THE ROYAL SOCIETY

Quantum-enhanced information processing

By M. Mosca[1,3], R. Jozsa[2], A. Steane[1] and A. Ekert[1]

[1] *Centre for Quantum Computation, Clarendon Laboratory, Department of Physics, University of Oxford, Parks Road, Oxford OX1 3PU, UK*
[2] *Department of Computer Science, University of Bristol, Woodland Road, Bristol BS8 1UB, UK*
[3] *Mathematical Institute, University of Oxford, 24–29 St Giles', Oxford OX1 3LB, UK*

Information is stored, transmitted and processed always by *physical* means. Thus the concept of information and computation can be properly formulated only in the context of a physical theory and the study of information processing requires *experimentation.* It is clear that if computers are to become much smaller in the future, their description must be given by quantum mechanics. Somewhat more surprising is the fact that quantum information processing can be qualitatively different and much more powerful than its classical analogue. In the following we will explain why.

Keywords: quantum computation; quantum communication; quantum cryptography; quantum algorithms; teleportation; quantum error correction

1. Introduction

Computers are physical objects and computations are physical processes. This sentence, innocuous at first glance, has remarkable and far-reaching consequences. To start with, computers are getting smaller and smaller as technology moves from gears to relays to valves to transistors to integrated circuits and so on. Today's advanced lithographic techniques can squeeze logic gates and wires of submicron size onto the surface of silicon chips. Soon they will yield even smaller parts and inevitably reach a point where logic gates are so small that they are made out of only a handful of atoms. On the atomic scale, matter obeys the rules of quantum mechanics, so, if computers are to become smaller in the future, new, *quantum* technology must replace or supplement what we have now. The point is, however, that quantum technology can offer much more than the miniaturization of gates and increased clock-speed of microprocessors. It can support entirely new kinds of computation with qualitatively new algorithms based on quantum principles and new modes of communication with many remarkable features. In recent years, a new quantum theory of information has been developed making it clear that fundamental questions regarding computability, data security and computational complexity are questions about physical processes rather than being purely mathematical abstractions.

In the following we will describe how quantum mechanics gives rise to new modes of computation that appear to be vastly more powerful than the capabilities of any conventional (classical) computing device. We will show how quantum physics allows us to communicate with absolute security and we will outline the process of quantum *teleportation*, according to which physical properties may be transferred from one object to another. Finally, we will discuss how these theoretical developments relate to current experimental technology and outline some of the formidable challenges that need to be overcome for practical realizations in the future.

2. Quantum information: bits and qubits

The handling of information is becoming an increasingly important part of everyday life. Anyone who has enjoyed listening to music on a CD, watching a video, browsing the Internet or has used an automatic bank teller machine has benefited from the recent explosion of developments in information processing, storage and transmission.

A given piece of information may be expressed in many different forms. For example, a cake recipe may be written in English or in Chinese. It may be spoken or stored in a computer memory coded as a sequence of 0s and 1s. Note that written words are arrangements of ink molecules on paper, spoken words are fluctuations in air pressure and a computer memory may be constructed from various kinds of electromagnetic components. In fact, all forms of information have a fundamental common feature: they all use physical objects to represent the information and processing is always performed by physical means. It follows that the possibilities and limitations of information storage and processing are ultimately dictated not by mathematical constructions, but by the laws of physics.

To use a physical system for information storage we must first identify and label a number of its distinguishable states. The basic unit of information is the *bit* (a contraction of 'binary digit'), which is represented by any physical system with just two distinguishable states, labelled 0 and 1. Any information may be represented by suitable sequences of bits. For example, the 26 letters of the alphabet may be unambiguously coded using some 26 of the 32 possible 5-bit strings (whereas the 16 possible 4-bit strings do not suffice), and any written text is then represented as a sequence of bits.

In a digital electronic computer, two levels of voltage are used to represent the bit values 0 and 1. One bit of information may also be encoded in two different polarizations of a photon or in two different electronic states of an atom. In the latter cases, the physical system is governed by the laws of quantum physics and is not well described by the formalism of classical physics.

This article will explore the fact that these fundamental differences between quantum and classical physics may be exploited to give rise to novel methods of information storage and processing, which, in principle, go well beyond the capabilities of current technology, which is based on classical representations of information.

Consider a bit coded in a quantum system such as the polarization of a photon. It is customary in quantum physics to label states using a curious notation of a half-pointed bracket enclosing a label. The two distinguishable states representing the bit values 0 and 1 are written as $|0\rangle$ and $|1\rangle$. This so-called *ket* notation was introduced by Paul Dirac in the 1920s to facilitate mathematical calculations. For us its odd appearance will serve as a constant reminder of the weirdness of quantum phenomena! According to the laws of quantum mechanics, the system may also be prepared in a *coherent superposition* of the two basic states. This is mathematically written as $|\ \ \rangle = a|0\rangle + b|1\rangle$, where   is the label of the superposed state and a and b are complex numbers. In such a state, the system is interpreted as being simultaneously in *both* state $|0\rangle$ and state $|1\rangle$ (to varying degrees depending on the values of a and b). Under physical evolution, the component parts $|0\rangle$ and $|1\rangle$ evolve separately and the system indeed behaves in a non-classical way, as though it were simultaneously in states $|0\rangle$ and $|1\rangle$. Any such quantum system that may encode the basic bit values as well as any possible superposition is called a *qubit* (pronounced 'queue bit').

An especially important feature of the quantum behaviour of qubits arises when we consider a string of several bits or qubits in a row. Consider first a row of two bits. There are four possible states: 00, 01, 10 and 11. If these are coded in a quantum system (i.e. we have two qubits), then, in addition to the basic states $|0\rangle|0\rangle$, $|0\rangle|1\rangle$, $|1\rangle|0\rangle$ and $|1\rangle|1\rangle$, we now also have general superpositions

$$|\phi\rangle = a|0\rangle|0\rangle + b|0\rangle|1\rangle + c|1\rangle|0\rangle + d|1\rangle|1\rangle.$$

Thus two qubits may simultaneously represent all four possible 2-bit strings (in a particular quantum way depending on the coefficients a, b, c and d). More generally, for n classical bits there are 2^n possibilities, but any single state of n classical bits may be described by a bit string of length n. In contrast, n qubits may simultaneously include all 2^n possibilities in superposition. In this sense n qubits are able to embody vastly more—exponentially more, in fact—information than n classical bits. The gap between n (classical) and 2^n (quantum) grows very rapidly with increasing n. This exponentially enhanced richness of multi-qubit states for representing information has profound consequences for information processing and communication, as described in the sections below.

Another fundamental non-classical feature of states of two or more qubits is the phenomenon of *quantum entanglement*. A general superposition state of two qubits (such as $|\phi\rangle$ above) has the property that each separate qubit cannot be assigned a separate state of its own. For example, the 2-qubit state

$$|\chi\rangle = \frac{1}{\sqrt{2}}|0\rangle|0\rangle + \frac{1}{\sqrt{2}}|1\rangle|1\rangle$$

cannot be expressed as a juxtaposition of 1-qubit states $(a|0\rangle + b|1\rangle)(c|0\rangle + d|1\rangle)$. In contrast, in any state of two classical bits (e.g. 01), each bit separately has a well-defined value and the whole is just the juxtaposition of the parts. In the quantum case, the information of the state does not reside locally in the separate qubits but is also distributed non-locally in a rich variety of possibilities of correlations between basic superposition components. Only the totality of all qubits together has a well-defined state. The qubits are thus said to be *entangled*. Note that classical bits may also exhibit correlations. For example, two bits might be prepared either as 00 or 11 but we do not know which. Then on examining the value of the first bit we immediately learn the value of the second bit as well; they are perfectly correlated. However, the quantum correlations involved in entanglement are far richer. It may be shown, for example, that the kinds of correlated behaviour exhibited by the entangled state $|\chi\rangle$ above cannot be reproduced by any model with just pre-assigned classical bit correlations.

A particularly interesting situation arises when entangled qubits are separated in space. This is not unusual in nature. For example, if an electron in a calcium ion is excited to a higher energy level and then allowed to fall back to its lower level, the excess energy is emitted in the form of two photons, flying off in opposite directions.

These photons are physically entangled in a state very similar to $|\chi\rangle$ above. The entanglement of the photons is independent of their spatial separation. They exhibit their peculiar quantum correlations even though there is no tangible physical connection between them. No physical process in the space between the photons can affect their entanglement. The correlation properties of such spatially separated entangled systems may be exploited for a variety of novel communication tasks, including the possibility of perfectly secure classical communication (impossible with classical bits) and the process of 'quantum teleportation', which will be described below.

So far we have discussed how information may be *embodied* in the state of a physical system. An important dual aspect of this physics of information is the question of how the information may be *accessed* or *read out*. In terms of physics, this is the question of what kinds of *measurements* are possible. Again we find a dramatic difference between classical and quantum physics. Information in classical physics may, in principle, always be read out completely and perfectly. By contrast, the laws of quantum physics (especially the uncertainty principle) imply that any attempt to read the information embodied in a quantum state will irretrievably disturb the state. Only a small amount of the potentially vast information content may be read out and most of it must remain inaccessible!

The full 'unknowable' information embodied in the identity of a quantum state is called *quantum information*. It has many peculiar properties. For example, we are all familiar with the idea of copying classical information: a cake recipe may be photocopied or copied out by hand or read out over the telephone, giving two or more copies of the information. In 1982 in a famous 'no-cloning' theorem, Wootters and Zurek showed that quantum information cannot be copied! For example, if quantum information is communicated from A to B, then the information is necessarily completely destroyed at A as it appears at B and no record of its identity can remain at A. The inaccessibility and non-clonability of quantum information may appear at first sight as entirely negative features, but these strange properties may be usefully exploited! For example, they can provide a means of perfectly secure communication, as we will elaborate later. Roughly speaking, any attempt to eavesdrop on the information must leave its imprint on the quantum state, which may be detected later by the legitimate communicating parties.

As any physical system evolves in time, the identity of its state changes. Thus, quantum physical evolution may be naturally viewed as the processing of quantum information. The laws of physics allow us to predict and *calculate* the changing identity of the state. Feynman (1982) made a remarkable and profound observation: if we calculate or 'simulate' the quantum evolution on any standard computer, then the amount of computational effort involved generally grows enormously as time passes. With each successive second of time evolution of the actual quantum physical system, the amount of computational effort needed to simulate it will grow so rapidly that soon the time and space requirements for the simulation will exceed all available resources. In fact, the resource requirements grow 'exponentially'; an important notion that will be elaborated in § 3.

Collecting the above ideas together, we arrive at a bizarre picture of the quantum world: in ordinary time evolution, nature processes quantum information at an astonishing rate that cannot be matched by any conventional computer simulation, yet when the processing is finished, most of the information is kept hidden and inaccessible to being read!

A quantum computer is any physical device that exploits the greatly enhanced information-processing power of quantum evolution for computational purposes. The very restricted accessibility of the processed quantum information provides a severe limitation on our ability to exploit the enhanced computing power. However, it does not annul it! As a basic illustrative example, suppose we have a quantum computer programmed to compute a function f. The computer evolves the labelled input state $|x\rangle|0\rangle$ to the labelled output state $|x\rangle|f(x)\rangle$. (Here, the second ket, initially set to $|0\rangle$, is the output register for the function value.) Now we may prepare the input register in an equal superposition $\sum_x |x\rangle$ of *all* possible input values x. Running the computer then yields the output state

$$|f\rangle = \sum_x |x\rangle|f(x)\rangle,$$

i.e. by evaluating the function *once* we evaluate *all* function values $f(x)$ in superposition. This process is called computation by quantum parallelism, as described by Deutsch (1985) in his seminal paper. The quantum information of the state $|f\rangle$ includes information about all of the $f(x)$ values, but, because of its inherent inaccessibility, we are unable to read them out. Nevertheless, small amounts of 'global' information—relating to all function values simultaneously—may be read out and this information may still require a vast (exponential) amount of computing effort to obtain on a conventional classical computer. For example, we may wish to discern simple patterns in the list of values, such as periodicity if the function is periodic (cf. § 3). In this way we may successfully exploit the greatly enhanced information-processing power of quantum evolution despite its inaccessibility. In § 3 we will elaborate on the idea of computational complexity and give some interesting fundamental applications of quantum computation.

3. Computational complexity and quantum computers

A computer is a piece of hardware that runs according to a program, or *algorithm*, that we specify depending on the task we wish to perform. The computer carries out the algorithm on a given input, producing the desired output. The hardware could be an old Commodore 64, for example, and the algorithm could be a simple spell-checking programme. This programme takes a file of text as an input and outputs a list of the misspelt words. Modern computers can take human voice input and translate it into voltages representing 0s and 1s, which, in turn, represent the sound patterns. These 0s and 1s are processed via the computer into other 0s and 1s representing the words corresponding to those sound patterns. These voltages are translated into various colours of light emanating from a colour monitor that visually displays the words. All of these translations and manipulations are controlled by hardware running according to algorithms in response to inputs.

The difficulty, or *computational complexity*, of a task is the amount of resources, such as time, space or energy, necessary to perform it. The computational complexity of a task depends, of course, on the size of the particular input. For example, the number of steps necessary to spell check a document with n bytes of text is proportional to n. Another example of a computational problem is multiplication. The input is a pair of n digit numbers, and the output is the product of these two numbers. The simple multiplication technique taught in primary school is an algorithm that uses roughly n^2 steps to compute the product. The reverse task, factorization, is to take an n-digit number (let us assume it is not prime) and to output two smaller numbers that multiply together to produce that number. The best-known rigorous classical algorithm for performing this task uses over $10^{\sqrt{n}}$ steps. The number $10^{\sqrt{n}}$ gets astronomically larger than n or n^2 as n grows, and performing this task for n even as small as 200 is beyond the computing power available on the earth today. When n is 400, $10^{\sqrt{n}}$ is 10^{20}. Even 1000 computers running at 1000 MHz would take three years to perform 10^{20} operations. When n is 900, it would take one million computers, running at 1000 GHz, 10 000 years to perform this many operations. As n grows only slightly, the difficulty of the problem quickly grows astronomically larger and quickly becomes intractable on any conceivable computing device. Multiplying numbers or spell-checking documents of these sizes, however, can be done in a fraction of a second. These latter two tasks are considered *tractable*, whereas the factoring problem is considered *intractable*. Another example of a problem that is considered intractable is that of deciding, for a given map of n countries, if it can be *properly* coloured using only three colours, that is if it can be coloured in such a way that no adjacent countries are coloured the same. For many maps the answer is simple, but there are huge families of maps for which this problem is believed to be very hard and the best-known algorithms require roughly 3^n steps.

The exact number of steps necessary depends on the details of implementation, which do not seem to qualitatively affect the difficulty of the problems. The reason we do not worry about the details of implementation is that we believe that any 'reasonable' computing device can efficiently simulate any other 'reasonable' computing device. For example, if Alice has a program that solves a problem in T steps on her computer, then we can efficiently translate that program to run on Bob's computer and solve it in at most T^k steps, where k is a constant that depends on the types of computers Alice and Bob have. It is thus convenient to consider a problem tractable if its computational complexity is no more than a polynomial in n (e.g. less than n, $n^4 + 5n + 1$ or n^k for some fixed number k), and intractable otherwise (i.e. if its complexity is super-polynomial, like $n^{\log n}$, $10^{\sqrt{n}}$ or 2^n).

It was widely believed that this crude notion of tractability does not depend on how you implement your computer, digital or analogue, mechanical, electronic, or optical, or in any other 'reasonable' way. Quantum computers, if they are indeed 'reasonable' computing devices (see § 5), pose a serious problem for this belief since we believe that they cannot be efficiently simulated by any classical device. This is the essential content of Feynman's observation mentioned in § 2. There are several tasks that are known to be tractable on a quantum computer yet are very strongly believed to be intractable on any classical computer. The most famous example is the factorization problem. Shor (1994) discovered a quantum algorithm for factoring an n-digit number, which runs for less than n^3 steps. It can be mathematically shown that the problem of factoring the number N is closely related to studying the sequence $1, 2 \bmod N, 2^2 \bmod N, \dots, 2^x \bmod N, \dots$, where $y \bmod N$ is the remainder when y is divided by N (e.g. $79 \bmod 35 = 9$). This sequence will eventually start to repeat and cycle through the same numbers. To factorize N, it suffices to find the period of sequences like this (replacing 2 with another number if necessary). That is we wish to find the smallest positive number r such that the sequence repeats or cycles after every r steps. We do not care about the actual values in the sequence, apart from the fact that they cycle. Quantum computers have an edge over classical computers

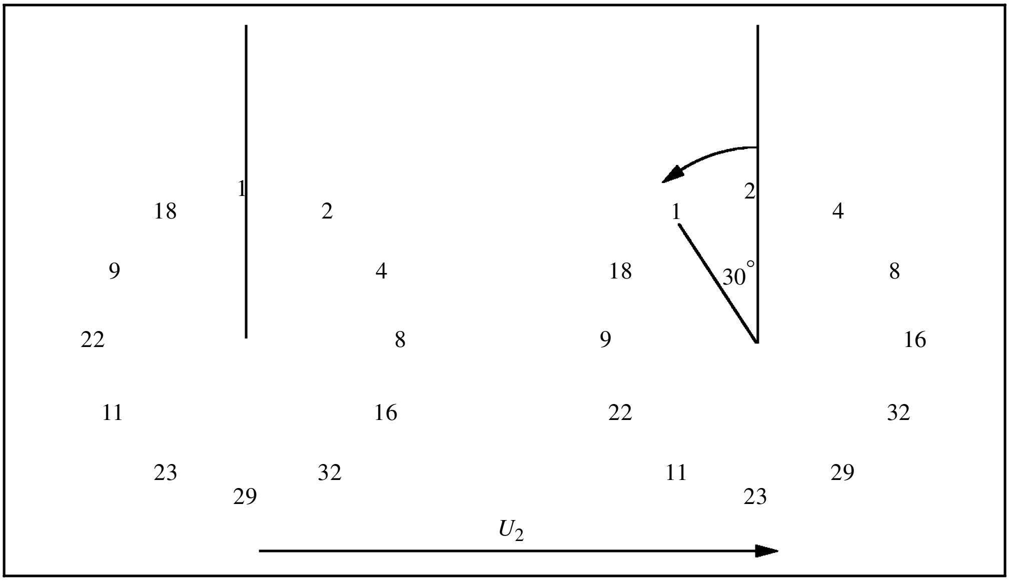

Figure 1. Quantum order-finding. This example uses $N = 35$ and $a = 2$, and we seek the period $r = 12$. The operation U_2, which multiplies by $2 \bmod 35$, sends 1 to 2, 2 to 4, ..., 9 to 18 and 18 to 1. Individually, these bits of information give little clue that $2^{12} \bmod 35$ is 1. Multiplying the above combination of all the powers of 2, however, will rotate the combination at a rate of $\frac{1}{12}$ revolutions (30°) per multiplication. This rate can be approximated by a simple quantum algorithm, and the denominator, 12, can be extracted. So now we know that $2^{12} \bmod 35 = 1$ and this information can be used to factor 35. The factors of $N = 35$ are no big secret, but try finding those of $N = 27\,997\,833\,911\,221\,327\,870\,829\,467\,638\,722\,601\,621\,070\,446\,786\,955\,428\,537\,560\,009\,929\,326\,128$ $400\,107\,609\,345\,671\,052\,955\,360\,856\,061\,822\,351\,910\,951\,365\,788\,637\,105\,954\,482\,006\,576\,775\,098\,580$ $557\,613\,579\,098\,734\,950\,144\,178\,863\,178\,946\,295\,187\,237\,869\,221\,823\,983$, the RSA 200-digit challenge. A quantum computer could.

in studying such global properties or patterns. Probing just a few values of this sequence gives little chance of finding the period: this approach requires looking at roughly $\sqrt{N}$ elements of the sequence. Quantum computers, however, can be in a state containing all the elements of this sequence. A physical realization of this sequence, when multiplied by 2, will shift and will cycle back to its original state after r shifts. An object that cycles after r steps must, in a sense, be rotating at a rate of k/r cycles per step for some integer k. A physical realization of this whole sequence, which a quantum computer can efficiently create, will rotate at such a rate, and a quantum computer can study this rate of rotation by looking at superpositions of elements in the sequence and not at individual entries. See figure 1 for an illustration.

Quantum computers also help with problems of less structure. Consider the problem we mentioned earlier, of properly colouring a given map using only three colours. This problem has the property that we can easily check solutions (since there are less than n^2 pairs of adjacent countries to check), but deciding if any solution exists can be very difficult (for some maps, no significantly better method is known than simply trying all 3^n possible colourings to see if any are proper). More formally, denote by f the function that takes as input such a colouring, x, and outputs 1 if it is proper

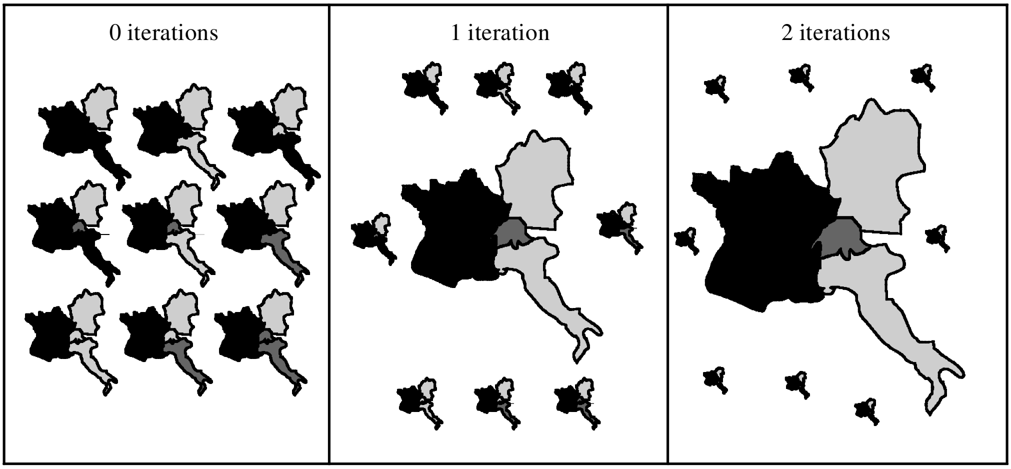

Figure 2. Quantum searching. By simply preparing a uniformly weighted combination of all possible colourings (we assume, without loss of generality, that France is coloured blue (the darkest colour) and Germany is green (the lightest colour), leaving $N = 9$ possible colourings to consider for Italy and Switzerland), we get the unique proper colouring with probability $\frac{1}{9}$. The first iteration of the quantum search iterate increases the probability of observing the proper colouring to *ca.* 73%, and a second iteration gives us the proper colouring with probability over 98%. This is a good time to stop and observe the colouring. Of course, for such a small map the answer is obvious; however, for even as few as $n = 50$ countries this problem can be a nightmare. Exhaustively trying all 3^{50} (this is over 10^{23}!) colourings is not feasible, whereas $\sqrt{3^{50}}$ (less than 10^{12}) repetitions is within the realm of possibility.

and 0 otherwise. We can evaluate this function on a superposition of all the possible inputs but we cannot force ourselves to observe a solution with $f(x) = 1$ using only one evaluation of f. However, using a quantum search algorithm developed by Grover (1996) (illustrated in figure 2), we can cleverly interfere the many superposed components in such a way that on the kth iteration we enhance the probability of observing a solution by roughly k/N. Thus, if we iterate this process roughly $\sqrt{N}$ times, the probability will be close to 1 (since $1/N + 2/N + \cdots + k/N \approx k^2/2N$). That is, we can drive our computer into a state containing almost only solutions to $f(x) = 1$ using only $\sqrt{N}$ evaluations of f, whereas randomly sampling inputs to f requires roughly N evaluations to find a solution with high probability (since we are only increasing our chance of finding a solution by $1/N$ each time).

In summary, some problems, like factorization, have an algebraic structure that quantum computers can exploit to a much greater extent than any known classical algorithm. These problems that were once thought to be intrinsically hard, that is not efficiently solvable by any reasonable computing device, we now know can be solved in very few steps on a quantum computer. Also, simple searching algorithms can also be speeded up by a square-root factor. For further reading on these and other quantum algorithms, see the papers by Deutsch & Jozsa (1992), Ekert & Jozsa (1996), Cleve *et al.* (1998), Steane (1998), the March 1998 issue of *Physics World*† or Vazirani (1997).

† Special issue on quantum information. *Physics World* **11** (March 1998).

4. Communication and information security

The previous section illustrates how exploiting the quantum nature of information can have a dramatic impact on the computational complexity of many problems. However, information is a valuable resource, which we may wish to share with our friends, or, perhaps, keep out of the hands of our foes or competitors. It is natural to ask if the quantum nature of information changes the rules of the game in the communication and security of information. The answer is 'yes', quantum physics has a dramatic effect. We will describe three examples: quantum teleportation, quantum key distribution, and quantum communication complexity.

As mentioned earlier, any interaction with a quantum system that extracts information about its state will disturb the system. In fact, any non-trivial interaction between the quantum state and its environment will alter the state. This means quantum information is extremely delicate, and to remain fully intact it must not interact with its environment. This makes storing and transporting quantum information a very challenging task. The error-correcting codes that we will discuss later use entanglement and redundancy as a means of keeping quantum information intact and resistant to interactions with the environment.

However, suppose that two people, Alice and Bob, are separated in space, and wish to communicate some quantum information. A powerful technique, known as *quantum teleportation*, developed by Bennett *et al.* (1993), allows Alice and Bob to communicate quantum information by sending only a small amount of classical information. The advantage is that 'classical' information is very robust and much less sensitive to interaction with the environment. To achieve this task, Alice and Bob must also be in possession of some shared entanglement. More precisely, at some point in the past when Alice and Bob were together, Bob created the entangled pair of qubits

$$|\chi\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$$

and carefully gave one-half to Alice, who took it away with her. With this resource, at any time in the future Alice can send Bob a quantum state $|\Psi\rangle$ by performing a special quantum measurement on $|\Psi\rangle$ and her share of $|\chi\rangle$ and sending Bob the measurement result classically. Bob can then reproduce the state $|\Psi\rangle$ and Alice no longer has her copy of the state. The protocol is described in figure 3.

The quantum teleportation of states of light has been realized in several laboratories around the world.

We next turn to the issue of secure communication. It is interesting to note that a commonly used method of secure communication—the method of public key cryptography—relies for its security on the computational intractability of certain computational tasks, such as factoring large numbers. As discussed in § 3, the computing power of quantum processes may be used to break such ciphers and render them insecure! However, we will now see that further quantum effects—the uncertainty principle and the inaccessibility of quantum information—may be exploited to provide new methods of communication that *are* unconditionally secure. They do not rely on any unproven assumptions about computational intractability.

The sensitivity of quantum information to interaction with its environment might seem like nothing but an inconvenience, but it is a very useful tool in the art of secret communication. One very useful primitive in cryptography is the distribution

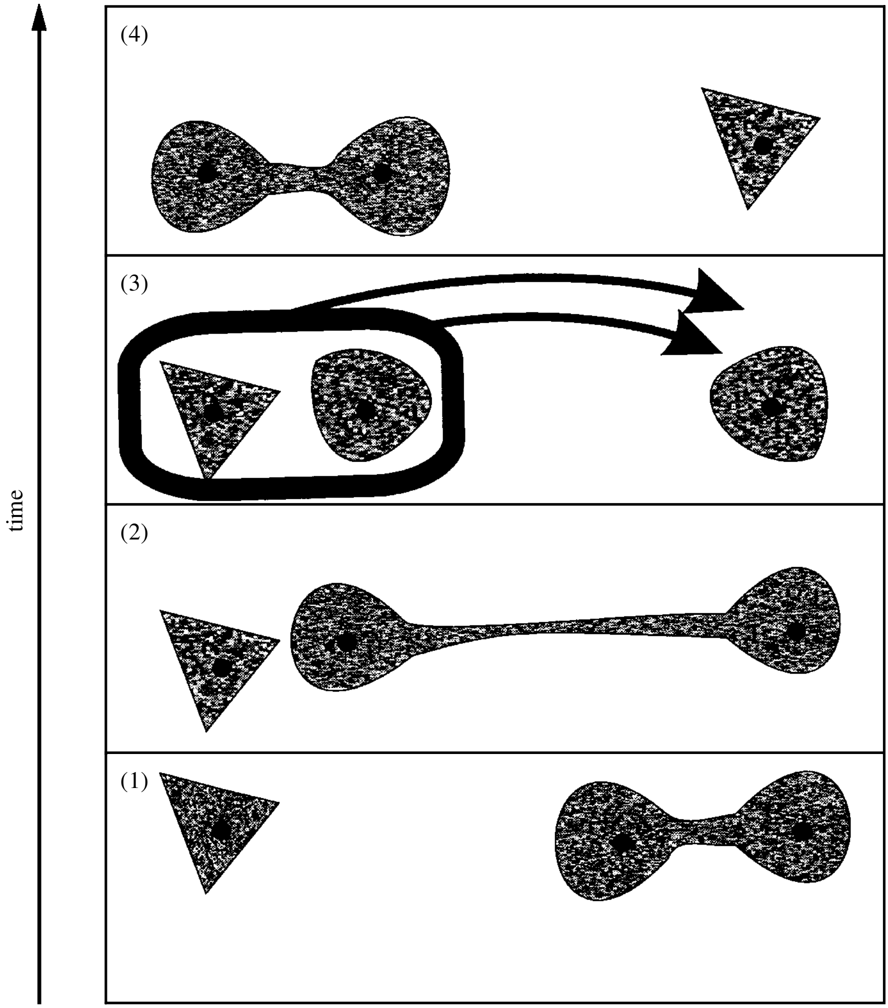

Figure 3. Quantum teleportation. In step 1, Bob prepares two qubits in the state $|\chi\rangle = (1/\sqrt{2})|00\rangle + (1/\sqrt{2})|11\rangle$. He then sends one of the two qubits to Alice (step 2), who has a special qubit $|\Psi\rangle$ she wishes to send Bob at some point in the future (she can, in fact, decide at any point before teleportation what state $|\Psi\rangle$ she wishes to send). Although the particles are far apart, their joint state cannot be described separately as they are correlated in a very strong quantum way. When Alice wishes to send this qubit to Bob, she interacts it with her share of $|\chi\rangle$ and performs a quantum measurement to obtain two bits of information, which she then sends to Bob by classical means (step 3). These two classical bits tell Bob which of four operations to apply to his share of the state, which was formerly $|\chi\rangle$. Once he applies the operation his qubit will be in the state $|\Psi\rangle$ (step 4), while Alice now has an entangled pair of qubits in a state very similar to the original $|\chi\rangle$ (the difference is described by the classical bits she sends Bob), and independent of the qubit she teleported Bob.

of a common secret key between Alice and Bob. The key itself is random, that is, it contains no valuable information. What is valuable is the fact that Alice and Bob have the same key (100% correlation, as between the two qubits of $|\chi\rangle$), and that no one else has any information about it (0% correlation). A shared secret key can be used in an array of private key cipher systems, including the well-known US data encryption standard (DES), and it also provides the only provably secure cipher, the so-called Vernam cipher.

For simplicity, let us assume that Alice and Bob have some way of confirming each others' identity and that they have access to some public communication channel that can be eavesdropped upon but not actively altered, such as a radio transmitter and receiver.

Any information that is stored classically and exchanged between Alice and Bob can, in principle, be copied by an eavesdropper, Eve, giving her a copy of the information (100% correlation) without a trace. Quantum information, however, cannot be copied or studied without a trace! Several protocols have been devised that allow Alice and Bob to produce a common key by the exchange of quantum information or the sharing of entangled particles. The advantage of doing the key distribution quantumly is that any eavesdropping or tampering will, by the uncertainty principle, affect the information being exchanged, and the correlation between the keys Alice and Bob share must be reduced. If Eve learns any significant amount of information about the key Alice and Bob share, then the uncertainty principle requires that with high probability the keys will not be identical. Alice and Bob can detect this effect of eavesdropping and bound the amount of information Eve has about their key. If the amount of information is small, there are methods, described in § 5, that allow Alice and Bob to distill their keys so that the remaining keys are almost certainly equal and, furthermore, there is now only a negligible correlation with any of Eve's information. If Eve has obtained too much information, they simply abandon the not-so-secret key. These sort of key distributions have been implemented over distances from 30 cm to tens of kilometres at several laboratories around the world. For further reading on quantum cryptography, see Bennett *et al.* (1992), Steane (1998) or the March 1998 issue of *Physics World*.

Quantum teleportation and key distribution deal with Alice sending information that she explicitly possesses to Bob, and vice versa. Suppose, on the other hand, that Alice and Bob need to communicate with each other to figure out some valuable piece of information. For example, Alice and Bob work in different places and their babysitter has just asked Bob if she would be needed on any of the next seven days as she wishes to go on vacation. Alice and Bob's task is to decide if there is some day next week when neither will be home, but they have no prior knowledge of each other's schedule. Neither Alice nor Bob individually possesses this information, but together they possess enough information to figure it out. Thus, they must communicate some information to each other to decide the answer. The answer could be a very small amount of information. The *communication complexity* of this distributed computation problem is the amount of information Alice and Bob must send back and forth in order to figure out the answer, and this is usually much more than the length of the answer. The communication complexity is the amount of long-distance telephone charges they must accrue in order to solve the communication problem. An easy solution is for Alice to fax Bob her entire schedule, and Bob figures out if they need the babysitter. In general, if Alice and Bob want to decide if they need a babysitter in the next N days, using this simple protocol uses in the order of N bits of information between the two. This seems like a lot of communication in order to compute just one bit of information ('yes' or 'no'), but if Alice and Bob only exchange information classically, it is necessary in some worst-case scenarios. What happens if

Alice and Bob can send *quantum* bits of information back and forth? Alice and Bob can, in fact, solve this problem by exchanging only roughly $\sqrt{N}\log N$ quantum bits of information, as shown by Buhrman *et al*. (1998).

The use of quantum information thus opens many doors for Alice and Bob. For certain communication or distributed computation problems—such as the scheduling problem described above, or the communication between different processors in a multi-processor computer—the amount of information that must be communicated back and forth can be greatly reduced by using quantum bits instead of classical ones. Sending such quantum bits can be done most safely by quantum teleportation. By using quantum information, their communications can also remain private without relying on assumptions about the computational intractability of any problems.

5. 'In principle' versus 'in practice'

We have already emphasized that the significance of quantum computers, and of quantum-information physics in general, is not just that it offers a faster way of doing some computing, but that it offers a qualitatively different way of conceiving of information and computing. Nevertheless, we should beware of a difficulty that often arises in the context of computers, namely that some seemingly promising idea turns out to be completely useless because the difficulties of realizing it in a system of useful size were vastly underestimated. Furthermore, quantum coherence is notoriously fragile, especially when complicated quantum systems are involved. Experimenters have to go to great lengths to preserve the coherence of even small quantum systems (such as a few atoms or photons) whose complexity is no greater than a few qubits' worth. So, when we contemplate the quantum computer, these considerations make us smell a rat. Is the striking computational power of the quantum computer actually illusory, since it is based on assuming a degree of precision in the construction of the computer, which, for all practical purposes, is impossible to achieve? The suspicion that this was the case certainly kept the interest and excitement in the field somewhat damped down in the early 1990s, and although we are about to discuss the tremendous progress that has been made towards understanding and tackling this issue, it remains the chief reason for caution regarding the future.

Let us examine the experimental-precision requirements for quantum cryptography and quantum computation. Cryptography protocols have to be carefully designed in order to allow a certain amount of experimental imperfection and noise, but it turns out that this is not too severe a problem. When setting up a cryptographic key, the main effect of noise is that Alice and Bob's data are going to differ to some extent, in a roughly random way, whether or not an eavesdropper is present. So how can they distinguish an eavesdropper from random noise? They can't, but as long as the noise level is low enough, they do not need to. If the noise (error probability) is below some level ϵ per bit, then they can rule out strong eavesdropping at least. They then go on to clean up their data by doing 'parity checks' (see below). When they thus correct errors they also shut out any remaining weak eavesdropping because the (limited) data gathered by Eve are effectively dropped by Alice and Bob.

The parity check is a very simple yet fundamentally significant concept of information science. The parity of a string of bits simply indicates whether the string contains an even or odd number of 1s. For example, 01100100 has odd parity (parity equal to 1), while 01100110 has even parity (parity equal to 0). In the case of cryptography, Alice and Bob have classical bit strings (the results of their measurements), from which they select agreed random subsets of bits. They each calculate the parity of the subset, and publicly report their results. If they found the same parity, then that part of Alice and Bob's data probably contains no errors. It might contain two or any even number of errors, which they can test by further random checking. It turns out that such procedures are successful for error rates in the apparatus at the level of a few per cent error probability per bit, which means a real working system not only could be built, but has been built with current technology.

The working cryptography systems are based on using photons as qubits. Weak light pulses containing single photons are sent through a specially designed interferometer at Alice's laboratory, and then transmitted several kilometres down standard fibre-optic telecommunications systems to Bob, who has a similar interferometer. Taking advantage of various ingenious methods to enhance the preservation of the photon's quantum states, the tolerated bit error rate of a few per cent can be achieved.

This is in contrast with the situation with quantum computers, where no computers exist nor will exist for some time. However, current technology does permit us to build few-qubit systems that allow the basic concepts of quantum information processing to be demonstrated and from which we can learn how to go further.

In principle, there are myriad ways one might conceive of an experimental system, but in practice there are only a few that meet the severe requirements of sufficient complexity and sufficient controllability. Currently, there are two systems in which three or so qubits can be manipulated: these are the ion trap and the nuclear magnetic resonance (NMR) spectrometer. The former takes advantage of high-precision atomic physics techniques, such as laser cooling, to allow one to manipulate the motion and internal states of individual charged atoms, or ions. The ions are confined by a set of electrodes in high vacuum and addressed by laser beams focused onto them. The latter uses a standard NMR spectrometer and manipulates the nuclear spins in a simple molecule using pulsed magnetic fields. These two approaches are to some extent complementary, in that they have different strengths and weaknesses.

The former allows complete interrogation of a single line of qubits, while the latter uses a liquid sample containing billions of molecules, and interrogates the average state. To date, thorough manipulation of two qubits is more or less routine for NMR work, and limited manipulation of three to seven qubits has been reported. Meanwhile, a single ion trap experiment has achieved manipulation of two and three ions, though not yet a completely general set of operations. Typically, the precision of these experiments allows *ca.* 100 quantum logic gates, such as exclusive-OR, to be applied before the state of the system is lost to noise and imprecision.

Needless to say, all this is a very long way from the level of quantum computing we would need to achieve a real rival to classical methods. In order to see how far, we can consider an example task for a quantum computer. We will take this to be the factorization of a thousand-digit number. Although we do not expect factorization to be the main use of quantum computers in the future, this is a quantum algorithm we understand and it may give us a feel for the size of machine we need to envisage in order to do useful processing that could not be done on classical computers. Shor's algorithm for this task would require *ca.* $L = 3000$ qubits to store the thousand-digit number to be factorized, and a further $4L = 12\,000$ as workspace, making $K = 15\,000$ in all. The algorithm requires $Q \simeq 300L^3 \simeq 10^{13}$ elementary steps (logic gates) for completion. In the course of such a computation, most qubits are involved most of the time, and we require that no qubit ever decoheres. In other words, the level of noise in the operations, and of coupling between the qubits and their environment, must be kept below one part in KQ, or *ca.* 10^{-17}. This level of precision is so difficult to attain that one can rule it out as more or less impossible. The reason for this is to do with the connection between legitimate operations and troublesome ones: in order to manipulate qubits, there must exist a physical interaction between them and the controlling machinery, but that implies that there is also a coupling between the qubits and other stuff, such as electrical noise, or, when all else has not failed, quantum mechanical vacuum fluctuations.

The situation looks at this stage like a house of cards: we can build a layer or two, but when contemplating a really tall house, the task seems hopeless.

Nevertheless, the estimate we just made could equally have been applied to classical computers. They routinely complete computations requiring this number of bits and gates without making an error (well most of the time anyway!). How is their remarkable reliability possible? The essential ingredient is that every bit in a classical computer is under scrutiny all the time! The on-chip transistor circuits play a role equivalent to that of the spring in a mechanical switch, forcing the switch, or in this case the bit value, one way or the other. Any small departure is 'detected' and forcefully suppressed by such strong 'springs'. Unfortunately, no such scrutiny is permitted in a quantum computer! To examine a qubit is to render it useless for quantum computation. What we need is a much more subtle approach, where we do not let our right hand know what our left hand is doing: we would like to detect erroneous changes in the quantum computer's state without learning anything about the state itself. We now know how to do this.

Quantum error correction (QEC), depicted in figure 4, is a set of powerful and elegant ideas that springs rather naturally from the union of information science with quantum theory. A central ingredient is the parity-checking operation, which is adapted to the quantum context as follows. In general, the state of a set of qubits might combine odd and even parity; take, for example, the three-qubit state $(a|100\rangle + b|110\rangle)$. However, we can choose to restrict the way we set up the computer, so that it only uses states of even parity. Suppose we use our qubits in groups of three, in such a way that the parity of any pair in the group is even. We thus only ever use the states $|000\rangle$ and $|111\rangle$ or the general combination $a|000\rangle + b|111\rangle$. This restriction means that every triplet of physical qubits in the machine is only able to perform the work of one logical qubit in the computation, but it enables us to stabilize the computer. How? A random noise process might change the state of a group of qubits to something like

$$a\sqrt{1-\epsilon^2}\,|000\rangle + b\sqrt{1-\epsilon^2}\,|111\rangle + \epsilon a|001\rangle + \epsilon b|110\rangle. \tag{5.1}$$

We now perform two parity measurements. A measurement of the parity of the first two bits reveals that it is even, as it should be, but a measurement of the parity of the second two bits can have two results. Such a measurement imposes a selection on the computer state, either selecting out all the even-parity results, $a|000\rangle + b|111\rangle$, or all the odd-parity ones, $a|001\rangle + b|110\rangle$. In the former case, the computer state has been restored to the noise-free one, while in the latter we know there is a problem because we just discovered the parity to be odd. However, the only error that can cause the parity of the first pair to be even and the second pair to be odd is a flip of the third bit. Therefore, we can deduce the error from our parity measurements, and undo it by deliberately flipping the third bit back again. We thus correct the computer by detecting the error, while at no time learning anything (such as whether b is greater than a, etc.) about the underlying quantum state $a|000\rangle + b|111\rangle$.

The noise might be more insidious, corrupting $a|000\rangle + b|111\rangle$ to $a|000\rangle - b|111\rangle$ for example. This will ruin the quantum computation just as surely as the first type of noise, but is not revealed by the parity check. However, another type of parity measurement involves using quantum states consisting of several terms added together, and measuring the number of minus signs in the total state. This is a new 'quantum' version of parity that does not have any analogue in classical computation. In combining both approaches, we need to use quite subtle and, as it turns out, highly entangled quantum states, but, fortunately, their exact construction can be based on the classical theory of error correction, which has been studied for 50 years. The extraordinary thing about all this is that the carefully constructed entangled states (called quantum error-correcting codes) do not require exponentially expanding resources in order to achieve exponentially suppressed noise levels.

At this point, the reader should still not feel altogether happy about building the house of cards. Although we introduced corrective measures, what if they themselves are faulty as they must be in any real system? Even this consideration can be met. We need to invoke a further set of new and subtle ideas, which together go by the name 'fault tolerance'. The basic problems we face are that a 'correction' based on bad parity information will actually make the situation worse not better, and even a perfect gate operation will couple previously occurred errors from one qubit to another, thus spreading the noise. The essence of the answer is twofold. Firstly, the parity information we require for QEC is essentially classical once we have obtained it, so it can be extracted repeatedly until a consistent result is obtained, and only then is the computer corrected. Secondly, because QEC is exponentially efficient, it can mop up not only the noise in the computation itself, but also that generated by the checking operations, as long as they do not amplify noise exponentially. Such an amplification is avoided by careful construction of the quantum networks, restricting routes for error propagation.

With all these ideas working together, the 'realistic' quantum computer looks very different from the idealized noise-free one. The latter is a silent shadowy beast that we must never look at until it has finished its computations, while the former is a bulky thing that we 'stare at' all the time, via our error-detecting devices, yet in such a way as to leave unshackled the shadowy logical machine lurking within it. For every elementary logic gate of the logical computation, the corrective procedures involve thousands of gates and thus dominate the machine, but we have gained a great deal because the machine can now tolerate noise in every one of these gates at the level of 10^{-5}. This degree of precision is achievable, in contrast to the figure of 10^{-17} that we had to contemplate previously. For further reading, see Steane (1998), Lloyd (1993), the March 1998 issue of *Physics World* or Vazirani (1997).

6. Future prospects

When the physics of computation was first investigated systematically in the 1970s, the main fear was that quantum-mechanical effects might place fundamental bounds on the accuracy with which physical objects could realize the properties of bits, logic

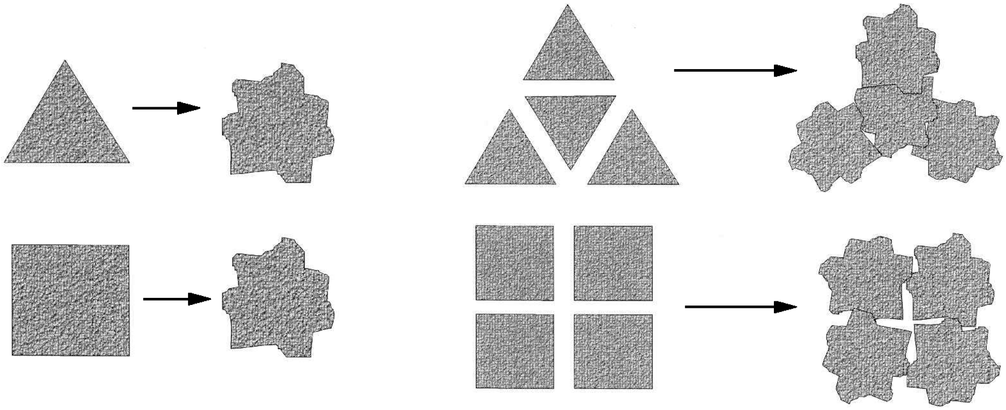

Figure 4. Quantum error correction. A single bit of quantum information is represented by two states, here symbolized by a triangle and a square. Noise will blur these states, making them indistinguishable. However, if, for example, we use four qubits together, we can construct joint states in which the quantum information is stored communally. Even though every qubit becomes blurred, the stored joint state is still detectable in the pattern among the blurred qubits (in practice, the pattern is parity information).

gates, the composition of operations, and so on, which appear in the abstract and mathematically sophisticated theory of computation. Those fears have been proved groundless. As we have explained, quantum mechanics, far from placing limits on what classical computations can be performed in nature, permits them all, and in addition provides whole new modes of computation, including algorithms that perform tasks that no classical computer can perform at all (secure key distribution) or can perform albeit not efficiently (factorization). Experimental and theoretical research in quantum information processing is accelerating worldwide. New technologies for realizing quantum computers are being proposed, and new types of quantum information processing with various advantages over their classical versions are continually being discovered and analysed.

The current challenge is not to build a full quantum computer right away but rather to move from the experiments in which we merely observe quantum phenomena to experiments in which we can *control* these phenomena. This is a first step towards quantum logic gates and simple quantum networks. The next challenge is to scale up quantum devices. The more components, the more likely it is that quantum computation will spread outside the computational unit and will irreversibly dissipate useful information to the environment. Thus the race is to engineer submicroscopic systems in which qubits interact only with themselves and not with the environment. New techniques, such as quantum error correction and fault-tolerant computation, together with new technologies, will allow us to achieve this task in the not too distant future. At this point, new devices, such as ultra-precise quantum clocks and entanglement-enhanced frequency standards, will supersede the existing ones. New quantum sources of light, better and cheaper photo-detectors, quantum repeaters, will make quantum cryptography a serious alternative to more traditional methods of encryption. Retransmission of information via satellite can even make quantum cryptography suitable for a long distance communication. The next millennium will witness computer technology departing from silicon and new quantum algorithms run on quantum processors. Quantum computers will eventually become a reality with a number of useful applications. There is no way we can predict them now. Imagine Charles Babbage being asked about the future of his analytical engine: would he have predicted word processors, games and the Internet? Whatever these applications might be, one of them will be an efficient simulation of complicated quantum phenomena. This will help to set up new experiments that will refute quantum theory and will let us learn more about the laws of physics. These breakthrough discoveries and their implications for computations will doubtless be nicely summarized by our descendants in the millennium edition of the *Philosophical Transactions* in the year 3000.

A.E. acknowledges the support of the Royal Society. R.J. acknowledges the support of the EPSRC. M.M. thanks the CESG and Wolfson College for their support. A.S. acknowledges the support of the Royal Society and St Edmund Hall.

We are grateful to Wim van Dam and Hitoshi Inamori for helpful conversations.

References

Bennett, C. H., Brassard, G. & Ekert, A. K. 1992 *Scientific American* (October), 50.

Bennett, C. H., Brassard, G., Crepeau, C., Jozsa, R., Peres, A. & Wootters, W. K. 1993 *Phys. Rev. Lett.* **70**, 1895.

Buhrman, H., Cleve, R. & Wigderson, A. 1998 In *Proc. 30th A. ACM Symp. Theory of Computing*, p. 63. New York: ACM.

Cleve, R., Ekert, A., Henderson, L., Macchiavello, C. & Mosca, M. 1998 *Complexity* **4**, 33.

Deutsch, D. 1985 *Proc. R. Soc. Lond.* A **400**, 97.

Deutsch, D. & Jozsa, R. 1992 *Proc. R. Soc. Lond.* A **439**, 553.

Ekert, A. & Jozsa, R. 1996 *Rev. Mod. Phys.* **68**, 733.

Feynman, R. P. 1982 *Int. J. Theor. Phys.* **21**, 467.

Grover, L. 1996 In *Proc. 28th A. ACM Symp. Theory of Computing*, p. 212. New York: ACM.

Lloyd, S. 1993 A potentially realizable quantum computer. *Science* **261**, 1569–1571.

Shor, P. 1994 In *Proc. 35th A. Symp. Foundation of Computer Science*, p. 124. Los Alamitos, NM: IEEE Computer Society.

Steane, A. 1998 *Rep. Prog. Phys.* **61**, 117.

Vazirani, U. (ed.) 1997 Special section on quantum computation. *SIAM J. Computing* **26**, 1409.

AUTHOR PROFILES

M. Mosca

Michele Mosca is the Robin Gandy Junior Research Fellow at Wolfson College, Oxford. He obtained a Bachelor of Mathematics from the University of Waterloo, Canada, in 1995 (alumni gold medal winner), and an MSc in Mathematics and the Foundations of Computer Science (with Distinction) from the University of Oxford, where he also completed his DPhil on quantum computer algorithms. He has made major contributions to the unification of quantum algorithms and is a co-inventor of the polynomial method for studying the limitations of quantum computing. Together with collaborators at Oxford he has implemented some of the first quantum algorithms. He will continue working on quantum computation as Assistant Professor of Mathematics, University of Waterloo. His other interests include rowing and languages.

R. Jozsa

Born in Melbourne Australia, Richard Jozsa studied at Monash University, graduating with first class honours in Mathematics. He obtained his DPhil in 1981 from the University of Oxford. Throughout the 1980s he held postdoctoral positions at Oxford, McGill, Sydney and other universities in Australia. Richard began working in quantum computation in 1989. In 1992 with David Deutsch he gave the first demonstration of the power of quantum computing over classical computing. In 1994 he co-invented quantum teleportation. In 1996 he was a Royal Society Leverhulme Senior Research Fellow and subsequently became Professor of Mathematical Physics at the University of Plymouth. Currently, Richard is an EPSRC Senior Research Fellow and Professor of Computer Science at the University of Bristol. Recreations include gastronomy and playing the violin.

A. Steane

Andrew Steane obtained a doctorate in Physics from Oxford University in 1991, in the area of experimental atomic physics. He was a Junior Research Fellow at Merton College, Oxford, and then a European Research Fellow at the Ecole Normale Supérieure, Paris, where he developed experiments in laser cooling and interferometry of atoms. In 1995 he took up a Royal Society University Research Fellowship at the Clarendon Laboratory, Oxford, and was appointed Fellow, by Special Election, of St Edmund Hall in 1996. Through studying the theory of quantum interference involving many particles, he established the basic principles of quantum error correction theory. He now combines this research with experiments on trapped ions. Aged 34, he is married and lives in north Oxford, leads a children's group at his church, and is a keen walker and singer.

A. Ekert

Artur Ekert is a Professor of Physics at the University of Oxford and a Fellow and Tutor of Keble College, Oxford. He obtained his DPhil from Oxford in 1991. In his doctoral thesis he introduced the entanglement-based quantum cryptography. From 1991 until 1998 he was a Research Fellow at Merton College, Oxford, and in 1993 he was elected the Royal Society Howe Research Fellow. Since 1992 he has been in charge of the Quantum Computation and Cryptography Research group (recently turned into a research centre) in the Clarendon Laboratory, Oxford. He has been a consultant on quantum information technology (both to industry and government agencies). For his work in quantum cryptography he was awarded the 1995 Maxwell Medal and Prize. He enjoys outdoor sports.

Standing (from left to right): Arthur Ekert, Andrew Steane, Michele Mosca.
Seated: Richard Jozsa.