The Age of Entanglement

Quantum Computing the (Formerly) Uncomputable

"... because nature isn't classical, dammit ..."

Richard Feynman, 1981

Quantum mechanics is a venerable field of study. The year 2000 marked the 100th anniversary of the original quantum hypothesis proposed by Max Planck in November of 1900. Few current fields in physics or engineering are as old as quantum mechanics. It predates relativity, both special and general. It predates nuclear and particle physics. Quantum mechanics even predates universal acceptance of the molecular hypothesis, that is, that all matter is made up of individual molecules in thermal motion. It may be hard to believe, but this happened only after Einstein's paper on Brownian motion was published in his miracle year 1905.

Quantum mechanics was a topic of study long before the beginnings of modern solid state physics, and indeed quantum theory formed the basis of the modern theory of solids. All of modern electronics, with semiconductor chips and computers, is a younger field of study than quantum mechanics. At the time of Planck's announcement, no-one knew that the Milky Way was a galaxy of stars. Nor did anyone realize that the so-called nebulae were other galaxies in a vast Universe of galaxies. The discovery of the Big Bang by Edward Hubble happened decades after Planck's announcement in Berlin, and evidence for the Big Bang came only several decades after that. All-in-all, most of the

fields of physics that hold our attention today are upstarts in comparison to quantum mechanics.

If quantum mechanics is so old and mature, why is it the focus of so much attention? When we read late-breaking science news in popular magazines, quantum devices are in the headlines, including the recent fervor about quantum communication and computation. For such an old field of study, one for which *all* theoretical aspects have long ago been verified experimentally, why does it hold onto popular imagination so strongly?

The answer is not that quantum mechanics is unintuitive. Ask any college freshman who is taking introductory physics whether they think physics is intuitive, and they will fill your ears with exasperation and frustration at the seemingly unintuitive subject. Even seasoned physicists are often stumped and surprised by ordinary physics. Systems as well-understood as electromagnetics or old-fashioned mechanics can be unintuitive. Spinning tops or collections of magnets can raise hour-long debates among highly educated and savvy physicists.

In the Physics Department at Purdue University, I can usually be found in the morning between 9:30 and 10:00 attending Professor Ramdas' Coffee Club in the Solid State Library down the hall from my office. This is a loose group of nuclear, high-energy, and solid-state physicists. We include both theorists and experimentalists among our members. One of the coffee club members, Marty Becker, is an Emeritus who travels around Indiana giving physics shows to school children. He is always coming to coffee with bars and rods, balls and magnets, pails of water, or whatever, all part of demonstrations he is developing for his show. Without fail, his demonstrations (all of them fundamental in nature and certainly classical (raise energetic arguments among the attendees. It is not unusual for there to be as many conflicting explanations of the phenomenon as there are people in the room. And we are a pretty erudite group! Needless to say, physics, even classical physics, is largely unintuitive even to physicists.

Quantum mechanics, in this sense, is no less intuitive than classical mechanics. This is not the reason why quantum theory so grabs our attention.

The reason why quantum mechanics holds a place apart from the rest of modern physics is that it is implausible. The physical behavior of extremely light-weight particles, like electrons and protons, defies Aristotelian logic. The logical problems of quantum mechanics are not even that deep. They run into trouble right at the beginning of Philosophy 101 with an apparently obvious tautology: an electron is *either* a particle, *or* it is not a particle. This sentence is clearly true. But in quantum mechanics I can also make the following true statement: an electron *is* a particle, *and* it is *not* a particle. This sentence is a contradiction in classical logic (violating the proposition {⊦.~(p.~p)} in the notation of Russell and Whitehead), but it strikes at the fundamental core of quantum behavior.

For indeed, an electron is a particle, and it is not a particle. It is a wave, and it is not a wave. It is found precisely where you observe it to be, yet it is nowhere before you observe it. You can know its momentum to infinite accuracy, yet only if it can be found anywhere in the universe. You can equally pin it down precisely inside an atom, yet its momentum can take almost any value at all.... Every statement seems to be either a contradiction, or a restriction. This is quantum mechanics! It is not the fact that quantum mechanics is unintuitive that gives it its allure, but that it lives in states that cannot logically exist.

But they do exist. It is an unassailable fact that every prediction of quantum theory has been experimentally verified, to date. There are no stones unturned. Every system, from solids to liquids to gases to plasmas to high-energy particles, bear out every aspect of quantum theory. Quantum mechanics is one of the most thoroughly tested theories in physics, and it has passed every test flawlessly. Therefore, its logical implausibility, though a nuisance to philosophers, causes no trouble for the practicing

physicist. We take the laws of quantum theory, derive their consequences, and look for those consequences in the laboratory.

Some of those consequences are dazzling (and potentially useful). With confidence borne of success, we say without fear of contradiction that quantum measurements performed on particles on our side of the Universe instantaneously affect the outcome of experiments performed on particles on the other side of the Universe. Unbelievable! We further assert that a quantum computer can simultaneously compute the answer to a million questions all at the same time, by performing only a single computation. Audacious! These are bold and implausible assertions, and I demonstrate their validity in this chapter.

INTERFERING PHOTONS

It is always best to start with those things with which we are most familiar. Therefore, before describing the quantum behavior of light, I begin with the interference of coherent light which has been discussed in several applications in the preceding chapters. We will see that much of what we understand about classical light (light made up of electromagnetic waves) can be used with only slight modification when we begin to talk about the quantum of light ⁄ the photon. For instance, we saw that interference of light inside nonlinear interferometers allows light to control light in the all-optical internet. And interference inside holographic crystals is the origin of imaging computers that use the full parallelism of visual images. Interference in these examples occurs because waves satisfy the *principle of linear superposition*, which states that any wave can be described as a sum of individual waves. For light, the electric fields of individual waves add together to produce a resultant wave that experiences constructive or

destructive interference. The addition and interference of light waves is our point of departure as we begin our discussion of quantum optical machines.

Long before Dennis Gabor thought of holography, the ingenious Thomas Young devised a simple and elegant experiment that demonstrated the wave nature of light. Young (1779-1823) was an English physician and physicist who, in his spare time, helped decipher the Rosetta Stone and demonstrate once and for all that Egyptian Hieroglyphics were phonographic in nature, shattering the romantic notion that the magic symbols could be an instance of Leibniz's universal "character". As a physician, his principal interest was in the physics of visual perception. He was the first to measure the change in curvature of the eye as it focused at different distances, and he discovered the cause of astigmatism. The three-color theory of color perception (that only the colors of red, green and blue are needed to perceive all the colors of the rainbow, and is the basis for every color computer screen today (was also one of Young's significant accomplishments.

It was through his interest in the perception of light that he came to study the effects of light passing through tiny holes in opaque screens. When he passed light through two such holes and allowed the light transmitted from each to overlap on a distant screen, he observed bands of light alternating with bands of darkness. This was an astounding discovery that defied perceived common sense at the time: that light added to light could produce darkness. Yet this is precisely the interference effect that I employed to described holography in the last chapter. Young was able to explain the effect as a consequence of the wave nature of light. He went on to explain the colors of soap films based on this theory, as well as to explain polarization of light waves. Despite his genius, he was disparaged by the professional English physicists of his time, principally because Isaac Newton had proposed that light was composed of particles. In England, to disagree with Newton was sacrilege and heretical. On the other hand, the continental physicists were not so loathe to debunk Newton, and Young's work gained

wider acceptance after work by the French physicist Augustin-Jean Fresnel (1788 - 1827) confirmed Young's hypothesis.

An idealized experimental arrangement of Young's double pinhole experiment is shown in Fig. 9.1. A single pinhole emits a single color and illuminates two pinholes situated a small distance apart in an opaque screen. The light emitted from each pinhole illuminates a distant viewing screen, and the field of illumination of each hole overlaps with the other. On the observing screen, bands of light alternate with bands of darkness, demonstrating the coherent interference of the light coming from both pinholes. A bright band on the screen is obtained when the difference in the distances from the two holes allows the waves to add constructively. Conversely, a dark band on the screen is obtained when the difference in the distances from the two holes allows the wave amplitudes from each pinhole to subtract to produce destructive interference. If either pinhole is blocked by an opaque obstruction, the interference pattern disappears and is replaced by an even illumination from the unobstructed pinhole.

Up to this point, the discussion has been purely classical. But now we do a simple experiment to take us out of the classical regime and into the quantum realm. We reduce the intensity of the source so that it becomes extremely weak (much weaker than even moonlight. At this stage we need to relate the intensity of light to the flux of photons, Photons carry quantized units of light energy. For instance, a single photon of light has an energy of a few electronvolts (1 eV is the energy an electron would gain after accelerating across one Volt). A light beam with an intensity in units of energy per second per area can therefore be described as a stream of individual photons, like drops of rain in a shower, in units of photons per second per area. In the pinhole experiment we can reduce the intensity of the source so low that there can only be a single photon in flight at a time between the source and the observation screen. We replace the screen with a photosensitive plate that records the arrival of photons, allowing the photon hits to

accumulate over time on the plate. This plate is something analogous to the photodetector called a CCD (charge-coupled device) used in your video camera.

When we turn on the experiment, the photo-plate responds at the positions where it records a photon. After only several photons have been detected, the photo-plate may look like Fig. 9.2a. No discernible pattern of bright and dark bands can be seen. As we continue the experiment, the plate may look like Fig. 9.2b. Now, there are the beginnings of a pattern, but it is still rough. However, after continuing for a longer time, the photo-plate looks like Fig.9.2c. In this case, the interference bands are easily visible and begin to look like the bands of bright and dark that we see in the classical experiment.

This new version of the experiment is fully in the quantum domain. The light travels as photons and arrives at the photo-plate as photons. The photo-plate responds at specific positions that are hit by single photons. There is no room in this description for classical electromagnetic waves, nor even of the interference of electric fields. The absence of classical interference is made clear by the conditions of the experiment that allow only a single photon to be in flight between the source and the screen at a time. Since only one photon is present, its electric field neither adds constructively to nor subtracts destructively from the electric field of any other photon. Yet the interference pattern slowly develops on the photo-plate, just as in the classical interference experiment. Where does this interference pattern come from if the photons cannot interfere with each other?

The quantum answer is that the photon interferes with itself: an answer that warrants considerable discussion. Though the photon has an electric field associated with it, we cannot view the quantum experiment as an interference of the electric field of the photon with its own electric field. Instead, something else must be interfering to generate the interference pattern that we see accumulating on the photo-plate. To understand what is interfering, we need first to understand something of quantum wave mechanics.

Wave mechanics for quantum systems was developed in 1927 by Erwin Schrödinger (at age 40). Schrödinger was able to show that the behavior of quantum particles could be understood as special functions, called wavefunctions, that obeyed a straightforward wave equation that came to bear his name. The result of this theory was an understanding that, in the quantum world, particles behave like packets of waves. This is the famous the wave-particle duality that has so perplexed quantum philosophers (how best to understand objects that are both particles and waves at the same time.

The meaning of the wavefunction of an elementary particle, like an electron, was not initially obvious. That interpretation was supplied by Max Born, a German theoretical physicist at the University of Göttingen. He suggested that the squared amplitude of a quantum wavefunction at a place and an instant in time is proportional to the probability for finding an electron at that place and time. This interpretation was radical (equally as radical as the original quantum hypothesis. Whereas the wavefunction of an electron could be accurately and uniquely specified by giving the state properties, the electron's location could only be predicted by a probability governed by the amplitude of the wavefunction.

If you take a hundred atoms, all in exactly the same quantum state (that is, the electrons of each atom are all described by the same wavefunction), and measure the positions of the electrons in each of them, you will get a hundred different answers. But if you continue preparing more atoms in identical states and measure those, you will slowly build up a distribution of electron positions that tended to occur close to the nucleus of the atom. With enough measurements on enough atoms, you would eventually have a smooth distribution of electron positions that exactly matched the squared amplitude of the electron wavefunction for that quantum state. The important feature of this description is the difference between where the electron is *before* the measurement, and where the electron was found *during* a single measurement.

The wavefunction for an electron is a well-defined smooth function of position. At any radius away from the atom nucleus there is some value for the wavefunction. We therefore say that an electron *occupies* this wavefunction, meaning that the electron is simultaneously everywhere where the wavefunction has some non-zero value. In this sense, an electron surrounds the nucleus all the time. But in the act of measurement, let us say a measurement of the position of an electron using a microscopic probe of some kind, a single electron must be found at only a single location. The measured location of the electron is merely one place that the electron was before the process of measurement, but is certainly not the only place that the electron occupied. The electron occupied all locations (where the wavefunction was non-zero) prior to the measurement, and the measurement merely happened to find it at a specific spot.

Photon wavefunctions behave in the same way. When Young's apparatus contains only a single photon, that photon is governed by a single wavefunction. The wavefunction fills all the space inside the apparatus, just like the electron wavefunction filled all space around the nucleus of the atom. Part of the wavefunction passes through one pinhole, and part of the same wavefunction passes through the other. When these parts of the wavefunction overlap on the screen, the amplitudes of the wavefunction add and subtract in just the way that the electric fields of classical light waves would. When the path length differences make the crests and troughs of the waves line up, constructive interference occurs, and the squared amplitude of the quantum wavefunction is a maximum. Using Born's interpretation, this means that there is a high probability that the photon will be detected at this location. On the other hand, in regions of destructive interference, the squared amplitude of the photon wavefunction vanishes, as does any chance to observe the photon at that position.

This explains the results of quantum experiments that use one photon at a time. Even though a photon cannot interfere with any other photons, its quantum wavefunction at the observation plane has regions of constructive and destructive interference. In

regions of destructive interference, the photon can never be detected, no matter how long the experiment is carried out. This is why the photon detections accumulate only in the regions of constructive interference, which are precisely where the electric fields of a classical wave would interfere constructively and produce high intensity.

You will note that the quantum theory gives the same answer as the classical theory. It looks like a sleight of hand to say that a photon is governed by a probability wave, and that it is probability that interferes, while at the same time the classical interference of the light fields produces exactly the same intensity pattern. This starts to look like a metaphysical question. If the quantum theory predicts an outcome that is identical to classical theory, do we really care? And more importantly, can the quantum theory predict any behavior that is impossible classically? The answer is yes (in volumes! The entire field of quantum information rests on specific differences between classical and quantum behavior. But before I can move to those topics, the unnatural quantum behavior of photons is best understood with a process known as "quantum seeing in the dark".

QUANTUM SEEING IN THE DARK

In your mind's eye, envision a diabolical terrorist who places Young's apparatus in a crowded theater. Inside the apparatus there may be, or may not be, a bomb that will detonate any time it is hit by a photon. If the bomb is there, it is placed behind one of the pinholes. As a member of the Quantum Bomb Squad, you are called in to determine whether the apparatus contains the bomb or not. Your goal is to detect the presence of the bomb without detonating it. However, all you can use to detect the bomb is a source of photons. How would you use photons to detect a photo-sensitive bomb without detonating it? If you shine photons (let's say by opening the apparatus) on the bomb to

see if it is there, then it will go off and you lose your job (if not your life). But if you were a good quantum student in school and have full faith in the Born interpretation of the quantum wavefunction, you devise a way of using quantum interference to detect the bomb without detonating it, at least with odds you are willing to live with. This is what you do.

You take the apparatus, turn on the photo-plate, and then hold your breath as you send in a single photon from the source. The single photon can pass either through the open hole, or through the hole with the bomb behind it. It will do either with a 50% probability. If it passes through the hole with the bomb, then it detonates, destroying valuable property, and you lose your job on the Quantum Bomb Squad. On the other hand, if it passes through the open hole, it will register a flash on the photo-plate. This is where your understanding of quantum mechanics is crucial.

If the photon hits a location on the photo-plate that would be inaccessible when *both* pinholes were clear, that is, a position of destructive interference caused by the wavefunction interference from the two pinholes, then the bomb must be present and you should evacuate the theater. In this result, *you* have detected the bomb with a photon, yet the *bomb* detected no photon because it passed through the open pinhole. How does the photon detect the bomb without detonating it, or even touching it? The answer is that the photon wavefunction extends throughout the apparatus. If the bomb is blocking the pinhole, it also blocks the photon wavefunction and prevents interference at the photo-plate. Therefore, blocking the wavefunction is not the same as blocking the photon itself. The wavefunction just determines where the photon is *likely* to go.

Unfortunately, the odds are not great that the photon will hit exactly at a location of complete destructive interference. It is more likely that the result will be ambiguous (by hitting in a location that would be accessible whether the pinhole is blocked or not). Then you will need to send in another photon, with another 50% chance of detonation. And if that result is ambiguous, you need to send in yet another photon, until you are

most surely going to need to find a new job. Fortunately for you, there are higher-probability ways of detecting bombs in the dark.

One way is to replace Young's apparatus with the simple interferometer shown in Fig. 9.3. This configuration uses a half-silvered beamsplitter to split the possible photon paths. When photons hit the beamsplitter, they either continue traveling to the right, or are deflected upwards. After reflecting off the mirrors, the photons again encounter the beamsplitter, where, in each case, they continue undeflected or are reflected. The result is that two paths combine at the beamsplitter to travel to the detector, while two other paths combine to travel back along the direction of the incident beam.

The paths of the light beams can be adjusted so that the beams combining to travel to the detector experience perfect destructive interference. In this situation, the detector never detects light, and all the light returns back along the direction of the incident beam. Quantum mechanically, when only a single photon is present in the interferometer at a time, we would say that the quantum wavefunction of the photon interferes destructively along the path to the detector, and constructively along the path opposite to the incident beam. Again, the detector would detect no photons. It is clear that the unobstructed path of both beams results in the detector making no detections.

Now place the light sensitive bomb in the upper path. Because this path is no longer available to the photon wavefunction, the destructive interference of the wavefunction along the detector path is removed. Now when a single photon is sent into the interferometer, three possible things can happen. One, the photon is reflected by the beamsplitter and detonates the bomb. Two, the photon is transmitted by the beamsplitter, reflects off the right mirror, and is transmitted again by the beamsplitter to travel back down the incident path without being detected by the detector. Three, the photon is transmitted by the beamsplitter, reflects off the right mirror, and is reflected off the beamsplitter to be detected by the detector.

In this third case, the photon is detected AND the bomb does NOT go off, which succeeds at quantum seeing in the dark. The odds, now, are much better than for Young's experiment. If the bomb is present, it will detonate a maximum of 50% of the time. The other 50%, you will either detect a photon (signifying the presence of the bomb), or else you will not detect a photon (giving an ambiguous answer and requiring you to perform the experiment again). When you perform the experiment again, you again have a 50% chance of detonating the bomb, and a 25% chance of detecting it without it detonating, but again a 25% chance of not detecting it, and so forth. All in all, every time you send in a photon, you have one chance in four of seeing the bomb without detonating it. These are much better odds than for the Young's apparatus where only exact detection of the photon at a forbidden location would signify the presence of the bomb.

It is possible to increase your odds even above one chance in four. You do this by decreasing the reflectivity of the beamsplitter. In practice, this is easy to do simply by depositing less and less silver on the surface of the glass plate. When the reflectivity gets very low, let us say at the level of 1%, then most of the time the photon just travels back along the direction it came and you have an ambiguous result. On the other hand, when the photon does not return, there is an equal probability of detonation as detection. This means that, though you may send in many photons, your odds for eventually seeing the bomb without detonating it are nearly 50%, which is a factor of two better odds than for the half-silvered beamsplitter.

These are about the best odds you are going to get, but this is impressive in itself. To be able to "see" something without ever having the photon "touch" it is only possible in a quantum world. This serves to illustrate how one must reason when dealing with quantum systems, and it gets us in the habit of thinking about photons and beamsplitters, which will be useful when we begin our discussion of entangled photons. But before we talk about that, we need to understand a physical property of photons called photon

polarization, because quantum information can be stored in the two orthogonal polarizations of light.

PHOTON POLARIZATION

One of Thomas Young's innumerable contributions to physics was the idea of polarization. He correctly understood that the electric field of light has an orientation perpendicular to the direction of light propagation. This means that if you look directly at the source of a lightwave, the electric field of the wave will lie in a plane, as shown in Fig. 9.4. When the electric field points in a constant direction, say along a 45º diagonal, we say that the light has linear polarization.

In a plane, there are always two mutually orthogonal directions, like the x and y axes. Therefore, the polarization of any linearly polarized wave can be decomposed into components that point along the two mutually orthogonal directions. The choice of a direction, like a vertical axis or a diagonal axis, automatically defines the other. For instance, for the electric field in Fig. 9.4 we can describe this vector as a sum of two vector components, one vector, denoted by E_H, pointing along the horizontal axis, and another vector, denoted by E_V, pointing along the vertical axis. The sum of the two vector components yields the total field E. Alternatively, we can choose the mutually orthogonal axes V' and H'. These axes define the field vector E just as well as before, producing two projections $E_{V'}$ and $E_{H'}$ which, again, add up to yield the total field E. It is important to recognize from this example that there is nothing sacred about the word "vertical". We are free to choose and define any axis as "vertical", and the orthogonal axis as "horizontal". Any choice is arbitrary, yet equally valid as a way of describing the electric field.

When we stop thinking of light as classical electromagnetic waves and think instead of photons, the notion of polarization remains, but the interpretation changes. A photon has a polarization just like a classical light wave, but the polarization now is associated with the photon wavefunction. If a photon is originally polarized at an angle relative to the horizontal, we say that the wavefunction is a linear combination of *two* wavefunctions, one that has a polarization along the vertical and another that has a polarization along the horizontal.

The linear combination of orthogonal wavefunctions is one of the most important features of quantum theory. This concept is tied up closely with the idea of what it means to make a measurement (or an observation) on a quantum system. Let us say that we have a quantum wavefunction Ψ that is the sum of two orthogonal wavefunctions ψ_V and ψ_H such that

$$\Psi = a\Psi_V + b\Psi_H$$

where Ψ is the total wavefunction and ψ_V and ψ_H are the two orthogonal wavefunctions into which Ψ can be decomposed. The wavefunction is understood in terms of the Born interpretation that says that the squared amplitudes a^2 and b^2 are equal to the probability of observing the photon with a vertical polarization or horizontal polarization, respectively.

Nature has provided us with an ideal method for separating a flux of photons into the ones polarized along V from the others polarized along H. This is accomplished using a crystal of calcite that has the chemical name of calcium carbonate ($CaCO_3$). It is the most common constituent of limestone and marble. In its pure crystalline form it is transparent and colorless, and is noted for its property of double refraction; anything you look at through the crystal has a double image. This is because natural light has equal amounts of orthogonal polarizations, and the calcite crystal directs light of the two

polarizations along two different directions, as shown in Fig. 9.5. Light with polarization out of the plane of the figure (V polarization, where the dots on the beam represent the tips of electric field arrows pointing out of the plane at you) travels straight through the crystal, while light polarized in the plane of the figure (H polarization, shown with the electric field arrows in the plane) is refracted at the surface. The H beam is deflected at an angle of 6º (dictated by the specific crystal structure). At the exit plane of the crystal, the two orthogonal polarizations have been separated. Each beam has a pure polarization: one vertical, the other horizontal.

The calcite crystal is known as a polarization analyzer. It takes any input beam and breaks it down into its V and H components. As an optical device we say that it has one input port, and two exit ports. Detectors are placed at both of the exit ports. If the crystal is very pure, none of the light energy is absorbed, and all of the light is detected. For a classical light wave with a polarization angled at 45º relative to the horizontal axis, half of the intensity is detected in the V port, and half is detected in the H port because the electric field of the photon has equal parts of V and H.

Now let's consider the quantum behavior of the calcite when we send a single photon polarized at 45º. In the quantum case, the entire photon emerges whole from either one port or the other, but never both ports, and never as a fraction of a photon. For example, a photon with a polarization of 45º has a quantum wavefunction given by

$$\Psi = \frac{1}{\sqrt{2}}\left(\Psi_V + \Psi_H\right) \ .$$

which is an equal combination of H and V photon polarizations. It has 50% probability (the square of the $\frac{1}{\sqrt{2}}$ coefficient) of exiting the V port, and a 50% probability of exiting the H port. If it exits from the H port, it has 100% H polarization. Similarly, if it exits from the V port, it has 100% V polarization. In all cases, the photon is detected whole in one detector or the other, but never both. We do not detect half a photon in each detector.

When we consider the action of the calcite crystal on a single photon, we may ask an apparently simple question: Does the crystal simply observe the polarization of the photon, or does it modify it by rotating its polarization? Let us assume, for the moment, that the second option is true, that the crystal modifies the photon. Many materials rotate the polarization of a light beam. For instance a solution of corn syrup rotates the polarization of a light beam as it propagates through the liquid by a process known as optical activity. Corn syrup is optically active because the sugar molecules, called dextrose, in the syrup have only a right-handedness. Light polarizes the dextrose molecules, and the radiated light field is rotated slightly to the right. This effect accumulates over distance into a macroscopic rotation of the polarization of the light beam.

With this in mind, we can try to explain the effect of the calcite crystal as polarization rotation. For a 45º polarized photon passing through a calcite crystal, the crystal either rotates the polarization right by 45º, or left by 45º. But this is not deterministic as it was in the case of the corn syrup. For the syrup, the rotation was always to the right. In the case of the calcite, the photon polarization for a series of identical photons is rotated right for half of them (on average) and left for the other half. On any given instance it is impossible to predict which will occur. Since the result is indeterminate, it is impossible to assign a specific physical rotation mechanism to the process. Therefore, we have no choice but to accept that calcite is *not* rotating the polarization, but rather is making a quantum observation, i. e., determining whether the photon has V or H polarization.

This example illustrates the fundamental indeterminacy of quantum mechanics. It is impossible to predict with certainty which polarization will exit the crystal for a single incident photon. This type of indeterminacy was what Einstein was unwilling to accept. He viewed this type of experiment as evidence that quantum mechanics was incomplete. In this regard, it is important to make the distinction between "incomplete" and

"incorrect". Einstein never considered quantum mechanics to be incorrect. He was fully aware that quantum theory accurately predicted the outcomes of quantum experiments. In fact, he was the theoretician who, using quantum theory, correctly predicted many of those outcomes. Einstein's argument against quantum mechanics was rather that, in those areas where it could say nothing, as in the prediction of the result of a single observation of a single quantum particle, some deeper and more complete theory *could* predict the outcome. It is in this sense that Einstein considered quantum mechanics to be incomplete.

THE EPR PARADOX

It is perhaps fitting that the most imaginative and sustained attack on the completeness of quantum theory was devised by Einstein (along with Boris Podolsky and Nathan Rosen) in the EPR paradox of 1935. The paradox was introduced briefly in Chapter 2, but let me take the time now to describe the paradox in more detail because an understanding of the paradox is a necessary starting point for later discussions of quantum teleportation and quantum entanglement. To illustrate the paradox, I use a formulation along the lines proposed by the physicist David Bohm that is simpler to think through. This formulation begins with the self-annihilation of an atom-like entity called positronium into two photons.

Positronium is an electron bound to its anti-matter pair, a positron, in a quantum state similar to that of a hydrogen atom. Unlike hydrogen (which is stable for times at least as long as the age of the universe), the electron and positron annihilate each other in a flash of energy that produces two gamma rays. The atom lives for only about a tenth of a microsecond, on average, before annihilation. When the positronium is initially at rest, and is in its ground state, the atom has no linear momentum and no angular momentum

[NOTE: Angular momentum is associated with any rotating mass, like the rotation of a bicycle wheel or a spinning top.]. By the law of conservation of momentum, the final state must also have no net linear or angular momentum. We therefore immediately conclude that the two photons must travel in opposite directions, carrying equal amounts of energy and momentum, and the sum of their individual angular momenta must be zero.

Now consider a thought experiment in which many individual positronium atoms sequentially self-annihilate, and the linear polarization of the two decay-product photons are observed by two observers, we will call them observer A and B, who are located opposite each other and very far away from the source. These two observers make their measurements at times t_A and t_B, respectively. By varying their distance from the source, the observation made at time t_A can be either earlier or later than the observation made at time t_B. Finally, we insist that the difference between the observation times must be much shorter than the time it takes for photons to travel from the source to either observer. This ensures that no information about one measurement *causally* (that is, traveling at the speed of light) affects the outcome of the other measurement.

It has become a well-established tradition, in discussions of this sort, to give the name "Alice" to observer A, and the name "Bob" to observer B as a fairly convenient device. In Bohm's thought experiment, therefore, Alice and Bob each have a crystal of calcite with single-photon detectors placed at both the H and V output ports of the crystal [NOTE: Calcite cannot be used to analyze the polarization of gamma rays emitted from the annihilation of positronium. In most table-top experiments that study the EPR paradox visible or near-infrared photons are generated by Calcium atoms or by special nonlinear crystals called down-conversion crystals.]. They choose any angle θ_α (by Alice) and θ_β (by Bob) from observation to observation. They can also move their apparatus farther away from the source, or nearer to it, thereby altering their measurement times. By doing this, measurement A sometimes will be first, and at other times measurement B will be first. The measurement time and the crystal orientation are

all chosen independently and randomly for each photon detection event. Neither of the observers knows what the other is doing. Each observer makes a large number of observations, recording in their notebooks the time and angle of the crystal, as well as whether the H detector or the V detector flashed for each case. When a flash is observed, the value of the measurement is written down as a "1", and otherwise as a "0". Each observer has two observed values for each event, one for the H port of the crystal, and the other for the V port.

What each observer sees locally, as the experiment is progressing and as they randomly choose the measurement time and angle, is a perfect anticorrelation between their own H and V ports of their calcite crystal. Whenever their H detector flashes, the V detector does not flash, and vice versa. Each photon exits the crystal in either one port or the other, but never both and never none. The observers also note that there is equal probability for the photon to appear in the H port as the V port. In other words, their local data is extremely uninteresting; they just see a long random string of photon hits in either one detector or the other with a 50/50 probability for each detection regardless of what angles they choose for their crystals. No other structure is visible in their data.

When the experiment is over, the observers pack up their equipment and travel back to the source to compare their seemingly random data of "1"s and "0"s. To perform the comparison, they multiply Alice's first number with Bob's first number. Then they multiply their second numbers together, and so on for each individual measurement. Once they have a list of products, they group the products according to the difference in their chosen angles, $\theta_\alpha - \theta_\beta$, and average all the results for a given angle difference. Finally, they plot the averages against the difference in the measurement angles. This final function is written as $P(\theta_\alpha - \theta_\beta)$. What they find is shown as the solid curve in Fig. 9.6. It has a very simple shape ⌠ that of a perfect sinusoid. The function is simply $P(\theta_\alpha - \theta_\beta) = \cos^2(\theta_\alpha - \theta_\beta)$.

This function has a simple interpretation. Regardless of the value θ_α chosen (at random by Alice and unknown to Bob) for V, the probability that Bob will detect a V photon is simply equal to $\cos^2(\theta_\beta - \theta_\alpha)$. Equivalently, regardless of the value θ_β chosen (at random by Bob and unknown to Alice) for V, the probability that Alice will detect a V photon is again simply equal to $\cos^2(\theta_\alpha - \theta_\beta)$. All that matters in the experiment is the randomly chosen difference between the measurement angles. This looks innocuous enough (but a deep mystery is actually contained in this simple function. To expose this mystery, we need to look more closely at the photon polarizations that come out of the positronium decay.

We know that the two decay photons must have equal polarizations. When $\theta_\alpha = \theta_\beta$ the correlation function is $P(\theta_\alpha - \theta_\beta) = \cos^2(\theta_\alpha - \theta_\beta) = 1$, which means that when Alice and Bob both pick the same angle (accidentally), they both see the same polarization. This result is true no matter what common value $\theta_\alpha = \theta_\beta$ they chose for the angles. This means that if Bob chooses the angle θ_β and observes his photon emerging from the V port, then if Alice has also chosen $\theta_\alpha = \theta_\beta$, she observes her photon in the V port, and vice versa. This result holds whether Bob measures first or Alice measures first as they randomly vary their measurement time by changing their distance from the source.

One way to interpret this result is that when a polarization measurement is made on one photon, the twin photon *instantly* acquires the identical polarization. The effect is instantaneous, which means that no matter how far apart the two photons are when the first measurement is made, whether they are at opposite sides of an experimental optical bench in a laboratory, or are at opposite sides of the universe, as soon as the first measurement is made, the second photon instantaneously assumes the identical polarization. This "influence", being instantaneous, must therefore occur at speeds exceeding the speed of light. Such an "influence", or effect, is called "nonlocal" to contrast it with conventional forces that only exert their influence at speeds limited by the

speed of light. It is precisely this nonlocal nature of the effect that Einstein and his EPR colleagues objected to, and that violated their sense of physical reality.

This paradox, to Einstein's thinking, was further evidence that quantum mechanics was incomplete. Just as he was unwilling to accept that each quantum event occurred at random, he also believed that nonlocality was unphysical. To banish nonlocality, as well as randomness, from the interpretation of quantum theory required the existence of some unknown element that determined ahead of time what polarization a photon would assume during a measurement. This unknown element is called a hidden variable.

Hidden variable theories sprang up in abundance in the early days of quantum mechanics in attempts to solve the randomness and nonlocality problems. One idea was that each quantum particle carried along with it some hidden variable that determined whether it would pass through the V port or the H port. Such a hidden variable would solve the nonlocality problem because the photon polarizations are predetermined. Each photon would already know how their twin would pass through a polarizer and therefore would require no influence traveling faster than the speed of light to tell them.

Imposing a hidden variable, like a common (but unknown) polarization on twin photons sounds like a condition we could impose in an real-world experiment. Let's try this route, and see where it leads us.

THE EPR DEMON

Consider a black box that contains a source of positronium. It also holds two calcite crystals oriented at common angles, and a miniature Demon whose job it is to select the angle randomly for each positronium decay [NOTE: Demons are a common device used to describe certain paradoxes in physics. One of the earliest and most

famous uses of a Demon was by Maxwell, called Maxwell's Demon, to describe apparent violations of the second law of thermodynamics]. The arrangement of the black box is shown in Fig. 9.7. The common angle the Demon chooses for the calcite crystals is denoted by θ_γ. There is an A side and a B side to the black box. The A side sends the photon to Alice and the B side sends the photon to Bob. When the demon selects a specific θ_γ for a decay, the photon on the A side will either exit through the V port or the H port of the Demon's calcite crystal. Only when the photon goes through the V port does it escape from the box through the opening. Because the photons from the positronium are randomly polarized, there is a 50/50 probability that the photon will escape through the V port. What happens on the B side?

The B-side calcite crystal has the same angle θ_γ as the A-side crystal. Without making any assumptions about the properties of the two photons coming from the positronium, we know that if a photon exits from the B-side, that it has the same polarization as the photon that was emitted from the A-side. Therefore, when two photons are emitted from the black box, we know that they both have the same polarization given by θ_γ. For the moment we will not worry whether both photons always are emitted, or whether a photon might be emitted on one side but not the other. We are concerned only with the case when two identically polarized photons are emitted from both sides [NOTE: If we wish, we could construct a more sophisticated black box that used photodetectors on the H port that would open a shutter on the exit port only when both detectors detected no photon. This would guarantee that the black box emitted only pairs of photons that had identical polarizations.]. Now we come to the specific problem of the Demon.

The Demon knows what angle θ_γ he chooses each time, but neither Alice nor Bob know. Therefore θ_γ is a variable that is hidden from the observers but that uniquely determines each photon's polarization. If the Demon were to communicate with the observers, they could set their own crystals to the common angle θ_γ and would have a

100% probability of detecting each photon. But the Demon is not so inclined. Therefore Alice and Bob randomly choose their own angles θ_α and θ_β for each measurement. After a long series of measurements they meet and compare their data, just as they did when they were looking at the bare positronium decay.

What have we gained by putting the Demon in the black box? We have a source of two photons which have identical polarizations. But we know from momentum conservation that the two photons coming from a bare positronium decay also have identical polarizations. Aren't these the same thing? Since the Demon does not communicate with the observers, all he has seemed to accomplish is to substitute one set of random polarizations for another. But let's look at what Alice and Bob see when they compare their data.

The probability for both getting a detection for a single event, including the hidden variable θ_γ, is given by $P(\theta_\alpha, \theta_\beta, \theta_\gamma) = \cos^2(\theta_\alpha - \theta_\gamma) \cos^2(\theta_\beta - \theta_\gamma)$, which is just the product of the individual intensities coming out of Alice's and Bob's polarization analyzers. Because they do not know the angle set by the Demon, they must average over many events in which the Demon is choosing his angle randomly. When they do this, they get the dashed curve shown in Fig. 9.6. It is not the same as for bare positronium! It has the same general shape as the quantum prediction, but its amplitude is smaller. Most importantly, the probability for joint detection never goes to zero as it does for the quantum prediction. Sometimes, even when Bob and Alice chose orthogonal angles (which gives zero joint probability for the free-decay case) they can both detect a photon.

What has happened here? Both the bare positronium and the Demon produce pairs of photons that have identical polarizations. In both cases, the value of that polarization is unknown to the observers, so all they can do is make separate observations and compare their results. So why aren't these two situations the same? The only difference is the presence of the hidden variable θ_γ which neither Alice nor Bob know.

Can this have such a large effect? The answer is yes. The act of measurement performed by the Demon, even though it is unknown to observers A and B, alters the physical state of the two photons from the case of the bare positronium decay (without the Demon). This difference between having an intermediate measurement or not shows up in a very real way in the measurements performed by the observers.

The failure of our hidden variable model does not necessarily mean that *all* hidden variable models would fail to reproduce the behavior predicted by quantum mechanics. Indeed, a minor industry of hidden variable theories sprang up after the EPR paper and produced fairly ingenious theories that could explain quantum predictions (one of the most notable of these being from David Bohm. It was in the context of hidden variables that he proposed his alternative EPR paradox based on measuring polarizations. His hidden variable theory was considerably more sophisticated than the one I presented. Nonetheless, we can still ask whether any of these hidden variable theories might actually be able to complete the quantum picture of reality.

This was the question asked by John Bell, an Irish physicist working at CERN in the early 1960's. He proved, using arguments about probabilities, that all hidden variable theories (if they permitted only local interactions among particles) must be false [NOTE: John Bell]. The proof was surprisingly simple, and produced what has come to be called the Bell Inequality. Any local hidden variable theory must satisfy the inequality. Quantum systems, on the other hand, violated the inequality. Despite the simplicity of the argument, every theoretical prediction must be followed by experimental verification. Devising a physical experiment that unambiguously demonstrates a violation of Bell's inequality was a challenging prospect. The definitive demonstration came in 1981 - 1982 when Alain Aspect and his research group performed a series experiments of increasing sophistication that violated Bell's inequality with extremely high confidence [NOTE: Ref Aspect]. The most important aspect of these experiments, and the aspect that made them so difficult, was the need for the detection events of the

two photons to be separated far enough so that no signal moving at the speed of light could travel from one side of the experiment to the other during the time of the measurements. This condition was absolutely necessary to guarantee that no local interaction (defined as an interaction limited by the speed of light) could explain the correlation between the two measurements.

The experiments by Aspect used an atomic beam of calcium atoms in excited states that radiated two photons as they fell back to their ground state. The two photons carried away polarizations in the same way as the two photons from positronium. The use of calcium instead of positronium significantly simplified the experiments because the atomic beam produced copious numbers of visible photons that are relatively easy to analyze for polarization. The initial experiment was no more complicated than the problem of measuring individual polarizations with two analyzers. Already in this case they observed large deviations from Bell's inequality and hence firmly established the nonlocality of quantum mechanics. However, nagging suspicions of local influences persisted among the experiment's critics, leading Aspect and his team to devise an ingenious technique that allowed them to select the polarization angle after the photons were already in flight [NOTE: Ref Aspect delayed choice]. These experiments continued to agree with quantum mechanics and violate Bell's inequality. Since these experiments delayed the choice of polarization until after the photons were in flight, there was no way for the photons to have shared a local hidden variable when they were created.

These experiments unambiguously proved that *all* hidden variable theories that were concocted to solve the nonlocality problem *are wrong*. None of them will ever give results that agree with quantum theory. The inescapable conclusion is that quantum mechanics is nonlocal (the instant that one measurement is performed on one member of a pair of twin photons, the other photon's quantum state is immediately known, even if that state is on the other side of the universe. This statement is provably true, as John Bell demonstrated in 1964. Once nonlocality is accepted, the next most pressing question

is whether this nonlocality can be used to communicate faster than the speed of light. We will see that the answer to this question is "no." But we will also see that twin photons are useful for quantum communication and computation, and they even provide the basis for quantum teleportation. To understand these points, we need to look closer at the quantum properties of the twins.

ENTANGLED PHOTONS

We have just seen that there is a fundamental difference between the pair of photons that are prepared in a specific polarization by the Demon or any other deterministic process (even if that polarization is unknown), and the pair of photons that emerge from a single quantum event such as the annihilation of a positronium atom. In both cases the photons must have equal polarizations when they are measured along the same directions. Yet their behavior is distinctly different when the measurement angles are not the same.

Let's return to the case of the EPR Demon when neither Alice nor Bob know the hidden variable θ_γ. For a single decay event, assume that the Demon has chosen $\theta_\gamma = 45^o$, and Alice and Bob both choose $\theta_\alpha = \theta_\beta = 90^o$. In this case there is a 50% probability that Alice will observe her photon coming in her V port and an equal 50% probability that she will observe the photon in the H port. Bob experiences the same odds. But that is the point: Bob shares the same odds as Alice, but not the same experience. In any given event, Alice may observe a H photon while Bob observes a V photon, and vice versa. Though they have the same odds for detecting V or H, nothing says they will see the same result for the same event.

Now contrast this to the case of the photons emitted from the positronium in the absence of the Demon. Again let Alice and Bob both choose $\theta_\alpha = \theta_\beta = 90^o$. In this case

there is again a 50% probability that Alice will observe her photon coming in her V port and an equal 50% probability that she will observe the photon in the H port. Bob experiences the same odds. But now, Bob sees *exactly* what Alice sees. If Alice observes a V photon, Bob observes a V photon, and vice versa. Because they have chosen the same measurement angles, they obtain exactly the same answers.

The photons from the positronium have a redundancy about them. Once Alice makes her measurement, Bob's measurement is redundant. If we know Alice's results, then we can say with certainty what Bob will see if he chooses the same angle for his crystal. Because of this redundancy, the quantum pair of photons are said to be "entangled" in a *single* quantum state. Rather than each particle having its own quantum wavefunction, both particles share a single quantum wavefunction. Performing a single measurement on a single photon already constitutes a measurement of the whole quantum wavefunction, so performing the second measurement on the second particle is not needed. If Alice sees her photon in her V port, then the vertical polarization is shared by both particles, so Bob's particle is immediately known to also have vertical polarization.

The quantum wavefunction of an entangled pair of particles is written as a combination (or sum, or linear superposition (these terms are all interchangeable) of two situations. One situation is that both Alice and Bob observe V photons. The other situation is that they both see H photons. The total entangled wavefunction Ψ^e is given by the combination of these two situations as

$$\Psi^e = \frac{1}{\sqrt{2}}\left(\Psi_V^A \Psi_V^B + \Psi_H^A \Psi_H^B\right)$$

where the superscript refers to the observer and the subscript refers to what they observe. This is notably different than the wavefunction for the pair of photons prepared by the Demon, which looks like

$$\Psi^e = \frac{1}{\sqrt{2}}\left(\Psi_V^A + \Psi_H^A\right)\frac{1}{\sqrt{2}}\left(\Psi_V^B + \Psi_H^B\right)$$
$$= \frac{1}{2}\left(\Psi_V^A\Psi_V^B + \Psi_H^A\Psi_H^B + \Psi_V^A\Psi_H^B + \Psi_H^A\Psi_V^B\right)$$

where the first two terms in the last line are equal to the terms in the entangled wavefunction Ψ^e, but there are two additional terms. These additional terms allow Bob to observe an H photon when Alice observes a V photon, and vice versa. These terms are specifically missing from the entangled wavefunction Ψ^e. The most important feature of the Demon's wavefunction is that it can be factored into a term that depends only on Alice (the first term in the parenthesis on the first line), and a term that depends only on Bob. The entangled wavefunction cannot be factored this way. These differences between Ψ^e and Ψ^d are what makes the solid and dashed curves different in Fig. 9.6 when correlated polarization measurements are made on the two types of pairs.

Apart from these mathematical details, there are severe metaphysical problems that entangled pairs of particles present to philosophers. Even if the two entangled particles are separated by the diameter of the universe, they still belong to the same quantum wavefunction. In this sense, the nonlocality problem is primarily the problem with a macroscopic quantum wavefunction. It is a challenge to think of a quantum wavefunction, something that is supposed to operate at atomic and subatomic scales, extending over the size of the universe. As we saw, one viewpoint is that the common polarization shared by the two particles is indeterminate until the moment of measurement. At that moment, as one particle assumes a specific property, the entangled twin instantaneously assumes the same property. This viewpoint is known as *wavefunction collapse*. If the wavefunction is macroscopic, extending over long distances, the common wavefunction shared by both particles collapses at the moment of measurement, regardless of who makes their measurement first.

By taking this view, we can convince ourselves that making a measurement here and now on our side of the Universe instantly affects the state of the twin member of an entangled pair made on the other side of the Universe, regardless of any limits imposed by the speed of light. This interpretation is exactly what Einstein and his friends objected to, and exactly what hidden variable theories had attempted to dispense with. To this day, no satisfactory agreement has been reached between the pragmatists who merrily perform their experiments free from any guilt about philosophical ramifications, and the quantum philosophers who worry about the "real" meaning of entanglement.

From the pragmatic point of view, the instantaneous nature of wavefunction collapse does not provide a means of sending information faster than the speed of light. It is tempting to try to construct a quantum communication system in which Bob and Alice receive a steady stream of entangled particles from some central source. Bob chooses his crystal angles to be either 0º or 90º, with 0º corresponding to a "0" bit and 90º corresponding to a "1" bit. By making successive measurements on his particle, he collapses the wavefunction instantaneously at Alice's location. If he sees the photon come out of his V port, then she will also see her photon come out of her V port (that is, if she has happened to choose the same crystal angle as Bob. If she chooses a different angle, the results of her measurement are only predictable statistically.

Unfortunately, even if Bob and Alice decide ahead of time to make only 0º and 90º measurements, they cannot send information back and forth instantaneously. The local measurements made by Alice and Bob look completely random. Photons emerge with 50/50 half of the time out the V port , and the other half of the time out of the H port. It is only when they meet to compare their results that meaning emerges from their measurements. This is not to say that no information is sent, only that the information cannot be recovered unless they meet. Alternatively, they may send auxiliary information to each other using conventional means (that travel at or below the speed of light). In fact, by sending just two additional (classical) bits of information that describe

the results of their quantum measurements, it is possible to transport whole quantum states from one location to another. This is called quantum teleportation.

QUANTUM TELEPORTATION

"Beam me up, Scotty," is an echo of pop culture that has reverberated since the Star Trek TV series first aired in the mid 60's. Captain Kirk of the Starship Enterprise is requesting Scotty, his chief engineer, to teleport him out of danger from the surface of some planet where he may have too boldly gone where no man had gone before. On the set of the TV show it was cheaper to "beam" a body to and fro with the low-budget of the original episodes than to have to film expensive landing and launch scenes of shuttle craft. But the transporter has become etched in popular culture, and remains one of the lasting icons of science fiction. The question is: What fundamental laws of physics does teleportation violate?

Maybe none, if the teleporting speed is slower than the speed of light. The aspect that makes a teleporter look so far-fetched is the scale of the task (and issues of scale are usually issues of technology rather than fundamental problems. If given enough time, clever engineers can often tackle scale as long as the fundamental physics is allowed. Sending a man to the Moon was a project of immense scale that surely must have seemed like science fiction to writers only a single century ago. With many centuries ahead of us (let us hope), perhaps the scale of teleportation will be surmounted.

Nonetheless, the scale of the problem is daunting because the human body contains something around 10^{28} atoms and nuclei, and about fifteen times that many electrons. These would all need to be transported to maintain the complete being. There is furthermore the question of the quantum states of all those particles. Would it be enough to transport the physical electrons and nuclei and place them in identical

locations, or would the exact quantum states of the particles need to be preserved in order to preserve the intangible essence of the human soul? This is a point that is hotly argued. Some say that as long as all the neural synapses are identically configured, it would not matter whether the exact quantum states were reproduced. Others argue that consciousness is a fundamentally quantum phenomenon that would be destroyed if the quantum states were scrambled during the teleportation.

If the quantum states do matter, there is a fundamental hurdle that must be overcome to measure those quantum states and transmit the quantum information to the destination. Quantum measurement is a violent act because it destroys delicate quantum superpositions. It also destroys quantum information because it projects an unknown state, which is in a superposition of states, into only *one* of those states. The "presence" of those other states in the superposition is lost forever to that quantum particle. Quantum measurement is such a disruption of quantum information that theorists were able to prove a quantum non-cloning theorem. This theorem states that it is fundamentally impossible to clone a quantum state because the act of quantum measurement would disrupt the original. This law would seem to place teleportation forever out of the reach of reality.

But there is a small loop-hole in the law that is just big enough to let teleportation wiggle through. The non-cloning law forbids the cloning of a particle without disrupting the original. But if the original is discarded, the law says nothing about the ability to recreate the original at the same or even a different location, leaving that possibility open.

Quantum teleportation is still faced with a conundrum. Quantum measurement of an isolated particle destroys quantum superpositions and hence destroys quantum information. The direct task of measuring the quantum state of even a single particle and reconstructing that exact state is impossible, because the act of measurement only projects out one state of the many-state superposition. Therefore, even though the quantum non-cloning law would seem to allow the possibility of quantum teleportation,

the process cannot be done by direct quantum measurement. Some alternate approach must be found.

That alternate was proposed in 1993 by Charles Bennett of IBM and Gilles Brassard of the Université de Montreal with their collaborators [NOTE: Bennett and Brassard Ref]. They showed that Alice could start with the unknown quantum state that is to be teleported, and then use an entangled pair of particles as a quantum resource. She takes one of the entangled pair, and the other is sent to Bob. Alice makes a quantum measurement of a joint property belonging to both her entangled particle and her unknown quantum state. By doing this joint measurement, the other particle of the entangled pair would assume some of the quantum properties of the original unknown quantum state. Then Alice sends two bits of information through a classical channel to Bob, telling him how to rotate his entangled particle to reconstruct the original unknown state. The beauty of this approach is that the quantum state remains unknown to both Bob and Alice, even after teleportation. Therefore, if it had been in a delicate superposition of states before teleportation, it remains in that superposition after the teleportation. Also, the process of teleportation destroys the original state when the joint properties are measured with the entangled state, thereby obeying the non-cloning theorem.

The key to quantum teleportation is the ability of Alice to perform a measurement that provides Bob with enough information to recreate the original quantum state ⎛ but without having Alice actually measure the individual properties of the unknown state. This sleight of hand is performed through a process known as a Bell State Measurement (BSM), named after John Bell. This is a quantum measurement of an unorthodox kind where the joint properties of two particles are measured relative to each other, but no direct measurement of the individual properties of each particle is needed. The fundamental goal in the BSM is to project the properties of the unknown state onto four possible Bell States. These Bell States have the important feature that the unknown

quantum state can be completely described through a linear combination of only these four states. The BSM is the measurement process that projects the unknown state onto one of these four. Even though the measurement process has caused a collapse of the joint properties of the unknown state with the entangled particle (and by the properties of entanglement Bob's particle collapses at the same time), knowing this one Bell State does not tell Alice anything about the actual individual properties of the unknown state. On the other hand this information is all that Bob needs to know to perform the rotation on his entangled particle to get the original unknown quantum state.

The schematic arrangement for quantum teleportation is shown in Fig. 9.8. Particle 1 is the unknown state that is to be teleported to Bob. Alice and Bob share an entangled pair of photons; Alice has Particle 2 and Bob has Particle 3. Alice performs a Bell State Measurement on the joint properties of Particles 1 and 2, projecting her unknown quantum state of Particle 1 onto the entangled Particle 2. At the instant of the BSM, Bob's Particle 3 collapses into the same joint state as Particle 2 and 1. But Particle 3 is not yet in the exact state as Particle 1. To put Particle 3 into the state of Particle 1 Bob has to make one of four possible rotations on his particle. Which rotation to make depends on the results of Alice's BSM. Since there are four Bell states, Alice needs to send two bits of information classically to Bob ($2^2 = 4$). When Bob receives which Bell state Alice observed, he then knows which of the four different rotations to perform on his particle. Once he performs the rotation, his Particle 3 is identical to Particle 1 in the unknown quantum state.

After the teleportation, neither Bob nor Alice know what the unknown state is. Both the BSM and Bob's rotation provide them with no information about the state of the particle. Yet by the laws of quantum mechanics and entangled states, Alice and Bob can be certain that the particle has been successfully teleported. Because Bob needs to know which rotation to perform, and he only gets this information from Alice through a conventional communication channel, quantum teleportation cannot occur faster than the

speed of light. Even though the wavefunction collapse of Bob's particle is instantaneous with Alice's BSM, no information is sent until Alice and Bob communicate through classical means. Quantum teleportation therefore satisfies relativity and hence causality, and it also satisfies the non-cloning theorem (hence violates no known physical laws.

The first quantum teleportation experiment was performed in 1997 in Vienna using nonlinear optical crystals to generate entangled pairs of photons and using simple beamsplitters and photon detectors to perform the Bell State Measurements [NOTE: First teleportation Ref]. In this experiment, the quantum state could be teleported correctly (and verified) only one time out of four. But it was a start. The challenge facing teleportation experiments is the same challenge of the Star Trek transporter: one of scale. Teleportation has been accomplished in the laboratory using only one or a few quantum states. Pushing the number of teleported states, and the distances over which they are being teleported, is a severe challenge. Going from one (or a few) teleported states to teleporting 10^{30} quantum states of the human body may be beyond reach. The data rate for such teleportation, even if it took an entire century to transmit all the quantum information of a single human body, would still be a data rate in excess of 10^{20} bps. Comparing this data rate to the simple *classical* rate of 10^{12} bps that we are struggling with today, tells us it would take about the age of the Universe to teleport a single human. Even with incredible improvements in data rates, teleportation of people does not look promising.

On the other hand, setting our sights on teleporting a human is probably not the best use of the technology. Quantum information contained in small systems of a few particles has potential that goes far beyond classical information. A small ensemble of quantum particles can be in a superposition of hundreds or thousands of quantum states all at the same time. Transporting these states using quantum teleportation therefore becomes an important resource, especially for a quantum computer. Teleportation can become the data bus that ports quantum information from the output of a quantum logic

gate to a quantum memory device where the quantum information is stored until it is needed by another logical operation.

Aside from quantum logic gates there is a much more immediate need for quantum information transmission, especially if the information needs to be unassailably secure, free from any hint of an eavesdropper. Quantum effects guarantee absolute channel security through the simple fact that an eavesdropper must make quantum measurements to extract information, and the act of measurement fundamentally disturbs the information content, just as the EPR Demon disturbed the entangled states. The presence of the eavesdropper can therefore be uncovered through simple measurements of the photon statistics, and the channel can be abandoned before any important information is sent.

QUANTUM CLOAK AND DAGGER

Every time you make a purchase over the internet with your credit card, the pertinent information is scrambled using an encryption scheme that multiplies two large prime numbers together. Multiplying large numbers together is easy, but it is very difficult to factor them apart again. For instance, see how long it takes you to find the two prime factors of the number N = 152,399,021 [NOTE: 152,399,021 = 12343*12347]. This number takes 27 bits to describe in binary notation, and it takes my old computer (Motorola 86040 processor) about 15 seconds to factor using a simple sequential search algorithm [NOTE: Ref Conway pg. 464]. But the problem is that the time to factor a number increases exponentially as $2^{B/2}$, where B is the number of bits needed to express the number. A number with 40 bits (the limit imposed by law on all foreign export versions of web browsers until 1996) would take my machine about 20 minutes to factor. But 128 bits would take my machine about 50 billion years (the Universe is only about 20

billion years old). Of course, much faster computers and much more efficient algorithms are available. As we will see shortly, 429-bit keys have already been factored, although the technology that is needed to do this is hard to come by.

Therefore, forcing potential eavesdroppers to factor large products of primes is an excellent way to ensure privacy and is the basis of an encryption scheme called RSA (named after Rivest, Shamir and Adleman who invented the scheme in 1977 [NOTE: RSA Ref]) that is used almost universally for the transmission of electronic data. With this scheme, the person who wishes to receive a message, let it be Alice, publishes two public numbers. One is the product of two large prime numbers and the other is any number of choice. Using these public numbers, Bob constructs a message that he sends publicly back to Alice. Because the encryption key is completely public, as is the subsequent coded message, this scheme is called public key cryptography. Yet the encoded message can only be broken by someone who can succeed in factoring the large key into its prime factors.

As an example of the difficulty factoring large numbers that are the products of primes, Martin Gardner, writing for *Scientific American*, published the 129-digit number

N=114381625757888867669235779976146612010218296721242362562561842935706935245733897830597123563958705058989075147599290026879543541.00

plus an addition number M = 9007, and a message encrypted by the original RSA team using these numbers. A cash award of $100 dollars was promised to anyone who could crack the code. This was known as the challenge of RSA-129. A 129-digit number can be represented by 429 bits, and 512-bit encryption was (and still is) commonly being used in commercial RSA schemes.

A decade passed before the mathematical and computational tools were available to crack RSA-129, but it finally fell to a sophisticated attack mounted by researchers at

the Bellcore research labs in 1994. They mustered a coordinated effort that used 1600 separate workstation platforms distributed internationally. They succeeded in deciphering the message: THE MAGIC WORDS ARE SQUEAMISH OSIFRAGE. Today, even 512-bit encryption is susceptible to such concentrated attack, which has raised the level of suggested security to 768-bit keys for personal security, 1024-bits for corporate security and 2048-bits for ultimate security [NOTE: Ref: Brown pg. 170]. Even with the powerful mathematical tools in use today, it would take a time larger than the age of the universe to factor the 2048-bit encryption. Yet even these numbers or greater can fail as advances are made in mathematical techniques in number theory. The fundamental problem is that the public key is always susceptible to attack.

On the other hand, quantum cryptography provides a means of sending information that is impervious to eavesdropping. What is needed is a quantum channel, for instance a fiber carrying single photons, between Alice and Bob. A third person, conventionally named Eve (a play on the word "eavesdrop"), is the suspected eavesdropper. How can quantum effects, especially quantum entanglement, be used to guarantee the security of the communications between Alice and Bob and keep Eve in the dark?

In cryptography by entanglement Bob and Alice receive entangled photon pairs from a central source. They each perform a long random sequence of polarization measurements along three different directions that they agreed upon publicly in advance. Each makes measurements that are completely random and completely independent of each other. The outcome of each measurement produces a photon in half the cases, just as in the case of the EPR experiment. Afterwards, Bob and Alice publicly send each other the polarization directions they chose for each measurement. They identify for which cases they had each used *different* measurement directions, and they then publicly send the results of only those measurements. If Eve (the EPR Demon) is eavesdropping, then the quantum correlations will be perturbed. In that case, Bob and Alice abandon the

channel. On the other hand, if the correlations are correct, then they conclude that Eve is not present. In that case they use their remaining data, obtained when they had chosen the same directions, as a random encryption key. Because of entanglement, they each have exactly the same random key. They use this to encrypt a message that they send over a completely classical channel.

This approach to quantum cryptography is virtually immune to attack. The quantum correlations in the EPR experiment are highly sensitive to any attempt to eavesdrop. Furthermore, once Alice and Bob have their random key, it is virtually impossible for the public encrypted message to be decoded because the encoded message has perfectly random statistics based on their random measurements. There is no handle for a code-breaker to grab onto.

Practical implementation of cryptography by entanglement is the closest of all the quantum information technologies to becoming a "real" enterprise. An experiment conducted in Geneva, Switzerland in 1997 succeeded in sending entangled photons 10 km over a fiber without losing quantum correlations [NOTE: entanglement over 10 km Ref]. More recent demonstrations have succeeded in sending quantum information over conventional fibers installed for local-area networks, and also through several kilometers in air [NOTE: Los Alamos LAN Refs]. In addition, the dense part of the atmosphere near the Earth's surface is only about 10 km thick, which means that quantum communication with satellites is a clear possibility. These recent advances point to the feasibility of quantum communication and cryptography as real-world applications of quantum information. Given the growing importance of information security in a world that is progressively operating on-line, quantum cryptography is poised to become the first commercial quantum technology.

But the potential of quantum computing is closely tied to quantum cryptography (for instance the parallel quantum information contained in superpositions of quantum states can be used to perform calculations that are intractable on any conceivable classical

computer. One important problem like this is prime factorization. Quantum parallelism rises exponentially (like the problem of prime factorization) with the number of quantum states, providing an enormously parallel resource. This potential is so vast, and the threat to RSA so great, that the field of quantum computing has become one of the fastest growing fields of science and technology. Quantum computing operates on units of quantum information called qubits.

QUBITS

A bit can be constructed from any system that has two states. A light switch can be on or off. A door can be open or closed. For a transistor in an electronic logic gate, the control voltage can be passing current or blocking it. These are all examples of classical bits. Even though light switches can have dimmer knobs, doors can be partly open and transistors can pass continuous values of current, these possibilities are disallowed explicitly for the expression of binary information. In the decision tree of a binary search (Fig. 5.2 of chapter 5), the choice is purely binary. The hidden item is found either in the right-hand branch or the left-hand branch. The number of branches is equal to the number of bits needed to specify the information content of the hidden value. No fuzzy answers are allowed [NOTE: Fuzzy logic is an alternative logic system that allows mixtures of answers.].

Quantum information has similar constraints that are subtly different. A quantum bit, called a qubit, is a quantum entity that has two orthogonal states. For instance, in a two-level atom, the electron can be in the upper state or in the lower state. For a photon, the polarization can be H or V. So far, this sounds like the classical case. But now we can take one step further by considering a 45º photon. This photon is a linear combination of equal amounts of H and V. It is tempting to think of this case as a door

half open, but this is where the difference between the classical world and the quantum world is crucial. For the half-open door, it cannot be viewed as a door that is open and closed *at the same time*. For the 45º photon, on the other hand, it is both H and V at the same time, just as a photon passes through *both* slits in Young's double slit experiment in Fig. 9.1.

The qubit exists in both states at the same time. Even though a photon may be observed in one state or the other, this is merely the projection of the state onto a specific axis, as in Fig. 9.4. The key property of a qubit is that we can always choose a specific direction to rotate our detector that guarantees we will observe it along that direction. Even though it may be a linear combination of the two possible states for one choice of the coordinates, another set of coordinates can be found for which the photon is in a pure state. This means that the qubit is a single well-defined quantity. Even though we cannot predict which port of the calcite crystal it will emerge from *in general*, we can always find some rotation of the axes that will guarantee only one port will emit the photon.

Therefore, when we say that a qubit holds the answers of 'yes' and 'no' at the same time, we are saying something that must be interpreted very specifically. The qubit is in a coherent superposition of 'yes' and 'no', just as waves can be in coherent superpositions as we talked about in Chapter 2 and again in Chapters 6 and 7. The coherence of the superposition allows us to find a rotation that makes the answer a pure 'yes' or 'no'. The coherence of quantum states is what makes the pair correlation function of Fig. 9.6 go to zero where the classical result does not. It is quantum coherence that allows us to see a bomb in the dark. And it is quantum coherence that allows us to use qubits to perform massively parallel computations on quantum information.

The power of qubits arises not from a single qubit, but from collections of qubits. For instance, we can contrast the information content of 2 classical bits with 2 qubits. For a pair of classical bits, there are four possible arrangements of the bits: 00, 01, 10 and 11. However, there is only *one* possible arrangement *at a time*. Therefore, to enumerate

all four possible arrangements, we need to step through the bits four times to produce them all. On the other hand, for 2 qubits, all four combinations exist at the same time. To enumerate all possibilities, we only need to produce a single quantum superposition: $\Psi = a\psi_{00} + b\psi_{01} + c\psi_{10} + d\psi_{11}$. It is important to keep in mind that, because this is a coherent superposition, Ψ is a single entity. To prove this, all we need to do is rotate our axes by the appropriate amount that would make Ψ a pure state. Since Ψ is a single entity, we only need to express it, or produce it, once. Yet it contains all the possible information of two bits. This represents a 4-to-1 savings in effort.

The value of qubits becomes obvious once we start to increase the number of qubits in our collection. If we have N qubits, we can describe 2^N different configurations *all at the same time*. But a classical system would need to enumerate those configurations one at a time. If N = 100, there are $2^{100} = 10^{30}$ distinct configurations. A classical machine would need to define all 10^{30} distinct configurations, while a quantum machine could do it by defining only 100 qubits. The exponential increase of 2^N configurations relative to N qubits becomes a resource of tremendous (literally astronomical (magnitude. Quantum memory systems of only a modest number of qubits have a potential for memory storage that makes our newest and largest classical RAM chips look infinitesimally insignificant. For this reason alone quantum information sciences have received considerable recent attention.

Quantum logic gates and quantum computers benefit fundamentally from the unique parallelism of quantum superpositions. A single input to a quantum logic gate contains a superposition of all possible input states. The output of the gate operating just once on this input contains all possible answers. To peek inside a quantum logic gate, I begin with a classical example to describe one of the basic logic gates called a controlled-NOT gate. However, as we shall see, classical versions only go so far before we need to turn to the unique features of the quantum world.

COMPUTING WITH SPINNING COINS

Consider a classical coin for which "tails" represents "0", and "heads represents "1". When this coin is flipped and lands, it will end in the "0" state or the "1" state. In other words, the coin is a classical bit. Its value, furthermore, is determined randomly, depending on the outcome of the coin toss. Once the coin lands, it is in a definite state of either "0" or "1". But there are times when this classical coin is both "1" and "0" at the same time, such as the time it is in the air, or if it happens to land momentarily on its side. At the instant it lands it can tip either way, but which way would be impossible to predict. You could say that there is a 50% probability that it will fall tails, and an equal probability that it will fall heads.

Yet it is possible to cheat the odds. For instance, if a small additional mass is added to one side, this side will have a slightly higher chance to land down, while the other side has a slightly higher chance to land up. In this case, perhaps there is a 48% chance to land tails and a 52% chance to land heads. With enough weighting the odds could be pushed far from 50/50 (perhaps as far as 90/10. Even then it is impossible to say with certainty what the coin will do. There is still that 10% chance to land tails.

When such a weighted coin is flipped, it is in an indeterminate state while it is in the air. Only the probability of how it will land can be stated. But this changes as soon as you observe the state of the coin. You do this by grabbing the coin in mid-flight and smacking it on your forearm. The instant you observe the coin you know exactly what state it is in. With 100% probability you will see either a head or a tail. Furthermore, after the observation the coin will remain in this completely defined state until it is flipped once again into the air.

Now let's think how we might try to do logic with this weighted coin. The simplest type of logic is conditional logic (if the coin lands heads up then you do A,

otherwise you do B. The action A may be as simple as turning a second coin over, while the action B would be to do nothing. What you have in this case is a two-coin logic gate. There is a control coin that is flipped, and a second coin that is operated on. If the control coin lands heads up, you flip the second coin over, otherwise you do nothing.

This logic gate is known as a controlled NOT , also known as C-NOT. The NOT operation by itself is simply the turning of the coin over: NOT HEADS = TAILS and NOT TAILS = HEADS. In our example, whether we apply the NOT operation or not to the data coin is conditional on the value of the control coin. Logic gates are visualized in terms of lines and nodes and are drawn as a diagram, such as the C-NOT gate in Fig. 9.9 The control bit passes through on its line unaffected, but it connects to the data line where it causes the data bit to switch if the control bit is "1". This is the meaning of the circle with the "x" in it. The logic table, known as a truth table, for the C-NOT is shown in the figure.

Let's consider how we might implement the C-NOT logic gate using real coins. Flipping coins is not the best way to do this. When they are in the air they have the nasty tendency of falling and hitting something. And if we tried to maintain the indeterminacy by placing the coins on their sides, this will last for only a short time before they tip over.

On the other hand we can spin the coin on a flat table. Or if we are concerned with it wandering while it spins we could place it in a slightly curved bowl. Furthermore, to allow the coin to spin as long as possible, we could make the coin and the bowl out of frictionless material, and we could place the whole thing inside a bell jar and evacuate the air. Under these conditions the coin could continue spinning for a long time (all that time maintaining its indeterminate state. Remember too that the coin can be weighted, giving the control coin uneven probability to land heads up or down. In spite of the weighted coin the truth table for the C-NOT does not change. When the indeterminacy of the control coin is removed by the act of observation it will be either a head or a tail. No

other possibility exists. Then the data coin will be turned over or not. The odds of which action is taken depends on the weight on the coin.

It is hard to see how this coin-operated logic switch could do anything useful. But it is not so far-fetched. For instance it is possible that the weight on the control coin was placed there by some earlier logical operation. And when this coin stops spinning let's say some weight is added to some later control coin in some subsequent coin-operated logic gate. As these gates are cascaded, complex logic operations can be implemented that go far beyond the capabilities of the single C-NOT gate.

It is important to make the distinction between a single realization compared with a collection of realizations. For instance, let's take a weighted coin that has only a 10% chance to land heads up. The probability for this to happen are low, but it is still one of the allowed outcomes for a single realization. However, if we repeat the experiment 100 times then even though all possibilities will occur, the cases when the control lands heads up will be far fewer than when it lands heads down. Therefore, for a weighted control coin in this classical logic gate we would need to perform the logic operation many times to get a clear measurement of the weight on the control coin.

The biggest problem is that each calculation by the network of gates is only a single realization. Because this computer operates probabilistically, we would need to run it over and over again. However, the drawback of this computer is not that it is probabilistic, but that it is classical. Every single realization is distinct from any earlier or later realization. Each calculation gives only a single answer. The final coin will be either heads up or tails up. There are no other classical possibilities.

Now let's make a modification in how we operate this computer. In the way we first described it, it was necessary at each stage to observe the spinning coin. In a network of such gates the result of all the previous gates must be made determinant (heads or tails) before taking action on the data coin. But what if we could relax this requirement? What if the output state of the data coin is also indeterminate? In other

words, the control coin would operate on the data coin without it ever being made to stop spinning. Since the control coin state is indeterminate, the data coin state is indeterminate as well, but in a special way (the data coin is perfectly correlated with the control coin.

This is where we need to leave our classical coin-operated computer and enter the quantum domain. There is no conceivable classical operation that can do what I just described (to have one indeterminate state operate with a predefined set of rules on another indeterminate state. Qubits, on the other hand, are perfectly happy to operate this way. Qubits, and especially their superpositions, are therefore at the heart of quantum logic gates.

QUANTUM LOGIC

The quantum Controlled-NOT, or C-NOT, gate has the same circuit diagram as the classical diagram shown in Fig. 9.9. However, its behavior goes beyond classical capabilities, and is one element out of which universal quantum computers can be constructed. The C-NOT gate has two important features that make it fundamental: it is conditional, and it is reversible.

Being conditional means that the qubits on the two lines *interact*, i. e., what comes out of the data line depends on what went into the control line. This interaction among qubits is exactly what causes quantum entanglement of the EPR-type. To become correlated, two particles need to interact with each other. The interaction can be the process of creation, as when the positronium decays and creates the two entangled photons. Or two particles that are already in existence can interact with each other. We saw in chapter 6 that photons interact with each other through intermediate electrons, such as electrons on atoms. It is therefore conceivable that a quantum C-NOT gate could

be constructed using the quantum states of a single atom to couple the quantum states of two photons.

Being reversible means that the input information of the quantum gate can be reconstructed based on a knowledge of the output. Reversibility was shown by Rolf Landauer of IBM in the early 1960s to be a necessary requirement for dissipationless computation [NOTE: Landauer reversible computation Ref]. The importance of removing dissipation from classical computers is obvious. For instance the heat caused by the increasing density of transistors on microprocessor chips is one of the principal obstacles to achieving even higher densities. This is because transistor logic uses voltages and currents that produce heat, just like a resistive heat pad or a thermal electric blanket. What Landauer showed was that dissipation of energy during computation was only necessary if information is destroyed during the computation.

An AND gate is an example of irreversible logic. It has only one output for two inputs. The output is equal to 1 if and only if both inputs are equal to one. The output is zero otherwise. But information is lost here. If the output is equal to zero, that could be because either Line #1 was zero or Line #2 was zero or both. Knowing the output therefore does nothing to enable us to reconstruct the input, indicating that one bit of information was destroyed by this logic operation, causing the emission of a minute amount of heat in the process. The C-NOT, on the other hand, is completely reversible because the input states can always be reconstructed just by knowing the output states. Therefore, in principle, a C-NOT gate could be constructed that dissipated no energy and hence produced no heat.

Reversibility in a classical logic gate is hard to achieve (although not impossible) because large numbers of electrons need to be transported from one location to another in an electronic device. If superconducting wires are used that have no resistance, it is possible to construct a reversible logic gate that produces no heat, although the engineering involved is highly challenging. On the other hand, reversibility is completely

natural for quantum systems. Reversibility in a quantum system is equivalent to rotating coordinate axes, such as when observing photon polarizations with calcite crystals. Quantum logic gates like the C-NOT therefore automatically satisfy the requirements of reversible computation to be performed without dissipation of heat. This feature of quantum logic makes it a candidate for the ultra-small scales and high densities that will be needed in computers of the future.

In addition to being conditional and reversible, which can also be satisfied by classical logic gates, quantum C-NOT gates go beyond classical capabilities when coherent superpositions of states are used at the inputs. For instance, the control state can be in a superposition of both 1 and 0 in the coherent state $\Psi_c = \psi_0 + \psi_1$. The output is then also a coherent state. If the input data state is 0, then the output of the C-NOT is $\Psi_{out} = \psi_{00} + \psi_{11}$. Notice that there are two subscripts for the output functions. The first subscript stands for the control-out value, which is just the control-in value. The second subscript stands for the data-out value. Since the data input was 0, it remains 0 when the control is 0, and it flips to 1 when the control is 1. The quantum C-NOT has therefore performed *two* controlled-not calculations in a single step. The output $\Psi_{out} = \psi_{00} + \psi_{11}$ is also an entangled state between the control line and the signal line. It does not matter that the control bit goes through unchanged. The interaction between the qubits correlates their values. A C-NOT gate is therefore a source of entangled pairs of photons that can be used in quantum teleportation.

More complicated C-NOT gates are obtained in a natural manner by allowing the control line and data line to accept collections of qubits. For instance if the control line has N qubits, and the data line has M qubits, then NxM combinations of qubits are produced by the operation of the gate. Because these are qubits, and not classical bits, this means that 2^{NxM} combinations of calculations are performed all at once. The values of N and M do not need to be very large before unimaginably immense calculations are performed in a single step by the gate.

The C-NOT gate is as complex a gate as is needed to construct a universal quantum computer. In 1994 it was realized that by using only a 2-qubit C-NOT gate in combination with a 1-qubit gate that performed a simple rotation, that any quantum computation could be performed. Two simple quantum logic gates are therefore all that are required to build a universal quantum computer, much as the AND and XOR gates are sufficient to build a universal classical computer.

QUANTUM COMPUTING

In the description of quantum logic gates, there has been a nagging problem that has remained unspoken. In the examples, great emphasis has been placed on the ability of the logic gates to calculate all answers at the same time for a single operation of the gate. But there is a problem: how do we read out all those answers? By now, you have gained enough of an understanding of quantum systems to know that a measurement on any coherent superposition of states produces only a single answer. It does not matter whether the final quantum superposition contains 2 or 10^{30} answers. When the superposition is measured, only one answer is projected out. To be able to see all 2^N answers in the gate output, we would need to make measurements on at least 2^N identical systems. But this is exactly the number of operations we would need to perform on a classical computer to get the same number of answers!

This is a serious problem. It is serious enough to make the potential of quantum computing look like a smoke and mirrors. If we can access only one answer from the quantum superposition at a time, what have we gained? What is the value of the vast parallelism of quantum computing if it evaporates as soon as we try to observe it? The value looks completely metaphysical, like the sound of the tree falling in the woods when someone

is not there to hear it. Does the information really exist if we cannot read it? And if we cannot read it, who cares whether the information is there or not?

This was the state of quantum computing around 1985. The potential of quantum computing was tantalizing, but it looked beyond grasp, and hence beyond usefulness. Quantum computing was an interesting exercise practiced by esoteric theoreticians to answer metaphysical questions. Then one of the esoteric theoreticians by the name of David Deutsch working at Oxford University found a way to use all the answers at once before destroying them. The key was not to try to make the quantum computer simply answer questions in parallel, but rather to get all the answers to interfere with one another to produce a single collective answer.

This is something like quantum seeing in the dark, but at a much larger scale. The two answers are that Path #1 has no bomb and Path #2 has no bomb. The coherent superposition of these two answers produces complete destructive interference at the detector (meaning "no bomb". But when a path *does* contain the bomb, the removal of the interference changes the collective answer to "bomb" (at least some of the time). It is this ability to utilize the coherent interference among all answers to produce a collective single answer that speeds up certain problems that are intractable by classical means. On top of this, Deutsch was able to show that quantum computers could be universal computers, in the sense of quantum Turing Machines.

This was an astounding breakthrough in the field, although his publication in 1985 [NOTE: Deutsch Ref] was largely overlooked (initially. Part of the continuing problem was finding the right kind of problems where the quantum parallelism, combined with interference, could be used to produce a single result. In other words, what quantum computing needed to bring it into the big time was a killer application, or "killer app", that was too important to ignore.

That "killer app" was provided in 1994 by a reclusive genius named Peter Shor working at Bell Laboratories in New Jersey. The problem he decided to tackle was the

problem of finding the factors of the product of two large prime numbers. If such an application could be achieved, then the RSA cryptographic security systems could be broken. This would be catastrophic to world commerce and political stability. A quantum computer, using quantum parallelism and interference, could solve the problem. Hence, here was a "killer app" that no one could afford to ignore. What did Shor think up?

Shor realized that a central part of the problem of factorization involved finding repetitive patterns in sequences of numbers. Any time that repetitive patterns are found in a signal, they can be analyzed in terms of waves. And waves interfere. This brought Deutsch's concept of quantum interference into the problem of prime factorization. In a feat of intellectual brawn, Shor was able to apply the aspects of quantum parallelism and interference to the problem. In the end, he was able to define a specific quantum algorithm that would be able to break all codes currently in use, and any codes likely to be used in the future. In a single stroke of genius he had toppled the entire world of RSA cryptography (almost.

The problem with quantum algorithms is that they need quantum computers to execute them. And they don't exist yet. Therefore, for the moment, cryptography and privacy are secure as long as quantum computers remain in the future. But the future has the nasty habit of becoming the present. Quantum computers may still be a long way off, but quantum logic gates have already appeared in selected laboratories around the world, making the first primitive steps towards the third generation of the Machines of Light.

THE THIRD GENERATION

There comes a time when theoretical speculation must meet reality and be reduced to practice. Qubits need to be more like flesh and blood than spirit of thought. They need a physical existence that can be touched and felt (if not by human hands, then by surrogate

means. Not only must qubits be created and supported by physical systems, they must also be protected from outside disturbances that destroy the fragile quantum superpositions. Furthermore, physical operations must be devised to make the qubits interact to entangle themselves. What type of laboratory systems allow qubits to control other qubits? What would these machines be like? More importantly for our interests, how does light play a role in these prototype quantum logic gates?

Our discussions in Chapter 7 of the control of light by light used classical Mach-Zender interferometers to allow one light beam to modify another. The interferometer allows light beams to interact using nonlinear materials to mediate the interaction. If we let the light beams get progressively weaker, the effects of single photons become more pronounced. In the quantum limit there may be only a few photons interacting at a time in the interferometer. To what extent does the interferometer become a quantum-optical logic gate?

This question was first asked in 1989 by Gerald Milburn, an Australian theoretical physicist working in quantum optics at the University of Queensland [NOTE: Milburn PRL]. He considered just such a case, as in Fig. 7.2b, where a control beam swaps the output ports if the imparted phase shift Φ can be made as large as π (180^o). The device was called a "quantum optical Fredkin gate", named after Ed Fredkin of MIT, who was one of the first computer physicists to recognize the importance of reversible gates for computing. A Fredkin gate has three inputs and three outputs. One of the inputs is the control line and carries its bits through unaffected, just as in the C-NOT gate. The information on the other two gates is unaltered if the control is 0, but is swapped if the control is 1. In our discussion of optical control in Fig. 7.2b, we neglected a second possible input port. This is the other input into the first beamsplitter. A beam can impinge on the upper port just as well as in the side port. With two beams entering, the beam in the side port exits Port #1, the beam in the upper port exits Port #2 in the absence of a phase shift Φ. On the other hand, if the phase shift is 180^o, then the two beams swap their outputs. This precisely describes a classical

Fredkin gate. What ⌠ Milburn wanted to know ⌠ happens in the quantum regime? Can this operate as a quantum optical Fredkin gate for quantum computing?

The idealized quantum optical Fredkin gate could indeed operate as a quantum logic gate, under the right conditions. For instance, if the important switching condition $\Phi = \pi$ could be achieved under the action of a single control photon, then a signal photon could be switched from one port to the other. Of course, there would be technological hurdles. For instance, the interferometer could have no losses, otherwise a photon may fail to appear at the output. Furthermore, any fluctuation in the intensity of the control beam would produce fluctuations in the phase Φ, which would occasionally, erroneously, cause the photon to exit the wrong port. These obstacles did not originally appear insurmountable ⌠ until further analysis highlighted a fundamental quantum limit on the switching of such a quantum optical gate.

Such a beautiful idea meets ugly reality in the need to achieve the appropriate phase shift $\Phi = \pi$ to perform error-free switching of single photons. The problem arises in the nonlinear interaction of single control and signal photons. It is impossible to get arbitrarily large interactions between them. From classical physics, we saw that weak nonlinear interactions can always be increased by increasing the duration of the light beam interaction. This is most easily achieved in fibers where the beams propagate together over long distances. One might think that the same condition would hold true for quantum interactions between two photons. It does not. It was Milburn himself, working with colleagues, who showed that the nonlinear interaction between photons is independent of the length of interaction. All that matters is the fundamental strength of the nonlinear susceptibility that links the two photons. This would need to be inordinately large to allow the photons to produce large phase shifts. Even more troublesome in the theoretical analysis was the prediction that even under the most ideal conditions and the strongest possible susceptibilities, that it would take at least π control photons to generate a π phase shift in the quantum interferometer. This means that a single photon would generate a phase shift of

only $1/\pi$, which is a phase shift of 60^o rather than the required 180^o. This is the best that the quantum interferometer can do (even in an ideal world, which is never the case in practice.

Here, as so often throughout this book, we see that the control of light by light is a daunting challenge. Photons to not like to interact with each other. Even when we coerce classical light beams into a nonlinear medium, we require long interaction lengths to get large effects. Now we see that in the quantum limit even this ploy fails, and single-photon control of a quantum gate looks fundamentally impossible. Without the ability to operate a quantum-optical logic gate, has the future of optics in quantum computing seems to have a dimmed?

A timely theoretical discovery changed this situation in 1995 when Seth Loyd of the Information Sciences Department at MIT was able to show that almost any quantum logic gate is universal, i. e., that almost any quantum gate with two or more inputs is computationally universal, making it possible to produce any desired quantum logic circuit [NOTE: Loyd, PRL 1995]. The question then was whether Milburn's faulty Fredkin gate can be considered "almost any quantum gate"?

Actually, the quantum optical Fredkin gate is a step too complicated to be useful. It turns out that a much simpler gate would satisfy Loyd's scheme that simply crossed a control beam with a signal beam to induce a mutual phase shift Φ, just as the control and signal photons were doing in the nonlinear crystal in the Fredkin gate. The good news was that Loyd's scheme did not require a full π phase shift in the operation. Almost any phase shift would do the trick, as long as it was robust and sizable relative to measurement errors. This type of quantum gate is called a conditional phase shifter. The phase is shifted only if both photons have the appropriate polarizations, and is not shifted otherwise.

The same year, in one of the first experimental demonstrations of quantum logic, a group at Cal Tech under the direction of H. J. Kimble constructed a conditional phase shifter and measured sizable phase shifts that were conditional on the polarizations of the photons [NOTE: Kimble, 1995]. Furthermore, they operated their logic gate using only single

photons. Their apparatus used an atomic beam of cesium atoms as the "nonlinear medium". These atoms passed through a very small optical cavity that was much like a miniature laser cavity only 56 microns long. Two laser beams with slightly different frequencies carried the qubit information in their respective polarizations and were transmitted through the cavity mirrors. The cavity allowed the photons to bounce back and forth many times, letting them interact with the atomic states. The gate was read out by simply measuring the polarizations of the photons that emerged from the cavity.

The performance of the conditional phase shifter was moderate but sufficient to cause excitement as a real-world demonstration of quantum conditional dynamics. They were able to observe a phase shift of 16º per photon. The speed of the operation was also reasonable, clicking along at a rate of 75 MHz entangling the state of one channel with the other. However, the experiment also highlighted continuing difficulties with quantum optics as the sole conveyors and processors of quantum information. Kimble described the qubits in the experiment as "flying qubits" because the quantum information encoded on the photons went flying through the apparatus at the speed of light. There was no form of quantum information storage in this configuration, meaning that the entangled photons emerging from the apparatus needed to be used immediately for the next stage of a large quantum computation. Also, the small phase shift of 16º was not as big as one might want. Even though larger phase shifts up to 60º would be possible under ideal conditions, the microcavities that could produce this phase shift would slow the interaction down because the photons would need to spend more time in the cavities before being emitted.

An alternative experimental demonstration in the same banner year was performed by a group under the direction of D. Wineland at the National Institute of Standards and Technology (NIST) in Boulder, Colorado, that swapped the roles of the photons and the atoms. As opposed to the Cal Tech experiment where the photons carried the qubits, and the atomic medium played a passive role as mediator of the interaction, in the NIST experiment the atoms carried the qubits, and the photons performed the operations on those atomic states.

Furthermore, the NIST experiment went farther than the Cal Tech experiment by actually implementing and demonstrating a C-NOT quantum gate.

Rather than using a beam of atoms that passed through a cavity, the NIST group captured and held a single beryllium atom inside a delicate apparatus known as an ion trap. The two qubits were represented by independent states of the atom. One qubit consisted of an internal electronic state of the atom, while the second qubit consisted of a mechanical (albeit quantum) vibration of the atom in the trap. The input qubits were prepared by a laser pulse that placed the atom in a superposition of internal and vibrational states.

The action of the C-NOT gate on the qubits was performed using laser pulses of specific frequencies and durations. The control pulses did not carry information. All the information in this implementation resides in the quantum states of the atom. The control pulses simply provide the physical mechanism that manipulates the qubits and entangles them, producing the output values of the gate. The logic gate performed with high fidelity at a processing rate of 20 kHz. Rates as high as 50 MHz were predicted as plausible.

Here, compared with the conditional phase shifter, we see optics playing a fundamentally different role. In the phase shifter the photons carry the qubits, while here the photons perform the operations that manipulate the quantum logic gate. In the atom trap the photons cease to interact directly. The qubit interactions are left to the atoms. We have seen this situation before in the hybrid optoelectronic computers of the first generation of the Machines of Light. As we noted then, light performs admirably as a courier of information, while matter performs best to control information. The advantage that matter has over optics is the electrons with their electric charge and strong Coulomb interactions. Control of information seems to work more effectively when the act of control is left completely to the electrons, although optics still plays an essential role. In the hybrid optoelectronic computers, as well as in the quantum gates, optics provides the communication channel. For instance, the qubits in the atoms have to be prepared in appropriate initial states by photons. The photons carry the qubit to the atoms, which store the qubits for the duration of the

computation. Photons then supply the quantum "program", changing the internal state of the atom and telling it to perform a C-NOT operation. The subsequent interaction between the qubits is carried out entirely within the atom. Finally, the qubits are read out by photons that pass the information downstream to the next gate. In this way, the photons are used as messenger and programmer, while electrons are used best for control. They work together in a photo-electric quantum network performing quantum logic.

The year 1995 was a watershed year for quantum computing. Loyd simplified the requirements of universal quantum computers, allowing Kimble to perform a conditional phase shifter as one element of such a computer. At the same time, Wineland demonstrated a two-qubit C-NOT gate using atoms to support the qubits and allow them to interact. By the end of the year, quantum computing had ceased being an exclusively theoretical science and had forged a beach-head in the laboratory.

But much remains to be done. Since 1995, improvements in the experiments have come slowly. The breakthrough experiments were sufficiently difficult that even after 5 years few additional experimental groups have had notable impact on the field. A chief obstacle to greatly improved performance of quantum logic gates is the difficulty of increasing the number of qubits to large enough numbers to be useful or interesting. The recent demonstrations by the NIST group of six trapped ions represents the state of the art [NOTE: NIST CLEO report]. It has taken Herculean efforts to achieve even this low level of parallelism in the number of qubits.

Another problem endemic to all quantum computing schemes is something called decoherence. The key element of quantum computing is the linear superposition of quantum states and the coherent interference among the states during computation. Yet the real world is constantly buffeting the quantum system, causing the coherence to decay in time. In the laboratory demonstrations, the computations were completed before the quantum states could decohere, but that was just for a single operation. Quantum computations of interest would require many operations to take place before decoherence could destroy the quantum

interference. Therefore, the best candidates for quantum computing are those that have the longest decoherence times. Single trapped ions are reasonable candidates because the decoherence times are around a millisecond. However, as the number of qubits, and hence the number of trapped ions, increases, the decoherence time decreases. This trend goes in the wrong direction and is troubling for the prospects for ion-trap quantum computing schemes.

One extremely important recent theoretical breakthrough has made the problem of decoherence a little less severe by allowing quantum computers to make mistakes, yet still arrive at useful answers. This breakthrough is in the area of quantum error-correction. New protocols that use entangled states repair the information carried by quantum states that have been damaged by decoherence. These error-correction schemes are a life-saver for realistic quantum computing because they make real-world implementations (with their unavoidable flaws and dissipation) candidates for realistic quantum computers.

The technological challenges faced by quantum computing are very difficult, just as they are for large-scale integrated electronic circuits based on single-electron transistors, and for holographic computers that use images as the unit of information. The problems that need to be overcome in each of these technologies will take years of concentrated effort by scientists and engineers. The beauty is that we have time. We do not need quantum computers tomorrow. We can live our lives without quantum parallelism. But quantum parallelism is inevitable because the problems that are faced today are chiefly technological and not theoretical. The fundamentals have been hammered out by fifteen years of imaginative search and discovery that have largely outlined the shape that the quantum Architecture of Light will provide to quantum computing. What remains is the hard work.

IN DEFENSE OF OPTIMISM

In the past hundred years pundits have generally underestimated technological growth. The history of technology teaches us that what is possible often becomes real. Furthermore, when things become real, demand for performance grows, and the technology grows faster than anyone expected. Just as estimates of the need for bandwidth on the internet have always underestimated the load, estimates for the uses of quantum communication and computing are likely to be conservative.

Certainly, counter examples abound. For instance, despite billions of dollars spent, nuclear fusion is no nearer being a source of energy today than it was thirty years ago. The technological difficulties turned out to be more complex than expected. Which points to a trend: that it is easier to work with smaller technology than larger. Feynman's proclamation "There's plenty of room at the bottom" is encouraging, because we have historically done well with miniaturization. Small scales are intrinsically easier to tackle than large scales because it is easier to work with less and less, than to need more and more to get bigger and bigger. The recent successes driving the burgeoning field of nanotechnology are ample proof of this.

One has to be careful to make a distinction between science fiction and science possibility. Warp speed and macroscopic teleportation are both probably impossible. But molecular computers that use quantum teleportation could be real. Because they are possible, and because they are small, they probably will happen. When? That is a different story. It is easy to be optimistic that something *will* happen, because there is a lot of time ahead of us. Saying exactly *when* is a lot tougher, because revolutionary technology needs to come along first. Sometimes this cannot be rushed. Often a critical mass of understanding has to be established for the next advance to take place. Ideas before their time are generally not useful.

This is why the pessimism of John Horgan's "The End of Science" is largely unfounded. Horgan claimed with this controversial book that science had already discovered most of what it could, and that all scientists now are merely hashing out the

details. In his sequel "The Undiscovered Mind" he goes further to claim that some things are fundamentally beyond the grasp of the human mind (like the human mind. Admittedly, his arguments are all true if learning and technology stop today. But they won't. There is a long time ahead of us in which to find ways around our current impasses. Even the question whether something is knowable is similar to the questions asked by computer scientists whether something is computable. Indeed, some things are not. But even these verdicts change. Quantum computing makes some things computable that previously were not. So our ideas of what is knowable may change as well.

We can therefore be optimistic that technology will continue. Advances will come. What was previously unknown will be understood. We have a long time ahead of us. It does not bother me if things come slowly, if I do not see them in my lifetime. Just knowing that things are possible is enough. And it keeps me excited.