $x \ln x$, and you get the result by summing.³⁷ Using this result you can derive a slightly sharper version of the inequality $H(\vec{p}) \leq \ln n$, as well as—by considering the relative entropy between a joint probability distribution and the probability distribution for two independent events—a slightly sharper version of the subadditivity inequality (191).³⁸

We get an interpretation of relative entropy if we go back to the probability of obtaining an outcome m times in N trials with two outcomes, eq. (17). Taking logarithms and applying Stirling's formula we find (after some calculation) that

$$\ln \binom{N}{m} \approx -m \ln \frac{m}{N} - (N - m) \ln \frac{N - m}{N} , \quad (196)$$

$$\ln (p^m (1 - p)^{N - m}) = m \ln p + (N - m) \ln (1 - p) . \quad (197)$$

³⁷Exercise: Do it!

³⁸Exercise: Do this as well!

Putting things together, and generalizing from the binomial to the multinomial distribution valid for n outcomes, we find that the probability to obtain the frequency vector \vec{v} in N samplings from a probability distribution \vec{p} obeys

$$P(\vec{v}|\vec{p}) \approx e^{-NH(\vec{v}|\vec{p})} . \quad (198)$$

This suggests that relative entropy is a good measure of the distinguishability of two probability distributions. It is highly asymmetric in its arguments however. How different is a fair coin from a biased coin that always gives heads? Pick one of them and start flipping it to see which is which. The number of flips you have to make before you feel sure which one you picked depends very much on the choice.³⁹

Error correction

So far we have focussed on data compression. We now change perspective, and observe that removing redundancy may be a very bad idea. If a message is phrased in a natural language it is hugely redundant—and this redundancy is useful in communication, because it means that errors in the transmission are easily corrected. The misprints in these notes may be annoying, but they are not fatal. In the real world transmission of information is always distorted by noise. If you want to send a message from a space probe far out in the Solar system, or if you want to build a quantum computer, this is a key issue. The way out is to introduce redundancies in the message, allowing us to spot and correct errors.

To be quantitative about this we need a model for the noise. A very simple, but useful, model is the binary symmetric channel (11). So we are sending bits, and each bit may flip with probability p . Suppose we want to send the message 1101. The message will come through correctly with probability $(1-p)^4$, which may be unacceptably low. The obvious way to deal with this is to repeat the message three times, so that we send 110111011101. Should an error occur in one place, we can correct it by taking a majority vote. In fact some double errors can also be corrected for. But the length of the message has gone up. This is bad in itself. Moreover, with increased

³⁹Exercise: Use eq. (198) to calculate $P(\text{fair}|\text{biased})$ and $P(\text{biased}|\text{fair})$ in the example.

length comes an increased risk of double errors, and if the error probability is high the repetition code will be useless.⁴⁰

It was once believed that the only way to combat noise is to reduce it. For the binary symmetric channel, this means that one tries to reduce p . But *Shannon* took the amount of noise as given, and asked for the minimum amount of redundancy that will ensure that the message can be corrected on arrival. He proved a sharp theorem about this. As for the noiseless coding theorem the proof is non-constructive. It shows that an optimal error-correcting code must exist, but it does not provide it.

We have passed over the question of how to actually encode messages so that you come close to the limits set by the noiseless coding theorem. For this you have to consult a book on information theory. But before we come to the theorem about noisy channels we will give a simple example of an error-correcting code that improves on the simple repetition code above. One reason for doing so is that it will suggest ideas for error-correcting codes to be used in quantum computers.

We again assume that we wish to send one out of 2^4 possible messages, but we view the messages as linear combinations of four basis vectors in a four dimensional subspace of a seven dimensional vector space, where the only numbers available for the linear combinations are 0 and 1. Thus the vector space is not \mathbf{R}^7 , it is \mathbf{Z}_2^7 where \mathbf{Z}_2 denotes the integers modulo two. To be precise, we let the four vectors be the four rows of the *generator matrix*

$$G = \left[\begin{array}{cccc|ccc} 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{array} \right]. \quad (199)$$

The message we wanted to send, 1101, is then sent in the form of the sequence obtained by adding the first, second and fourth rows together, namely as 1101001. In all the 16 messages that we can send the first four entries are the message, and the last three are to be used for correcting errors. If we want to send a sequence a, b, c, d of four binary integers, what we actually send is the sequence $a, b, c, d, b + c + d, a + c + d, a + b + d$, where addition is modulo 2. Now the subspace of \mathbf{Z}_2^7 in which the code words live has been very

⁴⁰Exercise: Show that the probability that the message will come out correctly is larger than $(1 - p)^4$ only for $p < 0.151$. [There is something fishy about this!!!]

carefully chosen. After staring at the basis vectors for some time one sees that in order to transform one of the 16 possible sequences into another one has to flip the value of at least three entries. This means that any single error can be corrected for by selecting the one out of the 16 that differs from the received message by a single bit flip. Hence all single errors will be corrected for, and none of the double errors.⁴¹

But the best part of the story is the way we correct the errors. Write the generator matrix in the form $G = [\mathbf{1}|A]$. Define the 3×7 matrix $H = [-A^T|\mathbf{1}]$. By construction

$$HG^T = 0 . \quad (200)$$

This means that $H\mathbf{u} = 0$, where \mathbf{u} is any linear combination of the four rows of G , that is to say if its components form one of the 16 correct messages. If an error occurs during the transmission the message received is given by the components of a vector $\mathbf{u} + \mathbf{e}$, and

$$H(\mathbf{u} + \mathbf{e}) = H\mathbf{e} . \quad (201)$$

The 3-component vector $H\mathbf{e}$ is known as the *error syndrome*. It is easy to check that if \mathbf{e} has a single non-zero component then it can be reconstructed uniquely from a knowledge of the error syndrome.⁴² Once this is known we can correct the message by subtracting \mathbf{e} from the received message (or adding it, which is the same thing since we work with integers modulo two all the time).

We have just described the simplest example of a *Hamming code*. For messages of length 4 it works better than the repetition code, and it is evidently less wasteful. We made do with less redundancy than we might have expected. But for us the really remarkable thing about it is that we can do the error correction without, in fact, reading the message. All we need is the error syndrome. Once we start sending quantum messages we have to measure the message in order to read it. This means that reading the message changes the message. Quantum error correction has to correct messages without reading them.

⁴¹Exercise: For what values of p is the probability that the message comes out correctly larger than $(1 - p)^4$ if we use this code?

⁴²Exercise: How many single component errors are there, and how many error syndromes? Check the correspondence in detail.

Channel capacity

A channel is any medium carrying a message. Physically it may be air, an optical fibre, or a printing press. The input of the message can be regarded as a random variable A , whose outcomes are the letters being sent. The output is another random variable B , whose outcomes are the letters received. The two random variables are connected by a conditional probability distribution $P(B|A)$, which tells us the probability that b_i is received given that a_j was sent. These conditional probabilities are the matrix elements of a stochastic matrix, in this context often called *transition probabilities*. We will therefore equate a *channel* with a stochastic matrix. This means that we assume the channel to be *memoryless*. This may not be an accurate model of the physical channel, because real noise often has a tendency to come in bursts, but it is a good first approximation. The physical details of the channel are left to the engineers.

In a *noiseless* channel the output is in one-to-one correspondence with the input, which means that the stochastic matrix is a permutation matrix. A *noisy* channel offers more of a challenge. Given a stochastic matrix describing a noisy channel, we want to know the maximal rate at which information can be sent over that channel.

The quality of the channel can be measured by the mutual information $I(A : B)$ between input and output. If the output is independent of the input the mutual information vanishes. If they are perfectly correlated the mutual information equals the information $H(A)$ in the input. We define the *capacity* of a noisy channel as

$$C = \max I(A : B) , \tag{202}$$

where the maximization is with respect to the probability distribution chosen for the input. The stochastic matrix defining the channel is kept fixed. Calculating the channel capacity is a constrained optimization problem that can be addressed using the Lagrange multiplier method. We write the mutual information in the form

$$I(A : B) = H(B) - H(B|A) . \tag{203}$$

Then we set

$$\frac{\partial}{\partial p_i} \left[I(A : B) - \lambda \left(\sum_j p_j - 1 \right) \right] = \frac{\partial}{\partial p_i} \left[\sum_{k,j} p_{k|j} p_j \ln \frac{p_{k|j}}{q_k} - \lambda \left(\sum_j p_j - 1 \right) \right] = 0 . \quad (204)$$

The conditional probabilities characterize the channel and are kept constant, but when taking the derivative we must remember that $\vec{q} = \vec{q}(\vec{p})$. Provided that no component of \vec{p} vanishes we obtain

$$\sum_k p_{k|i} \ln \frac{p_{k|i}}{q_k} = \lambda + 1 . \quad (205)$$

The “1” on the right hand side comes out when you take the derivative of q_k with respect to p_i . Fortunately we know that mutual information is concave, so these equations determine the unique maximum. Unfortunately they are difficult to solve.⁴³

There is one easy case: the binary symmetric channel (11). Then there are only two equations (204), and if you write them out you see that they imply $q_0 = q_1 = 1/2$. After that you quickly arrive at⁴⁴

$$C = \ln 2 + p \ln p + (1 - p) \ln (1 - p) . \quad (206)$$

It remains to interpret this formula. For this purpose we switch back to logarithms with base 2 (turning $\ln 2$ into 1).

The interpretation is provided by *Shannon’s theorem*. It is again concerned with a message containing letters chosen according to a probability distribution \vec{p} and coded into bits. It says that *a string of bits of length $NH(\vec{p})$ can be coded into a string of bits of length N and transmitted through a channel with capacity C with arbitrarily small error, provided that*

$$H < C . \quad (207)$$

If $H > C$ this is not possible. Actually the theorem says more, and as was the case with the noiseless coding theorem some fine print should be attached. It is important to realize that the theorem applies to very long and ‘typical’ sequences only. But when it applies, it tells us exactly how much redundancy

⁴³Exercise: Show that the equations imply that $\lambda = C - 1$.

⁴⁴Exercise: Do the whole calculation! If $p = 0.01$ and you use logarithms with base 2, what is the channel capacity?

that has to be added to the message in a perfect error-correcting code. With this remarkable statement we take leave of classical information theory.

The von Neumann entropy of a quantum state

We turn to the quantum case. In our discussion of measurements we came to the conclusion that a quantum state ρ will return a wide variety of different probability distributions, depending on what measurement we choose to do. Each measurement of ρ is associated to some Shannon entropy. Is there a single one that deserves to be called ‘the’ entropy of ρ ? The first observation is that if ρ is a pure state then there is a measurement that gives zero Shannon entropy, while other measurements give positive entropy. If ρ is a mixed state the issue is less obvious. If you stare at a point in the Bloch ball, and imagine the probability distribution returned during an arbitrary von Neumann measurement (defined by a pair of antipodal points on the surface), you see that the Shannon entropy is smallest if the measurement is performed in the eigenbasis of the state.

We just made a move that is typical for quantum Shannon theory. We start with a notion from the classical theory, and then we optimize it. In this case, we optimize over all possible measurements. One observation is that methods of optimization will be important throughout the subject.

By optimizing the Shannon entropy we arrive at the *von Neumann entropy*

$$S(\rho) = -\text{Tr} \rho \ln \rho . \tag{208}$$

Calculating the logarithm, or any other reasonable function, of a positive operator poses no problem of principle. The function is defined in the eigenbasis of the operator, by replacing its eigenvalues with the function of these eigenvalues. Taking the trace that occurs in the definition of the von Neumann entropy is even easier, because we can calculate the trace in any basis we want. In particular, we can calculate the trace in a basis in which ρ is diagonal. Its eigenvalues form the probability vector $\vec{\lambda}$, and the von Neumann entropy is the Shannon entropy of that probability vector,

$$S(\rho) = - \sum_{i=0}^{d-1} \lambda_i \ln \lambda_i = H(\vec{\lambda}) . \quad (209)$$

The suggestion to call Shannon's information 'entropy' was actually made by *von Neumann*, who had already studied its quantum version.

The von Neumann entropy is indeed a very distinguished Shannon entropy, singled out by the state itself. We can choose an arbitrary POVM and obtain a probability vector having many more components than the vector $\vec{\lambda}$. Let \vec{p} be any probability distribution returned by ρ in some measurement. Then one can show that

$$H(\vec{p}) \geq H(\vec{\lambda}) . \quad (210)$$

The proof is not particularly difficult, but it relies on just a little more background than we have introduced in our 'Lengthy Introduction'. Our Bloch ball argument should make it plausible, so we leave it unproved here.

The von Neumann entropy has the obvious property that

$$S(\rho) \geq 0 , \quad (211)$$

with equality if and only if the state ρ is pure. But the next obvious property is simply not true. We recall from our discussion of entanglement that we can have a pure state ρ_{12} in a composite Hilbert space, and reduced states ρ_1 and ρ_2 describing parts of the whole, such that

$$S_{12} = 0 , \quad S_1 = S_2 \geq 0 . \quad (212)$$

We have lost the natural analogue of the obvious inequality (187). This, however, is not a weakness. It is simply the way things are.

Strong subadditivity and quantum relative entropy

We have already noticed that some seemingly natural entropy inequalities fail in the quantum case. There is, however, a master inequality from which many other inequalities follow. It is called *strong subadditivity*. It states that

$$S(\rho_{123}) + S(\rho_2) \leq S(\rho_{12}) + S(\rho_{23}) . \quad (213)$$

This is a deep result due to *Lieb* and *Ruskai*. Since it was first proved in the 1970ies there have been many attempts to find a simple proof, but these attempts have not been very successful.

The inequality can be rewritten in an interesting form if we purify the state ρ_{123} by introducing a fourth factor Hilbert space such that $\rho_{123} = \text{Tr}_4 \rho_{1234}$. By assumption the four-partite state is pure, so we can rely on

$$S_{1234} = 0 \quad S_{123} = S_4 \quad \text{and} \quad S_{12} = S_{34} . \quad (214)$$

Using this (and changing the label $4 \rightarrow 1$ at the end, to make the formula look more natural) we find that the strong subadditivity inequality becomes

$$S(\rho_1) + S(\rho_2) \leq S(\rho_{13}) + S(\rho_{23}) . \quad (215)$$

For the classical Shannon entropy the inequalities $S_1 \leq S_{13}$ and $S_2 \leq S_{23}$ hold separately. In the quantum case they do not, but their sum does.

Many important results follow from strong subadditivity. For an easy example, let the Hilbert space \mathcal{H}_2 be one-dimensional so that $S_2 = 0$. The inequality (213) then collapses too

$$S_{13} \leq S_1 + S_3 . \quad (216)$$

This is the subadditivity inequality.

Finally we define the quantum relative entropy in analogy to the classical case,

$$S(\rho||\sigma) = \text{Tr}(\rho(\ln \rho - \ln \sigma)) . \quad (217)$$

If $[\rho, \sigma] \neq 0$ the two density matrices do not share any common eigenbasis. It follows that the relative entropy is an object that is hard to manipulate, even though it is every bit as important to quantum information theory as its classical cousin is to classical information theory. One reason for this is a theorem due to *Lindblad*, who used strong subadditivity to show that *quantum relative entropy is monotone under arbitrary CP maps*.

Quantum channels

A quantum channel may be an optical fibre through which we send polarized photons. In quantum information theory it will be modeled by a CP map. And a main aim of quantum information theory is to define quantum channel capacities, analogous to that of Shannon. Indeed more than one channel capacity, because we may ask how reliably qubits can be sent through the channel, or how reliably the qubits can be used to transmit classical bits through it. Some quite unexpected phenomena can occur. Thus two quantum channels may transmit more than twice the information transmitted by a single channel, if their inputs are entangled. [But we have to interrupt the story here, and go to the next chapter. It is scandalous that I did not arrive at the Holevo bound...]

QUANTUM COMPUTATION

With very little time left, we start a discussion of what may, or may not, turn out to be the most important part of the course. Either way, it illuminates what quantum theory is about.

Computation and complexity

In the 1930ies *Turing* analysed the meaning of *computability*. He came up with the notion of a *universal Turing machine* which is supposed to be able to compute anything that can be computed by means of algorithmic procedures. From a modern perspective you can think of a universal Turing machine as any existing computer, having a finite number of internal states and a finite program but equipped with an infinite memory. Indeed the computer can be stripped down a lot, making it very slow but still universal, in the sense that given enough time it can mimic the actions of any other computer. It is generally agreed that no machine, whether classical or quantum, can do better than the Turing machine in this regard.

The universal Turing machine is used to *define* what we are supposed to mean by ‘computable’. A function is said to be computable if there exists a program for the machine which takes the argument of the function as an input, and outputs the value of the function after a finite time. Since the input and the output can be coded in binary digits, the function is a function from the integers to the integers. Any question you may care to ask can be phrased in these terms. In ASCII encoding every upper or lower case Latin letter, Arabic numeral, question mark, and so on, is given by a 7 digit binary integer. Two symbols are then given by a 14 digit binary integer, and so on. Hence your question can be assigned a number, and the answer will be a function of that number.

The story changes if we are interested in obtaining the result in reasonable time. Suppose that the task is to multiply $n \times n$ matrices together. Proceeding as usual you see that this requires n^3 multiplications. This gives us some feeling for how the complexity of the calculation grows with n . We could say that it grows like n^3 . Actually it is possible to improve the algorithm by tra-

ding some multiplications for additions, which are cheaper.⁴⁵ One can set up matrix multiplication so that the number of multiplications grows like $n^{2.37}$, and perhaps further improvements are possible. Either way the growth is polynomial in the size of the input. We count time in terms of the number of computational steps needed, and say that matrix multiplication can be done in *polynomial time*, or equivalently that it belongs to the *complexity class* **P**.

As another example, consider factoring an integer N into primes, using an algorithm that applies to any N . If we use *Erathostenes'* sieve we simply check, for every integer less than or equal to its square root, whether it divides the given integer. Measuring the size of the integer in terms of the number n of binary digits needed to write it down, the size of the calculation grows like $2^{n/2}$, that is to say exponentially in the size of the input. This leads us to define the complexity class **EXP**, consisting of problems that can be solved in a time that grows exponentially with the size of the input. Clearly **P** is a subclass of **EXP**, and we do not have a proof that factorisation does not belong to that subclass. Still the distinction between the complexity classes is fundamental. An algorithm that belongs to **P** is regarded as 'tractable'.

You can object that, as a practical matter, it is not obvious what to choose if the choice is between an algorithm with a running time $10^6 + 10^6 \cdot n$ and another with running time $e^{10^{-6} \cdot n}$. But this situation rarely occurs in practice. A deeper answer is that one can build an interesting theory based on the distinction between **P** and **EXP**. For instance, once an algorithm has been shown to be tractable in the sense that it belongs to **P**, it can be called as a subroutine in a larger program without taking the latter out of **P**.

Returning to the factoring of integers, we can use this problem to introduce a few more complexity classes. Consider first the less ambitious question whether a given integer is a prime or not. This is a yes/no question. A yes/no question is said to belong to the complexity class **BPP**, spelt out as bounded-error probabilistic polynomial, if there is an algorithm that runs in polynomial time and gives the correct answer with probability equal to $3/4$.⁴⁶ It turns out that primality testing belongs to the complexity class

⁴⁵Exercise: For the matrices A and B in eq. (81), suppose you are given the seven products $(a_{00} + a_{11})(b_{00} + b_{11})$, $(a_{01} + a_{11})b_{00}$, $a_{00}(b_{01} - b_{11})$, $a_{11}(b_{10} - b_{00})$, $(a_{00} + a_{01})b_{11}$, $(a_{10} - a_{00})(b_{00} + b_{01})$, $(a_{01} - a_{11})(b_{10} - b_{11})$. Show that no further multiplications are needed to construct AB .

⁴⁶Exercise: Change the $3/4$ to any number larger than $1/2$. How does this affect the

BPP. Algorithms for primality testing that are in actual use are of this type, but in 2002 it was discovered that there does exist a deterministic algorithm that runs in polynomial time. So primality testing belongs to **P**. The lesson is that it may be difficult to decide what complexity class a given problem belongs to.

If you want to know what the factors of a non-prime integer are, the best existing classical algorithm is known as the Number Field Sieve. Its running time grows as $2^{n^{1/3}}$, which means that it requires exponential time. On the other hand, if someone suggests a factoring, you can check in polynomial time whether this is correct. For this reason this problem is said to belong to the complexity class **NP**, for *non-deterministic polynomial*.

Of course this does not prove that factoring a prime cannot be done in polynomial time on a classical computer. In fact, one of the main open problems in theoretical computer science is whether there exist problems that belong to **NP** without also belonging to **P**. What is known is that there are many problems that are **NP complete**, in the sense that if one could prove that one of them is in **P**, then the two complexity classes coincide. This said, there is close to a consensus that $\mathbf{P} \neq \mathbf{NP}$. (And also close to a consensus that this distinction will be unaffected by quantum computation).

The BQP complexity class

Turing's definition of 'computable' hinges on functions from the integers to the integers, and you may ask if an analog computer (operating, like most of classical physics, with real numbers) can bring changes. The standard answer is that a physical analog computer would need a precision increasing exponentially with the size of the input in order to give an advantage. A physical analog computer will always be subject to noise, and this prevents the analog computer to have much impact on complexity theory. A quantum computer shares some features with analog computers, since it takes a continuum to label its states, but on the other hand the theory of measurement means that the measured output consists of a discrete set of possibilities. It is not called 'quantum' for nothing.

Deutsch raised the question of what would happen if the Turing machine

complexity class of the problem?

is allowed to operate under the laws of quantum mechanics. One of his conclusions was that a quantum Turing machine can be always be simulated by a classical Turing machine, so the theory of quantum computation does not affect the definition of computable functions. But he also concluded that for certain problems the quantum Turing machine will be faster than its classical cousin, so we have to look at the complexity classes with fresh eyes.

Of course, the point we are driving at is that in 1994 *Shor* found an algorithm for factoring integers whose running time on a quantum computer grows like n^3 , and which returns an answer that is correct with a probability greater than $3/4$. This leads to the definition of a new complexity class **BQP**, bounded-error quantum polynomial. This has now to be placed somewhere in the hierarchy of classical complexity classes. The belief is that **BQP** is larger than **P** and larger than **BPP**, but not large enough to include all of **NP**. On the other hand it may include some problems outside **NP**. Evidently, since the question whether $\mathbf{P} \neq \mathbf{NP}$ is open, this is conjectural only, but the question what a quantum computer can do is presently attracting considerable interest.

We make three overall remarks before trying to define a quantum computer. The first is that in the 25 years that have passed since *Shor's* discovery, the number of genuinely new and interesting quantum algorithms that have been discovered is very small. The second is that the obvious objection to quantum computers, that they will be prone to errors that cannot be corrected, has been quite successfully met. The objection simply does not hold, or at least it is not obvious that it holds. The third and final remark is that as far as practical applications are concerned the most promising ones seem to be *quantum simulators*, that is to say quantum computers designed to simulate physically interesting quantum systems. This possibility was first raised by *Feynman*, and it seems quite plausible that it will have consequences for, say, quantum chemistry in the not too distant future. We will, however, spend the time that is left to us on the construction of a universal quantum computer. Whether useful or not, this is a machine that sheds light on quantum mechanics.

We have a fourth remark too. There are many choices to be made in the architecture of the universal quantum computer. We will use the *circuit model*, but other options exist. We will take it to operate on qubits, but this is not necessary. Eventually we will choose one out of many possible sets of

universal gates. And so on. So the fourth remark is that whenever we make a choice, we make it without much comment.

The circuit model

A classical computer operates on bits, represented by integers a, b, \dots , counted modulo two. That is, the integers take the values 0 or 1 and $1 + 1 = 0$. Physically, this may be no voltage, or some voltage. The aim is to calculate functions of strings of zeros and ones, taking the values zero or one. The action can be broken down into elementary logical components called *gates*. Examples of gates include AND, OR, and NOT, connected by wires that are allowed to bifurcate. The AND gate accepts two inputs a and b , and returns the output ab . The OR gate also accepts two inputs, and returns the output $a + b + ab$. The NOT gate accepts a single input a and returns $a + 1$. In all cases the arithmetic is modulo two.⁴⁷ From these simple ingredients one can build a universal computer. Of course there is some physics behind, but Figure 8 is our only comment on this. As *Turing* was saying, “being digital should be of more interest than being electronic”. (Actually, soon after he said that, developments occurred that made his remark look rather questionable.)




Figure 8: Two increasingly schematic picture of the NOT gate. A transistor is a (quantum!) device that conducts currents only in the presence of the voltage supplied by the input. A resistor is included as well.

A problem with the AND and OR gates is that they are irreversible. You cannot recover their inputs from their outputs. There are thermodynamic

⁴⁷Exercise: Write the AND, OR, and NOT gates as logical truth tables. Write the NAND gate (AND followed by NOT) using both notations.

reasons to worry about this, because irreversible evolutions generate heat. *Bennett, Fredkin, and Toffoli* showed that one can construct classical computers that operate entirely with reversible gates. This was an important step towards quantum computers.

A quantum computer operates by applying a reversible unitary transformation to a multi-qubit state described in a *computational basis* formed by product vectors. A measurement in the computational basis is performed at the end, in order to read the output. The unitary transformation must be built up from a finite number of more elementary unitary transformations, again called *gates*, in a way that can be efficiently described. We need a small set of gates, but large enough so that any unitary transformation can be well approximated. Then we have a *universal quantum computer*. On paper, that is.

The problem of finding a universal set of gates is solved in three steps. First we find a small set of unitary 2×2 matrices such that any unitary 2×2 matrix can be approximated, to any given precision, as a finite product of matrices from the set. In the second step we show that any unitary transformation acting on the full Hilbert space can be written as a string of unitary matrices that act non-trivially only on one or two qubits at a time. In the third step we show that it suffices to add a single two-qubit gate.

For the first step we use the set $\{H, T\}$, where H is the Hadamard gate (69) and T is the $\pi/8$ -gate

$$T = \begin{pmatrix} 1 & 0 \\ 0 & \sigma \end{pmatrix} = e^{\frac{i\pi}{8}} \begin{pmatrix} e^{-\frac{i\pi}{8}} & 0 \\ 0 & e^{\frac{i\pi}{8}} \end{pmatrix}, \quad \sigma \equiv e^{\frac{\pi i}{4}}. \quad (218)$$

The second equality is there only to explain the name. We can now construct other unitaries such as $Z = T^4$, $X = HT^4H$, and so on. The key step in the proof of universality is to compose T and H in such a way that the resulting unitary effects a rotation of the Bloch sphere through an angle that is an irrational multiple of 2π . Once this is achieved the argument proceeds along the lines that *Euler* used to introduce his Euler angles, and you can go on to approximate any rotation to within some arbitrarily small ϵ . Two unitaries U and V are said to approximate each other to within ϵ if

$$\|U|\psi\rangle - V|\psi\rangle\|^2 < \epsilon \quad (219)$$

for every unit vector $|\psi\rangle$. We cut this complicated story short by saying

that this is possible, and that the number of discrete qubit gates you need to approximate arbitrary qubit unitaries grows like a small power of the logarithm of $1/\epsilon$.

The second and third steps are actually easier than the first, but we simply give the answer. All we need to add to our generating set is the *controlled- X* or CNOT gate acting on two qubits according to

$$|00\rangle \rightarrow |00\rangle, \quad |01\rangle \rightarrow |01\rangle, \quad |10\rangle \rightarrow |11\rangle, \quad |11\rangle \rightarrow |10\rangle. \quad (220)$$

The action of a gate on any state is defined by its action on the computational basis. In words we describe the CNOT gate by saying that you apply the X gate to the second qubit if and only if the first qubit is in the state $|1\rangle$. If we have several qubits we can apply the CNOT gate to any pair, leaving the others as they are.⁴⁸ In the course of the calculation it can happen that the qubits become entangled, in which case they do not have pure states of their own, but we can still apply our gates to them. The logic is exactly the same as when we represent an operator as a matrix.

The conclusion is that the set $\{H, T, \text{CNOT}\}$ is a *universal* set of gates, in the sense that it can be used to approximate every unitary acting on n qubits. We skipped the proof, but you will probably trust the second and third step of the argument after looking at a few examples.

When we start to combine the gates we can use at least three different ways to describe things. We can use matrices (not recommended). One alternative is to use arithmetic modulo two to describe the action on the basis states. The one qubit gates are generated by

$$H|a\rangle = (-)^a|a\rangle + |a+1\rangle, \quad T|a\rangle = \sigma^a|a\rangle, \quad \sigma = e^{\frac{i\pi}{4}}. \quad (221)$$

Normalization factors are understood. We can now calculate the action of $X = HZH = HT^4H$ in three steps.⁴⁹ We obtain

$$|a\rangle \rightarrow (-)^a|a\rangle + |a+1\rangle \rightarrow |a\rangle + (-)^{a+1}|a+1\rangle \rightarrow |a+1\rangle. \quad (222)$$

⁴⁸Exercise: Write out (220) as a matrix. Compare it with the matrices representing $X \otimes \mathbf{1}$ and $\mathbf{1} \otimes X$. Then write out all 8×8 matrices that describe the CNOT gate applied to any two out of three qubits.

⁴⁹Exercise: Do all steps explicitly, and compare with the matrix notation.


We denote the controlled- X gate by C_X . This prepares the notation for handling controlled- U gates, which apply the one-qubit unitary U to a qubit if the control qubit is in state $|1\rangle$, but leaves things alone if the control qubit is in state $|0\rangle$. Compare Problem 3 for the Dirac notation. The action of C_X on the basis states is

$$C_X|a, b\rangle = |a, b + a\rangle . \quad (223)$$


The notation becomes more cumbersome once we have to specify which, out of many, qubit serves as the control qubit and which qubit serves as the target.

Circuit diagrams

We now switch to *circuit diagrams*, in which each qubit is represented by a horizontal line interrupted by boxes to keep track of which unitary operator that is acting. The input is on the left and the output on the right. Thus the sequence XXZ is written as




The controlled- U gate is⁵⁰



If you go through a few examples you will find that the circuit diagrams are close to self-explanatory.⁵¹ The first example shows that we can switch the role of control and target in the definition of C_X :


⁵⁰Exercise: Construct a C_Z gate from our universal set. Is it trivial to construct C_U for general U ? Give a reason. (If your answer is no, rest assured that someone has done it.)

⁵¹Exercise: Verify the following three circuit diagram equalities by tracing through what happens to the computational basis states.



At first sight, this looks wrong. Since $H^2 = \mathbf{1}$ it seems that we do nothing on the first qubit to the left of the equality sign, while we certainly do something on it to the right. The point, however, is that the matrices $H \otimes H$ and C_X do not commute.

In the second example we construct the SWAP gate, whose action can alternatively be written as $|a, b\rangle \rightarrow |b, a\rangle$:



Once we have constructed the SWAP gate it can be used as subroutine in larger circuit diagrams, which saves you from writing out three CNOT gates.


Note that, like the CNOT gate, the SWAP gate can be used in a classical computer too. Moreover one can use two SWAP gates to interchange the role of control and target in a CNOT gate. The quantum computer manages this interchange using only Hadamard gates, which act on single qubits only.

Our third example is the *Toffoli gate*, which has two control qubits and one target. The action is $|a, b, c\rangle \rightarrow |a, b, c + ab\rangle$. The construction to follow makes use of the fact that, in quantum theory, NOT has a square root:

$$\sqrt{X} = HT^2H \quad \Rightarrow \quad (\sqrt{X})^2 = HT^4H = X . \quad (224)$$

The circuit diagram that defines the Toffoli gate is⁵²

⁵²Exercise: In view of our discussion of open systems, what is the reason for the final CNOT gate?



Like the gates X and C_X the Toffoli gate can be defined also in classical computer circuits acting on bits. A reversible classical computer needs the Toffoli gate in order to be universal. But then it has to be defined as a primitive, because the square root of NOT does not exist classically. In the quantum case on the other hand one and two-qubit gates are enough to ensure universality. This is very good news if you want to build a quantum computer. One qubit gates are easier to fabricate than are two qubit gates. Fabricating a three qubit gate would be very hard.

Continuing in this way, we can build a circuit that approximates any unitary acting on an arbitrary finite number of qubits. (I did not say it is easy.) But we have still to discuss the *sine qua non* of the quantum computer: preparation and measurement.

Preparation and readout

We make the convention that the qubit register is initialized in the state $|0, 0, \dots, 0\rangle$. The circuit therefore begins by creating some more interesting state to act on. We may take it that the aim of the computation is to compute an integer valued function f , taking an integer x as its argument. Let us assume that $x < 2^n$ and $f(x) < 2^m$. We then divide the register into an n qubit *input register* and an m qubit *output register*. We may need additional work qubits to act on, and if so we must take care that the calculation does not leave them entangled with the input and output qubits.

We write the integer x in binary form, so that we get a sequence of no more than n zeros and ones. By means of a unitary transformation the register is transformed into the state $|x\rangle_n |0\rangle_m$, where $|x\rangle$ is an n -qubit state encoding the input. The m -qubit state $|0\rangle$ is there to ensure that the transformation is reversible, as we will see in a moment. By hook or crook, we find a unitary transformation U_f such that

$$U_f|x\rangle_n|y\rangle_m = |x\rangle_n|y + f(x)\rangle_m . \quad (225)$$

In particular

$$U_f|x\rangle_n|0\rangle_m = |x\rangle_n|f(x)\rangle_m . \quad (226)$$

This transformation is reversible because

$$U_f|x\rangle_n|f(x)\rangle_m = |x\rangle_n|f(x) + f(x)\rangle_m = |x\rangle_n|0\rangle_m . \quad (227)$$

A similar division of the register into an input and an output register is needed for a reversible classical computer to work.

Pause to make sure what is being meant. Suppose we wish to compute $f(5) = 3$. In binary this is $f(101) = 11$. It is enough to use a three qubit input register and a two qubit output register. We need to build a unitary transformation such that

$$U_f|101\rangle|00\rangle = |101\rangle|11\rangle . \quad (228)$$

It should now be clear why $|f(x) + f(x)\rangle_m = |0\rangle_m$ even though $f(x) + f(x) \neq 0$. We do no carrying on the ket labels.

So far the discussion would apply also to a reversible classical computer. But let us apply a Hadamard gate to every input qubit,

$$H|0\rangle = |0\rangle + |1\rangle , \quad H^{\otimes 2}|00\rangle = |00\rangle + |01\rangle + |10\rangle + |11\rangle , \quad (229)$$

and so on. (As usual we ignore overall normalisation factors.) When we do go on, we obtain the remarkable formula

$$H^{\otimes n}|0\rangle_n = \sum_{x < 2^n} |x\rangle . \quad (230)$$

A single application of the unitary U_f now has the effect that

$$U_f(H^{\otimes n} \otimes \mathbf{1}^{\otimes m})|0\rangle_n|0\rangle_m = \sum_{x < 2^n} |x\rangle_n|f(x)\rangle_m . \quad (231)$$

It seems as if we have computed every value of the function in a massively parallel computation. But the information has not yet reached its destination.


In the end we have to perform a measurement. In quantum computation this is not just an afterthought. It is an essential part of the computation.

If we start by measuring the input register in the computational basis the state collapses with equal probability to anyone of the states

$$|x\rangle_n |f(x)\rangle_m . \quad (232)$$

We now know x , and a measurement on the output register yields the value of the function $f(x)$. This quantum computation therefore gives as much information as the classical one. The one difference is that the choice of x was made at random after the completion of the calculation (which is a bit odd, but hardly an advantage).

But maybe we did not ask the right question? Perhaps we could design some measurement that extracts global information about the function f , rather than some special value? Before asking this in earnest we practice a little on how to prepare interesting input states. It is easy enough to design a circuit that affects the transition $|0, 0\rangle \rightarrow |0, 0\rangle + |1, 1\rangle$, namely



This is interesting because we have created an entangled state from a separable state. This supports the claim—whose proof I skipped—that we can approximate arbitrary unitary transformations using only the gates H, T , and C_X , so that we can produce any state whatsoever from the initial state $|0\rangle_n$.⁵³

Many quantum information protocols end with a measurement in some basis other than the computational one. In the quantum teleportation protocol Alice is asked to perform a measurement using a nice error basis, that is to say (if she teleports qubits) in the Bell basis

$$|\Phi^\pm\rangle = |0, 0\rangle \pm |1, 1\rangle , \quad |\Psi^\pm\rangle = |0, 1\rangle \pm |1, 0\rangle . \quad (233)$$

A moment's thought shows that our insistence that the result of the quantum

⁵³Exercise: Design a circuit that effects the transition $|0, 0, 0\rangle \rightarrow |\text{GHZ}\rangle$, where the GHZ state is $|\text{GHZ}\rangle = |0, 0, 0\rangle + |1, 1, 1\rangle$.

computation should be read out by means of a measurement in the computational basis imposes no restriction. We can still perform a measurement in the Bell basis if we first perform a unitary transformation that turns the Bell basis into the computational basis, and then measure.⁵⁴

The Deutsch–Jozsa algorithm

The *Deutsch–Jozsa* algorithm is a proof-of-principle, showing that there exists a problem where a quantum computer provides a speed-up compared to what a classical computer can do. The problem it solves is not particularly interesting in itself, in fact it is quite contrived, but the underlying idea recurs in more interesting algorithms. Suppose we have a function from $\{0, 1\}$ taking values 0 or 1, and a circuit that computes it, that is to say that

$$U_f|0\rangle|y\rangle = |0\rangle|y + f(0)\rangle, \quad U_f|1\rangle|y\rangle = |1\rangle|y + f(1)\rangle, \quad (234)$$

where $y \in \{0, 1\}$ is arbitrary. It may be expensive to run this calculation however. The question is: How many times do we have to run the calculation before we know whether $f(0) = f(1)$? This is an example of an *oracle* problem. Applying the unitary U_f is like asking an oracle for an answer. The question is how many times we have to ask the oracle before we know the answer we want to have. In a classical computer we would have to ask twice, but the Deutch-Jozsa algorithm allows a quantum computer to answer the question with only a single application of U_f .

The trick is to make a suitable preparation before the oracle is called. Thus

$$|0\rangle|0\rangle \rightarrow (H \otimes H)(X \otimes X)|0\rangle|0\rangle = (|00\rangle - |01\rangle - |10\rangle + |11\rangle). \quad (235)$$

Now we call the oracle. Denoting $f(x) + 1 = \bar{f}(x)$ we obtain

$$|0\rangle|f(0)\rangle - |0\rangle|\bar{f}(0)\rangle - |1\rangle|f(1)\rangle + |1\rangle|\bar{f}(1)\rangle. \quad (236)$$

It is a small exercise to check this, and then to check that if $f(0) = f(1)$ the result is


⁵⁴Exercise: Design the circuit we need for this.

$$(|0\rangle - |1\rangle)(|f(0)\rangle - |\bar{f}(0)\rangle) , \quad (237)$$

while if $f(0) = \bar{f}(1)$ it is

$$(|0\rangle + |1\rangle)(|f(0)\rangle - |\bar{f}(0)\rangle) . \quad (238)$$

Finally we apply a Hadamard gate to the input register, and make a measurement of the input register. If the qubit collapses to $|1\rangle$ the function is constant, if it collapses to $|0\rangle$ the function is not constant.⁵⁵ And the oracle was called only once, as you can see from the circuit diagram.



Notice however that if we measure the output register we get no information about the actual values taken by the function. We get an outcome which with probability one half is $f(0)$ and with probability one half is $\bar{f}(0)$, but we do not know which is which. So it is a trade off. We changed the question to be about a global property of the function. And in a sense this is the point: In a quantum computer we can ask a wider range of questions about the function.

Prime factorisation and RSA cryptography

The Deutsch-Josza algorithm answers no interesting question, but it serves as an inspiration for algorithms that do. The outstanding example remains *Shor's* algorithm for finding the two prime factors of a large number N . This is indeed an interesting problem, that may well not belong to the classical complexity class **P**. Shor's algorithm does the factorisation in a time that grows like n^3 or (with improvements) like n^2 , where n is the number of binary digits of N . More precisely, the factorisation relies on a subroutine that finds the period of a certain function. Shor's algorithm provides an exponential

⁵⁵Exercise: Check it all, starting from the small exercise.

speed-up of this subroutine. Once this period is known, a classical computer finishes the factorisation in polynomial time. Perhaps one can describe Shor's achievement by saying that he found an example of the kind of questions that quantum computers are good at, sitting inside a genuinely interesting mathematical problem.

Part of the excitement is that prime factorisation plays a role in *public key cryptography*. The idea here is to rely on a mathematical problem that is easy to solve one way, and hard to solve the other way. Using factorisation of integers for this purpose leads to *RSA cryptography* (for *Rivest, Shamir, and Adleman*, who were the first outside the British Secret Service to come up with the idea). Imagine a cryptographic protocol that requires a composite number N for encryption, but the prime factorisation $N = pq$ for decryption. Starting from two large primes p and q the receiver announces N publicly, so that anyone can use it for encryption. To read the encrypted message you must either know p and q in advance, which is true only for the intended receiver of the message, or calculate them, which will take an inordinate amount of time if you rely on an algorithm whose running time grows exponentially with the number of digits in N .

To see how this comes about we need some knowledge of arithmetic. First we note that *Euclid* provided a very fast algorithm for finding the greatest common divisor (a, b) of two integers. If $a < b$ you divide b with a to obtain the remainder r , and then observe that $(a, b) = (r, a)$. This reduces the size of the problem, and by repeating the process you reach the answer in polynomial time. Given any two integers a and b this algorithm also yields two integers m and n such that

$$am + bn = (a, b) . \tag{239}$$

If $(a, b) = 1$ the pair of integers have no common factor, and then they are said to be *relatively prime*.

Before we continue our journey through elementary number theory, let me say that although all the proofs we need are simple they can be quite exhausting to follow for the greenhorn. So I recommend that you take them on trust, as being part of the classical pre- and postprocessing of the quantum algorithm. We are tackling a very serious mathematical problem, and it is only to be expected that this will involve some serious mathematics.

With this warning we come to *modular arithmetic*. Two integers are

declared to be *equal modulo* N if they differ by a multiple of N ,

$$a = b \pmod{N} \quad \Leftrightarrow \quad a = b + nN . \quad (240)$$

Addition and multiplication modulo N is defined in the obvious way. Now consider two relatively prime integers a and N . From eq. (239), with N in the role of b , it follows that there always exist an integer m such that

$$am = 1 \pmod{N} . \quad (241)$$

Hence, provided $(a, N) = 1$, a has an inverse in arithmetic modulo N . In other words, the set of integers relatively prime to N form a group under multiplication modulo N . We will be interested in the *order* of this group, that is to say in the number of its elements. This is given by *Euler's totient function* $\phi(N)$, defined as the number of integers smaller than and relatively prime to N . If p and q are distinct primes the totient function is

$$\phi(p) = p - 1 , \quad \phi(pq) = (p - 1)(q - 1) . \quad (242)$$

We need only these two cases. Confusingly we will also be interested in the *order* of an integer a , that is to say in the smallest integer r such that

$$a^r = 1 \pmod{N} . \quad (243)$$

Given an integer a relatively prime to N the set of integers of the form $a^x \pmod{N}$ forms a subgroup of order r . *Lagrange's* theorem says that the order of a subgroup always divides the order of the whole group, so r divides $\phi(N)$. In every case it must be true that

$$a^{\phi(N)} = 1 \pmod{N} . \quad (244)$$

This is called *Euler's theorem*. It must hold, otherwise the number of elements in the group would be larger than its order.⁵⁶

In our examples N will be huge. It is nevertheless a quick affair to let a computer calculate a^x modulo N . The number of multiplications needed is kept modest if we do the calculation by repeated squaring,

⁵⁶Exercise: For $N = 4, 5, 6, 7, 9, 10, 15$, find the multiplicative inverse modulo N of every integer that has one, and count the latter. Verify that $\phi(pq) = (p - 1)(q - 1)$ when p and q are distinct primes. Find examples of integers having an order lower than $\phi(N)$.

$$a \rightarrow a^2 \rightarrow a^{2^2} \rightarrow a^{2^3} \rightarrow \dots . \quad (245)$$

Then a^x is created by multiplying a subset of these powers together. In this way a^x can be calculated in polynomial time on a classical computer.

Incidentally, this gives us a hint why primality testing is easier than factorisation. If there is an integer $a < N$ such that $a^{N-1} \neq 1$ modulo N then N cannot be a prime number. Unfortunately there are composite numbers that pass this test ($361 = 3 \cdot 11 \cdot 17$ is the smallest example), which is why the full story is much longer.

We now have all the number theory we need for RSA cryptography. The protocol starts when the receiver picks two large primes p and q , and another integer c relatively prime to $(p-1)(q-1)$. She also calculates the inverse of c in arithmetic modulo $(p-1)(q-1)$. This is an integer d such that

$$cd = 1 \pmod{(p-1)(q-1)} . \quad (246)$$

The product $N = pq$ and the integer c are made public, but p , q , and d are kept secret. Anyone wanting to send a message to the receiver converts the message to an integer a (using ASCII encoding, say), and checks that a is relatively prime to N . Then he calculates the encrypted message

$$b = a^c \pmod{N} . \quad (247)$$

The message is easily decrypted by anyone knowing the secret integer d , because

$$b^d = a^{cd} = a^{1+n(p-1)(q-1)} = a \pmod{N} . \quad (248)$$

All the calculations are done in polynomial time.

What can the eavesdropper do? To compute d she needs to factor the integer N . Mathematicians have worked on this problem since the days of *Eratosthenes*, and so far the best algorithm they have come up with runs in exponential time. If N is large enough this suggests that the message will be safe from eavesdroppers for many years to come. To learn what large enough means in practice, we note that in 2016 the NSA recommended keys with at least 3072 bits. Keep in mind that the NSA has a history of recommending cryptos that they themselves can just break.

Actually the eavesdropper can get by with somewhat less. She only needs to find the order of b , that is to say an integer r such that

$$b^r = 1 \pmod{N} . \quad (249)$$

The order divides $\phi(N) = (p-1)(q-1)$, but may not be equal to it. Now the publicly known integer c is relatively prime to the unknown integer $\phi(N)$. Since r is one of the factors of the latter it follows that $(c, r) = 1$. It then follows that there exists an integer d' such that

$$cd' = 1 \pmod{r} . \quad (250)$$

The next point to notice is that when two integers a and b are related the way they are in the RSA protocol then they have the same period. This is because we are assured that integers c and d exist such that $b = a^c$ and $a = b^d$. This means that the integers a and b belong to the same cyclic subgroup of the multiplicative group of integers relatively prime to N , and their respective orders are both equal to the order of that subgroup. Hence $a^r = 1 \pmod{N}$. Equipped with all these assurances, and assuming that she can find the period r in the time at her disposal, the eavesdropper performs the calculation

$$b^{d'} = a^{cd'} = a^{1+nr} = a \pmod{N} . \quad (251)$$

She can now read the message at her leisure.

What is the size of the calculation Eve must do to find the period r ? It is convenient to rephrase the problem slightly by defining the function $f_a(x) = a^x$. Then the element a is of order r if and only if the function $f_a(x)$ is a periodic function of period r in arithmetic modulo N ,

$$f_a(x) = a^x \quad \Rightarrow \quad f_a(x+r) = a^{x+r} = a^x = f_a(x) . \quad (252)$$

Now the question is how hard it is to find the period of this function. Classically, the way to find the period is to evaluate $f_a(x)$ for many values of x , until one finds a pair x_1 and x_2 such that $f_a(x_1) = f_a(x_2)$. The period will then be a divisor of $x_1 - x_2$, and by finding a few such coincidences it is likely that the period is the largest common divisor of them all. But since the difference $f_a(x_1) - f_a(x_2)$ can take N different values we need to examine about N different pairs to find a coincidence. The number of different pairs is roughly

the square of the number of evaluations of f_a that we perform, so it follows that we are likely to need \sqrt{N} evaluations of the function before we find a single coincidence. If $N \approx 2^n$ this means that the procedure is exponential in n . And this is provably the best a classical algorithm can do with the period finding problem. But a quantum computer offers an exponential speed-up. A constant number of evaluations, followed by some computational steps that take polynomial time only, suffice.

Before we turn to this, we should perhaps finish the factorisation question. Since we have no guarantee that the period r equals $\phi(pq)$, we are not yet able to solve for p and q . However, with some further number theoretical manipulations that we do not go into, it can be shown that if we pick the integer a at random then knowing the period of $f(x) = a^x$ modulo N will enable us to factor $N = pq$ with a probability no smaller than $1/2$. Hence, if the period finding problem belongs to the complexity class **BQP**, so does the factoring problem. I will focus exclusively on period finding from now on.

Shor's algorithm

I will be a little sketchy about Shor's period finding algorithm. The hope is that we can find a way of asking for the period of the function without actually having to ask for the values it takes. In the first step we realize the transformation

$$|x\rangle_n |0\rangle_{n'} \rightarrow \sum_{x=0}^{2^n-1} |x\rangle_n |f(x)\rangle_{n'} , \quad (253)$$

where $f(x) = a^x$ and n' is the smallest integer such that $2^{n'} > N$. This is not an oracle problem. We work with a concrete function. If the number of binary digits in N is 3072, or thereabouts, we need many qubits for the purpose. However, on paper, ingenuity supported by repeated squaring, as in eq. (245), enables us to do this in an efficient manner. I skip the details of this interesting step.

To reduce clutter we now perform a measurement on the output register, and obtain a value $f(x_0)$ for the function. This means that the n qubit input register collapses to a superposition of all the values of x that return this

value of f , namely

$$|\psi\rangle = |x\rangle_n \rightarrow \sum_{j=0}^{m-1} |x_0 + jr\rangle_n . \quad (254)$$

Here m is the smallest integer such that $x_0 + mr \geq 2^n$. The unknown period r is in there, but it is not so easy to get it out. A measurement on the input register would collapse the state to one of the m states

$$|x_0 + jr\rangle . \quad (255)$$

This is not helpful, because the value of x_0 is not known, and it would change if we run the procedure again. So we change the question.

We need a unitary transformation that moves the unwanted integer x_0 into an overall phase factor that does not affect the probability vector. The *discrete Fourier transformation* in dimension 2^n is defined by its action on the basis states, as

$$F|x\rangle_n = \frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} \omega^{xk} |k\rangle_n , \quad \omega = e^{\frac{2\pi i}{2^n}} . \quad (256)$$

Note carefully that in the exponent of ω the integers x and k are treated as ordinary integers. The product is nevertheless taken modulo 2^n because ω is a 2^n th root of unity. (And we could have made the definition in an arbitrary dimension. We choose it to be 2^n because this is what we want.)

When we apply F to the input register we obtain

$$F|\psi\rangle = \frac{1}{\sqrt{2^{nm}}} \sum_{k=0}^{2^n-1} \sum_{j=0}^{m-1} \omega^{(x_0+jr)k} |k\rangle_m = \frac{1}{\sqrt{2^{nm}}} \sum_{k=0}^{2^n-1} \omega^{x_0k} \sum_{j=0}^{m-1} \omega^{jkr} |k\rangle_n . \quad (257)$$

Now we perform the measurement on the input register. The probability of obtaining the k th outcome is

$$p_k = |\langle k|F|\psi\rangle|^2 = \frac{1}{2^{nm}} \left| \sum_{j=0}^{m-1} \omega^{jkr} \right|^2 . \quad (258)$$

This depends on the period r while x_0 has disappeared, just as we wanted.

There are two questions left to address. How do we physically implement the Fourier transformation, and how do we extract r from the probabilities? We will see the quantum speed-up when we address the first question, but first we deal—too sketchily—with the second. It ends with an interesting moral.

The easy case of the second question is if the unknown period r divides 2^n . Recalling that the period also divides $(p-1)(q-1)$, we see that this implies that the prime factors of N are of the form $2^x + 1$ for some integer x . (This includes the *Fermat primes* $2^{2^x} + 1$, of which the only known examples are 3, 5, 17, 257, 65537. Reports that $15 = 3 \cdot 5$ has been factored in the lab deal with the easy case.) Anyway, because $\omega^{2^n} = 1$ we see that the amplitudes in (258) will add up if kr is a multiple of 2^n , while destructive interference will occur otherwise. Thus, when the measurement results in the k th outcome it tells us the value of the integer $2^n/r$ up to an unknown multiple. If we run the whole procedure again, we will get the value of $2^n/r$ multiplied with another unknown multiple. Running the whole procedure a few times and taking the greatest common divisor of the received k -values determines $2^n/r$, and hence r . We are finally done, and if the Fourier transformation can be done in polynomial time we are done in polynomial time.

Although this was the easy case, it may be a little difficult to follow, so let us do an example. Set $N = 15$ and choose $a = 7$, which is relatively prime to 15. In arithmetic modulo 15 we find

$$a = 7, \quad a^2 = 4, \quad a^3 = 13, \quad a^4 = 1. \quad (259)$$

So the period $r = 4$. (If you want to do this calculation in your head, think $a^4 = a^3 \cdot a = 13 \cdot 7 = -2 \cdot 7 = -14 = 1$ modulo 15.) Even if we do not know the value of r , the machine does, in some sense. When we measure on the output register we might obtain the value $f_7(x) = 4$, in which case the input register is in the four qubit state

$$|\psi\rangle = |2\rangle + |6\rangle + |10\rangle + |14\rangle. \quad (260)$$

So $m = 4$. The probability to obtain the k th out of the 16 outcomes is

$$p_k = \frac{1}{16} \frac{1}{4} \left| \sum_{j=0}^3 \omega_{16}^{4jk} \right|^2 = \frac{1}{16} \frac{1}{4} \left| \sum_{j=0}^3 \omega_4^{jk} \right|^2 = \begin{cases} \frac{1}{4} & \text{if } k = 0, 4, 8, 12 \\ 0 & \text{otherwise.} \end{cases} \quad (261)$$

Here we relied on $\omega_{16}^4 = \omega_4$, as well as on the key formulas

$$\omega = e^{\frac{2\pi i}{N}} \quad \Rightarrow \quad \sum_{j=0}^{N-1} \omega^j = 0 \quad \Rightarrow \quad \sum_{j=0}^{N-1} \omega^{jk} = N\delta_{jk} . \quad (262)$$

Say that we obtain the outcome 8. We repeat the procedure, and obtain a different value, say 12. We compute the greatest common divisor $(8, 12) = 4$. Then we have determined the integer $16/r = 4$, so that we can calculate $r = 4$. Whatever the outcomes are, and even if N is large, the mathematics of the last step guarantees that with high probability we need only a few repetitions to determine r .

If the prime factors are not of the form $2^x + 1$ the story is more involved, and hinges on *continued fractions*. This is a chapter of number theory that allows us to approximate arbitrary numbers with rational numbers having the smallest possible denominators, given the degree of approximation we want. An interesting point emerges, which is that the analysis succeeds in identifying r precisely because the set of possible measurement outcomes is discrete, or if you like because in some ways the quantum computer is more digital than analog. Unfortunately, here we must cut the story short by saying that the conclusion is the same as in the easy case: if the Fourier transformation can be done in polynomial time then the probability that we will be done in polynomial time can be made as high as we please.

The Fourier transform: Fast and faster

We have arrived at the heart of the matter, where the quantum speed-up happens. Acting on the components of a vector, rather than on the basis vectors, the discrete Fourier transform is

$$\hat{\mathbf{f}} = F_N \mathbf{f} \quad \Leftrightarrow \quad \hat{f}_j = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} \omega_N^{jk} f_k . \quad (263)$$

It arises whenever a continuous signal has been sampled at discrete points. A first look suggests that it requires N^2 multiplication. However, classically, the Fast Fourier Transform achieves the same goal with only $N \log N$ multiplications, and the even faster quantum Fourier transform builds on that idea. It is convenient to assume that $N = 2^n$. (If not, pad the vector to be

transformed with zeros until the dimension reaches 2^n .) The $N = 2$ Fourier matrix is already familiar to us,

$$F_2 = H . \quad (264)$$

Moving on to $N = 2^2$ we decide to label the vector using binary digits. Then we see that

$$F_4 \mathbf{f} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & \omega_4 & \omega_4^2 & \omega_4^3 \\ 1 & \omega_4^2 & 1 & \omega_4^2 \\ 1 & \omega_4^3 & \omega_4^2 & \omega_4 \end{pmatrix} \begin{pmatrix} f_{00} \\ f_{01} \\ f_{10} \\ f_{11} \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & \omega_4^2 & \omega_4 & \omega_4^3 \\ 1 & 1 & \omega_4^2 & \omega_4^2 \\ 1 & \omega_4^2 & \omega_4^3 & \omega_4 \end{pmatrix} \begin{pmatrix} f_{00} \\ f_{10} \\ f_{01} \\ f_{11} \end{pmatrix} . \quad (265)$$

Remembering that $\omega_4^2 = i^2 = -1 = \omega_2$, we see that provided we reorder the components of the vector before applying the matrix, we can replace the matrix F_4 with

$$F_4 \rightarrow \begin{pmatrix} F_2 & D_2 F_2 \\ F_2 & -D_2 F_2 \end{pmatrix} , \quad \text{where} \quad D_2 = \begin{pmatrix} 1 & 0 \\ 0 & \omega_4 \end{pmatrix} . \quad (266)$$

We achieve this by placing the even columns of the matrix before the odd ones. In the vector, the even components f_{a0} are placed before the odd components f_{a1} .

The same trick works when we double the dimension. Before applying F_8 , we place the even vector components f_{ab0} above the odd components f_{ab1} , and we place the even columns of the matrix before the odd ones. This has the effect that

$$F_8 \rightarrow \begin{pmatrix} F_4 & D_4 F_4 \\ F_4 & -D_4 F_4 \end{pmatrix} , \quad \text{where} \quad D_4 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \omega_8 & 0 & 0 \\ 0 & 0 & \omega_8^2 & 0 \\ 0 & 0 & 0 & \omega_8^3 \end{pmatrix} . \quad (267)$$

(Yes, this becomes easier to see if you write it out explicitly, but typographically it takes a lot of space.) We now work recursively. We replace F_4 with

the expression we already have. We can do this if we perform the appropriate reordering within the two sets of four components of the vector. So, reordering in two steps,

$$\begin{pmatrix} f_{000} \\ f_{001} \\ f_{010} \\ f_{011} \\ f_{100} \\ f_{101} \\ f_{110} \\ f_{111} \end{pmatrix} \rightarrow \begin{pmatrix} f_{000} \\ f_{010} \\ f_{100} \\ f_{110} \\ f_{001} \\ f_{011} \\ f_{101} \\ f_{111} \end{pmatrix} \rightarrow \begin{pmatrix} f_{000} \\ f_{100} \\ f_{010} \\ f_{110} \\ f_{001} \\ f_{101} \\ f_{011} \\ f_{111} \end{pmatrix}. \quad (268)$$

The net effect is that we have ordered the components in *bit reversed* order.

We can start the recursion from a Fourier matrix of arbitrary size. Moving all the even numbered columns to the left we replace

$$F_N \rightarrow \begin{pmatrix} F_{\frac{N}{2}} & D_N F_{\frac{N}{2}} \\ F_{\frac{N}{2}} & -D_N F_{\frac{N}{2}} \end{pmatrix}, \quad (269)$$

where the definition of the diagonal matrix D_N should be clear. The fact that this works is known as the *Danielson–Lanczos* lemma, and leads to the following recipe for performing the Fast Fourier Transform: Start with a vector having 2^n components and write its components in bit reversed order. Apply F_2 to all pairs of components. In the second step we apply the matrix

$$\begin{pmatrix} \mathbf{1} & D_2 \\ \mathbf{1} & -D_2 \end{pmatrix} \quad (270)$$


to all quartets of components. In the third step we apply

$$\begin{pmatrix} \mathbf{1} & D_{2^2} \\ \mathbf{1} & -D_{2^2} \end{pmatrix} \quad (271)$$

to each set of 2^3 components. And so on. In each step we perform about N multiplications, and we are done after $n = \log N$ steps. That is to say that we have brought the number of multiplications down, from N^2 to $N \log N$. The preliminary bit reversal is computationally cheap, so this is a remarkable achievement. But it remains exponential in n .

We want more. We want the calculation to be polynomial in n . At the same time we are willing to settle for less, because we do not need to calculate the individual components of the Fourier transformed vector. All we need is to calculate $F|\psi\rangle$ in eq. (257). From this we can extract at most one component by means of a measurement.

Let us construct a circuit that does this. I first give the answer for a four qubit Hilbert space,



It should be clear from this what the answer is for any n . In particular you can write down the circuits for $n = 2$ and $n = 3$ and play with them, until you see why they work as advertised.

The circuit begins with a computationally cheap permutation effecting a qubit reversal of the input,

$$|x_3x_2x_1x_0\rangle \rightarrow |x_0x_1x_2x_3\rangle . \quad (272)$$

Then follows a number of Hadamard gates, and a number of controlled unitaries C_{R_k} , where the one qubit phase gates R_k are defined as

$$R_k = \begin{pmatrix} 1 & 0 \\ 0 & \sigma_k \end{pmatrix} , \quad \sigma_k = e^{\frac{i\pi}{2^k}} . \quad (273)$$

In the circuit the chosen integer k depends on the ‘distance’ to the control qubit. This is all.

It is time to count the number of gates that appear in the circuit for arbitrary n . Clearly there are n Hadamard gates, while the number of phase gates is

$$1 + 2 + \dots + (n - 1) = \frac{n(n - 1)}{2} . \quad (274)$$

Hence the number of gates needed to effect the quantum Fourier transform grows like n^2 , a remarkable improvement compared to the classical FFT which needs $2^n n$ operations. It is true that we do not get as much information out of the quantum Fourier transform as we get out of the FFT, but we do get out all the information we need for Shor's algorithm to work.

There is a critical question to ask. When n is large we will need to implement phase gates that rotate their qubits through very small angles. This suggests that the quantum computer is subject to the same objection that one can raise against classical analog computers: the precision needed will grow with the complexity of the calculation in a way that nullifies any advantages it may have. This question was analysed by *Coppersmith*. Unfortunately our discussion of the final steps of the algorithm was too sketchy for use to give the details, but the conclusion is that the success probability of the algorithm is quite insensitive to small errors in the phase gates. Indeed, one can fix a sufficiently large integer k_0 and decide not to implement any R_k with $k > k_0$. So the surprising answer to the critical question is that not only is there no problem, we can in effect do the calculation with a number of gates that grows like n rather than like n^2 . Ultimately it is the discreteness of the output from a quantum computer that saves the day.

Other algorithms

The ideas behind *Shor's* algorithm can be adapted and generalized to speed up the solution of a range of very interesting problems, such as finding principal ideals and unit groups in algebraic number theory. It can also be used to break several but not all public key cryptosystems. However, it should be remembered that the rapid development of classical computers happened because it was possible to make money out of each improvement. It is difficult to see anyone making money out of principal ideals. A quite different algorithm, due to *Grover*, gives a quadratic speed-up in a search problem, and there are also algorithms to speed up tasks like simulated annealing. At the moment it seems that the most likely practical applications lie in the direction of letting one quantum system simulate another, as originally suggested by *Feynman*.

Quantum error correction

A classical computer can easily perform Avogadro's number of operations without committing a single bit error. Its electronic switches are large on the scale of the thermal fluctuations, oscillations are quickly damped out, and whatever happens the bits remain unentangled with the environment. In a quantum computer error correction is essential. For several reasons it used to be thought that quantum error correction is impossible, but *Shor* and *Steane* surprised the world by showing that it is not.

The naive idea for error correction is to make sure that every qubit comes in triplicate. This founders on the fact that we cannot copy an unknown quantum state. A first attempt around this problem might be to use two ancillas and store the state $z_0|0\rangle + z_1|1\rangle$ as a three-qubit state,

$$|0\rangle|0\rangle(z_0|0\rangle + z_1|1\rangle) \rightarrow z_0|000\rangle + z_1|111\rangle . \quad (275)$$

This violates no rules, and is indeed easily achieved by means of two CNOT gates. We assume that at most one qubit at a time is affected by noise, so an error to be corrected for could be the 'bit-flip'

$$z_0|000\rangle + z_1|111\rangle \rightarrow z_0|010\rangle + z_1|101\rangle . \quad (276)$$


Now we run into the second problem. We cannot inspect a quantum state without changing it in an irreversible way. If we measure any of the qubits in the computational basis, the state collapses to either $|010\rangle$ or to $|101\rangle$. Then we know that an error occurred, but we cannot restore the original three-qubit state because all information about the amplitudes z_0 and z_1 has been lost. Indeed the problem is worse, because there may be a 'phase-flip' error

$$z_0|000\rangle + z_1|111\rangle \rightarrow z_0|000\rangle - z_1|111\rangle . \quad (277)$$

We would not even notice this if we measure in the computational basis. Still worse, in many ways the quantum computer behaves like an analog computer, so the errors are not discrete. There may be a small amplitude for an error, which then builds up during the computation.

Let us focus on bit flip errors to begin with. We are going to correct for single bit flips affecting one qubit only, without looking at the

state to be transmitted. The key observation is that the uncorrupted state $z_0|000\rangle + z_1|111\rangle$ belongs to a two dimensional *code subspace*. A single error transforms the state so that it sits in one out of three mutually orthogonal two dimensional subspaces, all of them orthogonal to the code subspace. This desirable situation occurs because the state to be transmitted belongs to a 2^3 dimensional Hilbert space. To do the correction we introduce two additional ancilla qubits set to $|0\rangle$ initially, and on which measurements will be performed. We will apply gates to the interesting qubits conditional on the outcome of those measurements. The circuit diagram is



In the final round we apply gates such as $X^{x\bar{y}}$, where as usual $\bar{y} = y + 1$ in binary arithmetic and y is the outcome of the measurement on one of the two ancillary qubits. You can easily check that the output is the desired—and still unknown—state $z_0|000\rangle + z_1|111\rangle$, also if the input is a state in which one of the qubits has been corrupted by a bitflip, such as $z_0|001\rangle + z_1|110\rangle$.

Of course this is only a partial success, because we still have to deal with phase flips, and the various kinds of continuous drifts that can occur. In fact the case we dealt with is a rather harmless one, in which the error is some unwanted unitary transformation of a qubit. The three qubit state stays pure. But in a system that is open to an environment the state can evolve in many ways that do not preserve this property.

It helps to look at this a little more abstractly. Suppose that the state we want to protect is encoded in an n -qubit state $|\Psi\rangle$, and that the relevant rest of the world starts out in some reference state $|0\rangle$. The state is corrupted by some unitary transformation acting on the whole Hilbert space. We can ‘discretize’ this by means of a unitary operator basis $\{U_I\}$, so that the state transforms according to

$$|\Psi\rangle|0\rangle_R \rightarrow \sum_I U_I |\Psi\rangle |\psi_I\rangle_R . \quad (278)$$

We can do this because we assume nothing whatsoever about the reservoir states $|\psi_I\rangle$. They are neither orthogonal nor normalized. Thus the ‘discretization’ is kind of a fake at this point. Nevertheless the unitary operator basis will soon earn its alternative name “error basis”.

To proceed we need to assume something about the noise, and something about the state $|\Psi\rangle$ that we are trying to protect. We need to ensure that

$$\text{Tr}|\Psi\rangle\langle\Psi|U_I^\dagger U_J = \delta_{IJ} . \quad (279)$$

If this can be arranged the error basis gives rise to a set of mutually exclusive alternatives, and a measurement can be devised so that the state collapses according to

$$\sum_I U_I |\Psi\rangle |\psi_I\rangle_R \rightarrow U_I |\Psi\rangle |\psi_I\rangle_R . \quad (280)$$

Once we know the outcome we can apply the appropriate operator U_I^\dagger to the state. The error is corrected, and never mind the reservoir state.

Concerning the noise, we assume that it acts on each qubit separately, and moreover that it does not affect more than w qubits. Thus we assume that

$$U_I = E_I^{(1)} \otimes E_I^{(2)} \otimes \dots \otimes E_I^{(n)} , \quad (281)$$

where at most w of the unitary operators $E_I^{(i)}$ differ from the identity. Most, but not all, people who worked on this agree that this is a physically reasonable assumption.

Now we look at (279) with new eyes. The last equality would hold if $|\Psi\rangle\langle\Psi|$ were the maximally mixed state. But now we have arranged that, in any term of the trace, at most $2w$ of the factors contain non-trivial error operators. Therefore it is enough if the partial trace over the remaining $n-2w$ factors (those that have not been affected by the noise) is a maximally mixed state.

We now specialize to the case $n = 5$ and $w = 1$, that is we aim to correct single qubit errors only. Consider the following interesting state, in which every computational basis state with an even number of 1s has been included,

$$\begin{aligned}
|v_0\rangle = & |00000\rangle + |11000\rangle + |01100\rangle + |00110\rangle + |00011\rangle + |10001\rangle - \\
& -|10010\rangle - |10100\rangle - |01001\rangle - |01010\rangle - |00101\rangle - \\
& -|11110\rangle - |11101\rangle - |11011\rangle - |10100\rangle - |01111\rangle .
\end{aligned} \tag{282}$$

It clearly has a lot of structure. This structure is actually coming from a discrete Heisenberg group, but we pass lightly over that. There is a corresponding orthogonal state $|v_1\rangle$, obtained by switching the 0s and the 1s in $|v_0\rangle$. Together they span a two dimensional code subspace, consisting of the states

$$|\Psi\rangle = z_0|v_0\rangle + z_1|v_1\rangle . \tag{283}$$

You can now check that if we take the partial trace over any three out of the five factors, there holds

$$\text{Tr}_{123}|v_0\rangle\langle v_0| = \text{Tr}_{123}|v_1\rangle\langle v_1| = \mathbf{1}_{45} , \text{Tr}_{123}|v_0\rangle\langle v_1| = \text{Tr}_{123}|v_1\rangle\langle v_0| = 0 . \tag{284}$$

I am being cavalier about normalisation factors here, and I singled out the first three factors as an example only. The calculation is most conveniently done by first calculating the eight scalar products $\langle abc|v_0\rangle$, where a, b, c are integers modulo 2. The scalar products $\langle abc|v_1\rangle$ are obtained by switching 0 and 1, and then you can do the sum that defines the trace. Given this result it follows that, for any state in the code subspace and regardless of which three factors you trace out,

$$\text{Tr}_{123}|\Psi\rangle\langle\Psi| = |z_0|^2\text{Tr}_{123}|v_0\rangle\langle v_0| + |z_1|^2\text{Tr}_{123}|v_1\rangle\langle v_1| = \mathbf{1}_{45} . \tag{285}$$

Given that at most one error occurs, this is all that is needed for eq. (279) to hold. We can code a qubit as a state in this code subspace, and we are assured that any single qubit error can be detected and corrected.

This is a promising start for the subject of quantum error correction. It is also where these notes end, since the ambition was to keep them to 100 pages. If they have no other merit, at least they are a success in this regard.