

Quantum information and computation

Charles H. Bennett & David P. DiVincenzo

IBM Research Division, T. J. Watson Research Center, Yorktown Heights, New York 10598, USA

In information processing, as in physics, our classical world view provides an incomplete approximation to an underlying quantum reality. Quantum effects like interference and entanglement play no direct role in conventional information processing, but they can—in principle now, but probably eventually in practice—be harnessed to break codes, create unbreakable codes, and speed up otherwise intractable computations.

Information and computation theory have undergone a spurt of new growth, and a renewal of their historic connection to basic physics, as they have expanded to treat the intact transmission and processing of quantum states, and the interaction of such 'quantum information' with traditional forms of information. We may wonder why this did not happen earlier, as quantum principles have long been accepted as fundamental to all of physics. Perhaps the founders of information and computation theory, such as Shannon, Turing and von Neumann, were too accustomed to thinking of information processing in macroscopic terms, not yet having before them the powerful examples of the genetic code and ever-shrinking microelectronics. Be that as it may, information until recently has largely been thought of in classical terms, with quantum mechanics playing a supporting role in the design of the equipment to process it, and setting limits on the rate at which it could be sent through certain channels. Now we know that a fully quantum theory of information and information processing offers, among other benefits, a brand of cryptography whose security rests on fundamental physics, and a reasonable hope of constructing quantum computers that could dramatically speed up the solution of certain mathematical problems. These benefits depend on distinctively quantum properties such as uncertainty, interference and entanglement.

At a more fundamental level, it has become clear that an information theory based on quantum principles extends and completes classical information theory, just as complex numbers extend and complete the reals. Besides quantum generalizations of classical notions such as sources, channels and codes, the new theory includes two complementary, quantifiable kinds of information—

classical information and quantum entanglement. Classical information can be copied at will, but can only be transmitted forward in time, to a receiver in the sender's forward light cone. Entanglement in contrast, cannot be copied, but can connect any two points in space-time. Conventional data processing operations destroy entanglement, but quantum operations can create it and use it for various purposes, such as speeding up certain classical computations and assisting in the transmission of classical information or intact quantum states. Part of the new quantum information theory is the qualitative and quantitative study of entanglement, and its interactions with classical information.

Any means, such as an optical fibre, of delivering quantum systems more or less intact from one place to another, may be viewed as a quantum channel. Unlike classical channels, which are well characterized by a single capacity, quantum channels have several distinct capacities, depending on what one is trying to use them for, and what auxiliary resources are brought into play.

New effects involving quantum information continue to be discovered, not only in the traditional areas of computation, channel capacity, and cryptography, but in areas such as communication complexity and game theory.

Theory of quantum data and data processing

Quantum data. How, then, does quantum information, and the operations that can be performed on it, differ from conventional digital data and data-processing operations? A classical bit (such as a memory element or a wire carrying a binary signal) is generally a macroscopic system, and is described by one or more continuous parameters such as voltages. Within this parameter space

Table 1 Comparison of classical and quantum information processing

Property	Classical	Quantum
State representation	String of bits $x \in \{0, 1\}^n$	String of qubits $\psi = \sum_i c_i x_i\rangle$
Computation primitives	One- and two-bit boolean operations	One- and two-qubit unitary transformations
Fault-tolerant computation	By classical fault-tolerant gate arrays	By quantum fault-tolerant gate arrays
Quantum computational speed-ups		Factoring: exponential speed-up; search: quadratic speed-up; iteration, parity: no speed-up; simulation of quantum systems: up to exponential speed-up
Communication primitives	Transmitting a classical bit	Transmitting a classical bit; transmitting a qubit; sharing an EPR pair
Noiseless coding techniques	Classical data compression	Quantum data compression; entanglement concentration
Error-correction techniques	Error-correcting codes	Quantum error-correcting codes; entanglement distillation
Noisy-channel capacities	Classical capacity C_1 equals maximum mutual information through a single channel use	Classical capacity $C \geq C_1$; unassisted quantum capacity $Q \leq C$; classically assisted quantum capacity $Q_2 \geq Q$; entanglement-assisted classical capacity $C_E \geq C$
Entanglement assisted communication		Superdense coding, quantum teleportation
Communication complexity	Bit communication cost of distributed computation	Qubit cost, or entanglement-assisted bit cost, can be less
Secret cryptographic key agreement	Known protocols insecure against quantum computer	Secure against general quantum attack and unlimited computing
Two-party bit commitment	Known protocols insecure against quantum computer	Insecure against attack by a quantum computer

two well-separated regions are chosen by the designer to represent 0 and 1, and signals are periodically restored toward these standard regions to prevent them from drifting away due to environmental influences, crosstalk, and finite manufacturing tolerances. An n -bit memory can exist in any of 2^n logical states, labelled 000...0 to 111...1. Besides storing binary data, classical computers manipulate it; a sequence of boolean operations (for example, NOT and AND) acting on the bits one or two at a time is sufficient to realize any deterministic transformation.

A quantum bit or 'qubit' in contrast, is typically a microscopic system, such as an atom or nuclear spin or photon. The boolean states 0 and 1 are represented by a fixed pair of reliably distinguishable states of the qubit (for example, horizontal and vertical photon polarizations: $|0\rangle = \leftrightarrow$, $|1\rangle = \updownarrow$). A qubit can also exist in a continuum of intermediate states or 'superpositions', represented mathematically as complex linear combinations of the basis states $|0\rangle$ and $|1\rangle$. For photons, these states correspond to other polarizations, for example $\nearrow = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, $\nwarrow = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $\curvearrowright = \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$ (right circular polarization). Unlike the intermediate states of a classical bit (such as voltages between the standard 0 and 1 values), these intermediate states cannot be reliably distinguished, even in principle, from the basis states. With regard to any measurement, the superposition $\alpha|0\rangle + \beta|1\rangle$ behaves like $|0\rangle$ with probability $|\alpha|^2$ and like $|1\rangle$ with probability $|\beta|^2$. More generally two quantum states are reliably distinguishable if and only if their vector representations are orthogonal; thus \leftrightarrow and \updownarrow are reliably distinguishable by one type of measurement, and \nearrow and \nwarrow by another, but no measurement can reliably distinguish \leftrightarrow from \nearrow .

A pair of qubits (for example, two photons in different locations) is capable of existing in four boolean states, $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$, as well as all possible superpositions of them. These include states such as

$$\frac{1}{\sqrt{2}}(|00\rangle + |01\rangle) = |0\rangle \otimes \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) = \leftrightarrow \nearrow \quad (1)$$

which is describable as a tensor product of states of the individual photons, as well as states such as $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ which admit no such description. Such 'entangled' states correspond to a situation in which neither photon by itself has a definite state, even though the pair together does.

More generally, where a string of n classical bits could exist in any of 2^n boolean states $x = 000...0$ through $111...1$, a string of n qubits can exist in any state of the form

$$\Psi = \sum_{x=00...0}^{11...1} c_x |x\rangle \quad (2)$$

where the c_x are complex numbers such that $\sum_x |c_x|^2 = 1$. In other words, a quantum state of n qubits is represented by a complex vector Ψ of unit length in a space ('Hilbert space') of 2^n dimensions, one for each possible classical state. The exponentially large dimensionality of this space distinguishes quantum computers from classical analogue computers, whose state is described by a number of parameters that grows only linearly with the size of the system. This is because classical systems, whether digital or analogue, can be completely described by separately describing the state of each part. The vast majority of quantum states, by contrast, are entangled, and admit no such description. The ability to preserve and manipulate entangled states is the distinguishing feature of quantum computers, responsible both for their power and for the difficulty of building them.

An isolated quantum system evolves in such a way as to preserve superpositions and distinguishability; such evolution, called 'unitary', is the Hilbert-space analogue of rigid rotation in real space, and is another important difference between quantum and analogue systems. Unitary evolution and superposition are the central principles of quantum mechanics.

Logical operations. Just as any classical computation can be expressed as a sequence of one- and two-bit operations (for example, NOT and AND gates), any quantum computation can be expressed as a sequence of one- and two-qubit quantum gates, that is, unitary operations acting on one or two qubits at a time¹ (compare with Fig. 1). The most general one-qubit gate is described by a 2×2 unitary matrix $\begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix}$ mapping $|0\rangle$ to $\alpha|0\rangle + \beta|1\rangle$ and $|1\rangle$ to $\gamma|0\rangle + \delta|1\rangle$. One-qubit gates can easily be implemented physically, for example, by quarter- and half-wave plates acting on polarized photons, or by radio-frequency tipping pulses acting on nuclear spins in a magnetic field.

The standard two-qubit gate is the controlled-NOT or XOR gate, which flips its second (or 'target') input if its first ('control') input is $|1\rangle$ and does nothing if the first input is $|0\rangle$. In other words it interchanges $|10\rangle$ and $|11\rangle$ while leaving $|00\rangle$ and $|01\rangle$ unchanged. Unlike one-qubit gates, two-qubit gates are difficult to realize in the laboratory, because they require two separated quantum information carriers to be brought into strong and controlled interaction.

The XOR gate is a prototype interaction between two quantum systems, and illustrates several key features of quantum information, in particular the impossibility of cloning an unknown quantum state, and the way interaction produces entanglement. If the XOR is applied to boolean data in which the second qubit is 0 and the first is 0 or 1, the effect is to leave the first qubit unchanged while the second becomes a copy of it: $U_{\text{XOR}}|x, 0\rangle = |x, x\rangle$ for $x = 0$ or 1 . One might suppose that the XOR operation could also be used to copy superpositions, such as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, so that $U_{\text{XOR}}|\psi, 0\rangle$ would yield $|\psi, \psi\rangle$, but this is not so. The unitarity of quantum evolution requires that a superposition of input states evolve to a corresponding superposition of outputs. Thus the result of applying U_{XOR} to $|\psi, 0\rangle$ must be $\alpha|0, 0\rangle + \beta|1, 1\rangle$, an entangled state in which neither output qubit alone has a definite state. If one of the entangled output qubits is lost (for example, discarded, or allowed to escape into the environment), the other thenceforth behaves as if it had acquired a random classical value 0 (with probability $|\alpha|^2$) or 1 (with probability $|\beta|^2$). Unless the lost output is brought back into play, all record of the original superposition $|\psi\rangle$ will have been lost. This behaviour is characteristic not only of the XOR gate but of unitary interactions generally: their typical effect is to map most unentangled initial states of the interacting systems into entangled final states, which from the viewpoint of either system alone causes an unpredictable disturbance.

Environmental interactions. Since quantum physics underlies classical, there should be a way to represent classical data and operations within the quantum formalism. If a classical bit is a qubit having the value $|0\rangle$ or $|1\rangle$, a classical wire should be a wire that conducts $|0\rangle$ and $|1\rangle$ reliably, but not superpositions. This can be implemented using the XOR gate as described above, with an initial $|0\rangle$ in the target position which is later discarded. In other words, from the viewpoint of quantum information, classical communication is an irreversible process in which the signal interacts *en route* with an environment in such a way that boolean signals pass through undisturbed, but other states suffer entanglement with the environment. If the environment is lost or discarded, the surviving signal behaves as if it had been irreversibly forced to choose one of the boolean states. Not only a classical wire, but any classical data processing, can be realized similarly by quantum processing supplemented by interaction with a quantum environment that is later discarded.

Paradoxically, entangling interactions with the environment are thought to be the main reason why the macroscopic world seems to behave classically and not quantum-mechanically². Macroscopically different states, for example, the different charge states representing 0 and 1 in a VLSI (very large scale integration) memory cell, interact so strongly with their environment that information rapidly leaks out as to which state the cell is in. Therefore, even if it were possible to prepare the cell in a superposition of 0 and 1, the superposition

would rapidly evolve into a complex entangled state involving the environment, which from the viewpoint of the memory cell would appear as a statistical mixture, rather than a superposition, of the two classical values. The spontaneous decay of superpositions into mixtures is known as decoherence.

Entanglement with the environment is thus a major obstacle to quantum computation. To avoid having a quantum computation decohere into a probabilistic classical computation (which could just as well be done on a classical computer) it is necessary, while creating and maintaining entanglement among the computational degrees of freedom, to avoid entanglement between them and the environment. Until recently it appeared that the feasible number of steps in a coherent quantum computation would necessarily be less than the ratio τ_d/τ_s of decoherence time to switching time characteristic of the elementary quantum systems used in the hardware. Even if all other problems in the design of a practical quantum computer could be overcome, currently attainable values of τ_d/τ_s are not high enough to make quantum computers competitive with classical ones; also, the search for systems with ever-higher τ_d/τ_s might ultimately be blocked by fundamental properties of available atoms and nuclei. Apart from decoherence, it also appeared that individual gate operations would have to be made more and more precise the longer the computation.

This pessimism has largely been dispelled by the discovery of quantum fault-tolerant computation³ (QFTC), the quantum analogue of von Neumann's discovery that unreliable classical gates can

be used to perform arbitrarily long classical computations reliably, provided the error probability per gate is less than some constant threshold. Because of QFTC, it appears that experimentalists need 'only' build quantum hardware with a per-gate decoherence that is below some finite threshold (variously estimated at 10^{-6} to 10^{-2} , with a similar precision for individual gate rotations) in order for quantum computers to do arbitrarily complex computations.

With this background we survey some of the main parallels and differences between quantum and classical information processing. **Quantum speed-up of classical computation.** This is potentially the most important application of quantum data processing. By using quantum gates and wires, with entangled states flowing through them in the intermediate stages of a computation, certain computations mapping classical inputs x to classical outputs $f(x)$ can be done in far fewer steps than by any known sequence of classical gate operations. Most famously, a quantum computer can factor large integers in time that is polynomial in the logarithm of the best classical time^{4,5}, thereby threatening the security of cryptosystems based on the presumed difficulty of factoring. This exponential speed-up depends on the quantum computer's ability to vastly parallelize the performance of a fast Fourier transform, using destructive interference among a number of parallel computation paths that increases exponentially with the number of physical qubits involved in the computation. Another class of problems for which quantum computers seem to provide exponential speed-up is the simulation of many-particle quantum systems^{6,7}. In contrast to these rather specialized problems, a much broader class of problems can be speeded up quadratically, that is, solved in a time proportional to the square root of the time that a classical computer would require. These include search and optimization problems (for example, given an algorithm for computing a function F , find an input s where $F(s) = 0$, or an input s where $F(s)$ is a minimum)^{8,9}. For some other problems there is no quantum speed-up. These include iterated function evaluation^{10,11} (for example, given an algorithm for computing F , compute the n th iterate $F^{(n)} = F(F(\dots))$ for large n) and computing the parity of a random set^{12,13}.

Quantum information theory. This generalizes the classical notions of source and channel, and the related techniques of source and channel coding, as well as introducing a new resource, entanglement, which interacts with classical and quantum information in a variety of ways that have no classical parallel.

As mentioned earlier, quantum channels have several distinct capacities, depending on what one is trying to use them for, and what auxiliary resources are brought into play. These include the following:

Classical capacity, C , equal to the maximum rate at which classical bits can be transmitted reliably through the channel;

Quantum capacity, Q , the maximum rate at which intact qubits can be transmitted reliably through the channel;

Classically-assisted quantum capacity, Q_2 , defined as the maximum rate at which qubits can be transmitted reliably through the channel, with the help of unlimited two-way classical communication between sender and receiver; and

Entanglement-assisted classical capacity, C_E , defined as the maximum rate for sending classical bits through the channel, with the help of unlimited prior entanglement between sender and receiver. These capacities obey the relation $Q \leq Q_2 \leq C \leq C_E$ for all known channels, but otherwise appear to vary rather independently, and are not easy to calculate from the quantum channel parameters, again, unlike the single capacity of classical channels.

Quantum data compression and error correction. The two central techniques of classical information theory, source and channel coding, have direct quantum analogues; a quantum source is an entity that emits quantum states ψ_i with probabilities p_i , and a channel is an entity, such as an optical fibre, that transmits quantum states more or less reliably from a sender to a receiver.

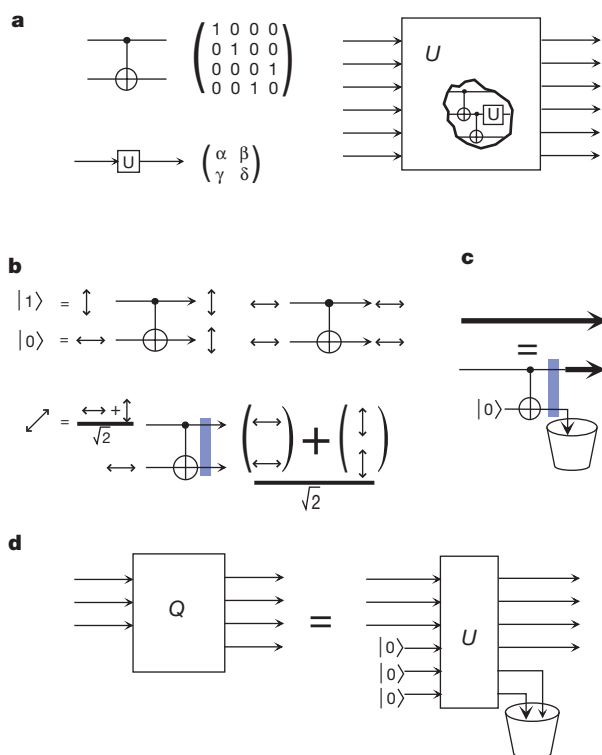


Figure 1 Quantum logical operations. **a**, Any unitary operation U on quantum data can be synthesized from the two-qubit XOR or controlled-NOT gate, and one-qubit unitary operations U . **b**, The XOR acts as a classical cloner on boolean valued inputs, but if one attempts to clone intermediate values, the cloning fails and an entangled state (blue) results instead. **c**, A classical wire (thick line) conducts 0 and 1 faithfully but not superpositions or entangled states. It may be defined as a quantum wire that interacts (via an XOR) with an ancillary 0 qubit which is then discarded. **d**, The most general treatment, or superoperator, Q that can be applied to quantum data is a unitary interaction with one or more 0 qubits, followed by discarding some of the qubits. Superoperators are typically irreversible.

The von Neumann entropy of a quantum source, $S = -\text{Tr}(\rho \log_2 \rho)$, where $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$, determines the minimum asymptotic number of qubits into which its signals can be compressed by a quantum encoder and still be faithfully recovered by a quantum decoder. This is the analogue of classical data compression or source coding, by which redundant classical data is compressed and faithfully regenerated, but quantum data compression¹⁴ differs in that it can be applied to non-orthogonal states (for example, equally probable horizontal and diagonal photons, as shown in Fig. 2a) which would be spoiled if one tried to compress them classically. Also, because the states are non-orthogonal, the encoder cannot retain a copy of them, or indeed any memory of them, if they are to be faithfully reconstructed at the receiving end. A quantum encoder is like a discreet telegrapher, who transmits messages without remembering them.

Source coding removes redundancy, allowing data to be sent more efficiently through a noiseless channel. Error-correction or channel coding, in contrast, introduces redundancy to enable data to withstand transmission through a noisy channel. The simplest classical error-correcting code is the triple repetition code $0 \rightarrow 000$, $1 \rightarrow 111$, which permits the encoded bit to be faithfully recovered after up to one transmission error in the three-bit codeword. Analogous error-correcting codes exist for quantum data, but they require more redundancy because they need to protect not only boolean states, but also arbitrary superpositions of them^{15–20}. Thus the simplest single-error-correcting quantum code (Fig. 2b) encodes an arbitrary input qubit $|\psi\rangle$ into an entangled state of five qubits, in such a way that if any one is corrupted *en route*, the decoder can funnel the effects of the error into the four ancillary qubits, while restoring the first qubit to its original state. Analogously to classical capacity, the quantum capacity Q of a noisy channel can be defined as the limiting ratio of faithfully transmitted qubits per noisy-channel use that is achievable by quantum error-correcting codes. This quantum capacity is usually less, and can never be greater, than the same channel's capacity C for transmitting classical bits. The inequality $Q \leq C$ holds for all channels because if a channel can faithfully transmit a general qubit, then it can certainly transmit the particular qubits $|0\rangle$ and $|1\rangle$.

The discovery of quantum error-correcting codes in 1995 came as a great surprise, probably because people were used to thinking of classical error correction in language unsuited to quantum generalization. For example, triple repetition, if it is taken to mean making three copies of the input qubit ($\psi \rightarrow \psi \otimes \psi \otimes \psi$), flies in the face of the well-known impossibility of exactly copying ('cloning') an unknown quantum state. In retrospect, the natural quantum generalization of triple repetition can be seen instead to be the mapping $\alpha|0\rangle + \beta|1\rangle \rightarrow \alpha|000\rangle + \beta|111\rangle$, which does not violate any quantum principle and indeed suffices to correct single qubit errors in any boolean input. As noted above, two more bits of redundancy are needed to extend the protection to non-boolean inputs. Although analogous in structure to classical discrete error-correcting codes, quantum error-correcting codes have the remarkable ability to protect a continuum of inputs from a continuum of errors. For example, in Fig. 2b, the input qubit might be a photon of any polarization state, and the error (red) might be a rotation of one of the five channel qubits' polarizations by an arbitrary amount; nevertheless, the error would be corrected. This is a beneficial side effect of the linearity of quantum mechanics: if a quantum error-correcting code protects a sufficiently rich discrete set of inputs from a sufficiently rich discrete set of error processes, then it will also protect any superposition of those inputs from any superposition of those errors. Besides the simple capacities C and Q , quantum channels have assisted capacities Q_2 and C_E mentioned above, which will be discussed later.

The oldest branch of quantum information theory^{21–23} concerns the use of quantum channels to transmit classical information. Even the seemingly pedestrian classical capacity C is not easy to calculate

for quantum channels, because it may depend on using a quantum encoder to prepare inputs entangled over multiple uses of the channel, and/or a quantum decoder to perform coherent measurements on multiple channel outputs. Unlike any classical channel, some quantum channels are superadditive, in the sense that more classical information can be sent through n parallel uses of the channel than n times the amount that can be sent through one use of the channel^{24–28}.

Entanglement-assisted communication. Two forms of quantum information transmission that have no classical counterpart, but are closely related to each other, are quantum teleportation²⁹ (Fig. 3a) and quantum superdense coding³⁰ (Fig. 3b). These involve an initial stage in which a pair of particles in a maximally-entangled state such as $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ (often called an Einstein–Podolsky–Rosen or EPR pair) is shared between two parties, followed by a second stage in which this shared entanglement is used to achieve, respectively, transmission of a qubit via two classical bits, or transmission of two classical bits via one qubit. Quantum teleportation illustrates the fact that transmission of intact quantum states requires two qualitatively different resources: a quantum resource that cannot be cloned, and a directed resource that cannot travel faster than light. In direct transmission of a qubit, these two functions are performed by the same particle. In teleportation the former function is provided by the shared EPR pair, the latter by the two classical bits. This situation may be summarized by saying that classical information theory involves one species of information, and one kind of noiseless communication primitive (transmission of a bit), whereas quantum information theory involves two species (classical information and entanglement), and three primitives (transmitting a bit, transmitting a qubit, and sharing an EPR pair) which are related through superdense coding and teleportation.

Superdense coding (Fig. 3b) is an example of entanglement-assisted classical communication, and shows that $C_E = 2$ for the noiseless qubit channel, while $C = Q = 1$. Surprisingly, the ratio

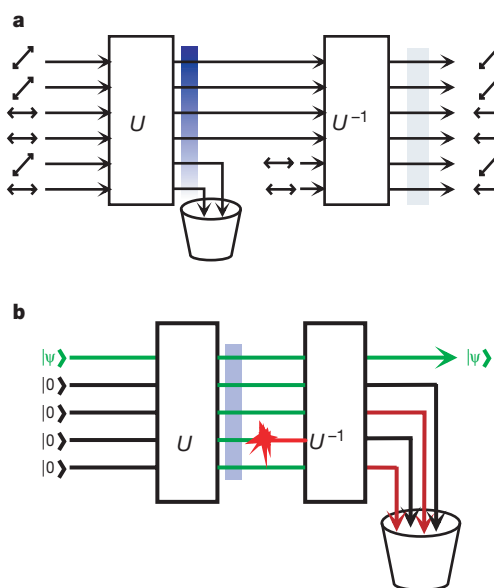


Figure 2 Quantum data compression and error correction. **a**, In quantum data compression, inputs from a redundant source (here, an unknown sequence of horizontal and diagonal photons) are unitarily transformed into an entangled state (blue) in which almost all the information has been concentrated into some of the photons, allowing the others to be discarded. At the receiving end of the channel the discarded photons are replaced by standard (horizontal) photons and the unitary transformation is undone, resulting in a close approximation to the original state. **b**, A quantum error-correcting code with unitary encoder and decoder. An arbitrary input qubit $|\psi\rangle$ is entangled with four standard $|0\rangle$ qubits in such a way that if any one of the five qubits is spoiled, the decoder can still restore the original state exactly.

C_F/C typically increases with increasing noise, and indeed can attain arbitrarily large values for channels so noisy that their quantum capacities Q and Q_2 both vanish³¹. Thus, unlike most quantum effects, entanglement enhancement of a quantum channel's classical capacity does not disappear in the limit of large noise. In this respect it resembles the ability of bulk nuclear magnetic resonance systems to carry out nontrivial quantum computations while remaining close to thermal equilibrium.

Superdense coding and teleportation have received much laboratory attention recently. The first work was by the Innsbruck group³², which implemented a version of superdense coding in which three distinguishable states (rather than the theoretical maximum of four) are created by manipulating one member of an EPR pair of polarization-entangled photons. Teleportation using these photon states was more recently achieved by the same group³³; by using these techniques several other protocols involving entanglement, for example, the creation of three-particle entanglement, has become possible. An experimentally different approach in which another attribute of one of the EPR photons (such as its position) is teleported has been implemented in Rome³⁴. This experiment is easier in that it involves two rather than three photons. A very recent experiment³⁵ has followed more directly a version of a teleportation scheme due to Vaidman³⁶, in which continuous quantum degrees of freedom are teleported. This work demonstrates that an arbitrary quantum state of one optical mode can be teleported with good fidelity; taking into account limitations on the intensity of the mode, one finds that roughly a one-million-state quantum system is involved, in contrast to the two-state systems used in the earlier work. Finally, the operations necessary to teleport a nuclear spin state have been performed using nuclear magnetic resonance³⁷, although the range of teleportation is only the distance across a single molecule.

Although entanglement by itself cannot be used to transmit a classical message, it can reduce the amount of classical communication required to perform a distributed computation^{13,38,39}. Classically, 'communication complexity' refers to the amount of communication needed to evaluate a function of several inputs in remote locations. For example, if Alice and Bob each have appointment calendars with n time slots, $O(n)$ bits of communication are required to determine if there is a time when they are both free. If they are allowed to share prior entanglement, or if they are allowed to communicate using qubits rather than bits, the communication complexity of this problem is reduced from $O(n)$ to $O(\sqrt{n} \log n)$.

Quantifying and distilling entanglement. Because of its usefulness in protocols such as teleportation, it is important to have quanti-

tative measures of entanglement, and to know whether all entangled states (those not expressible as products of states of their parts, or probabilistic mixtures of such products) can be converted into EPR pairs, and if so, how efficiently. In the case of bipartite pure states, entanglement is naturally measured by the state's entropy of entanglement, the von Neumann entropy of either subsystem considered alone. For such states^{40–42} the entropy of entanglement $E(\Psi)$ is equal both to the state's entanglement of formation—the number of EPR pairs asymptotically required to prepare one instance of the state by classical communication and local operations—and its distillable entanglement—the number of pure EPR pairs that can be asymptotically prepared from one instance of the state by classical communication and local operations.

For mixed states, and states of three or more parties, the situation is more complicated, and there are several non-equivalent kinds of entanglement. Multipartite states, pure and mixed, have been studied^{43–45}. Mixed states generally have a distillable entanglement that is less than their entanglement of formation, reflecting the irreversibility of the mixing process. An extreme form of this phenomenon is the existence of so-called "bound" entangled states⁴⁶—mixed states which are entangled, but from which no pure entanglement can be distilled.

Entanglement distillation is important not only for quantifying entanglement but as a distinctively quantum kind of error correction, complementary to the use of quantum error-correcting codes^{20,47,48}. Suppose Alice and Bob can communicate classically, and in addition have access to a noisy quantum channel. Now Alice wants to send an unknown qubit reliably to Bob. If the quantum channel is not too noisy, she can encode the input qubit into several qubits using a quantum error-correcting code as in Fig. 2b, send these through the noisy channel, and have Bob decode them. However for very noisy channels, such as a 50% depolarizing channel, this will not work, because such channels have zero quantum capacity $Q = 0$. In this case the best known strategy is for Alice not to send the input qubit through the channel at all, but instead prepare a number of pure EPR pairs, and share them through the noisy channel with Bob (resulting in noisy EPR pairs). Then, using their ability to communicate classically, Alice and Bob distill a smaller number of good EPR pairs from the larger supply of noisy ones. Finally, Alice uses one of the good EPR pairs, and additional classical communication, to teleport the input qubit safely to Bob. The ability of entanglement distillation to salvage such noisy EPR pairs gives rise to a fact noted earlier, that for many channels the classically-assisted quantum capacity Q_2 exceeds the direct quantum capacity Q . (However, this advantage depends on two-way classical communication between Alice and Bob—if they are limited to one-way communication, distillation is no more efficient than quantum error-correcting codes.) As a function of increasing noise, a typical quantum channel passes through two thresholds; a noise level beyond which Q vanishes but Q_2 and C remain positive; and a threshold beyond which Q_2 vanishes but C remains positive.

Quantum fault-tolerant computation. QFTC is both an expansion of research in the theory of quantum information processing and a practical necessity for implementing non-trivial quantum computation in the laboratory. Modern QFTC has been well reviewed (see ref. 3 and references therein); some of the basic ideas are sketched in Fig. 4. To avoid irrecoverable damage from a single error, an appropriate quantum error-correcting code is used to spread the logical state $|\psi_L\rangle$ being stored or processed over several physical qubits, carried by a bundle of parallel wires. Periodically the bundle passes through a restorative gate array R , where it interacts with clean ancilla qubits from the environment, in order to correct the errors by funnelling them into the ancillas, which are then discarded. Additional errors may occur during the restoration process itself, but if these are not too numerous they will be corrected by a subsequent restoration step (Fig. 4a). Such a regimen of active




Figure 3 Quantum information transmission between a sender (Alice) and a receiver (Bob). **a**, In quantum teleportation, prior sharing of an EPR pair, and transmission of a two-bit classical message from Alice to Bob suffice to transmit an unknown quantum state even when no direct quantum channel from Alice to Bob is available. **b**, In quantum dense coding, prior sharing of an EPR pair, and transmission of a single qubit from Bob to Alice, suffice to transmit an arbitrary two-bit classical message (x, y).

restoration can be used to implement a fault-tolerant quantum memory, able to hold quantum states reliably for much longer than the natural decoherence times of the hardware of which the array is built. Apart from the fact that it operates on quantum rather than classical data, this is entirely analogous to devices such as the dynamic random access memory (DRAM) used in today's computers, in which periodic signal restoration serves to delay the decay of the stored data almost indefinitely. To perform quantum computation fault-tolerantly, it is also necessary, besides storing the data, to perform gate operations on it without decoding it from its protected form. For some gates, such as the XOR, this can be done in a straightforward fashion, by applying the gate operation successively to wires in a bundle (Fig. 4b). Other gate operations, including some necessary single-qubit rotations, must be implemented in a more complex fashion, involving preparation and testing of special entangled states of a set of ancilla qubits, which are then brought into interaction with the encoded data to perform the desired logical transformation³.

The promise of quantum computation lies in the fact that, to perform a t -step computation fault-tolerantly, the number of gates and wires need only be multiplied by a factor polynomial in $\log t$. Therefore, for computations with a significant quantum speed-up, a quantum computer would still vastly outperform any classical computer on sufficiently large inputs.

Quantum cryptography. This is the art of applying the unique properties of quantum systems to cryptographic goals, that is, the protection of classical information from tampering or unauthorized disclosure in a multi-party setting where not all the parties trust one another. This adversarial element distinguishes it from the kinds of quantum information processing considered earlier.

One important quantum cryptographic task, quantum key distribution has as its goal the sharing of a secret random bit string K , called a cryptographic key, between the two protagonists Alice and Bob, who have at their disposal an insecure quantum channel and a public classical channel. (Purely classical protocols for key agreement exist and are in widespread use, but these result in a key that is not informationally secure—an adversary with sufficient computing power could infer it from the public messages exchanged between Alice and Bob. In particular, the most widely used classical key agreement protocols could be easily broken by a quantum computer, if one were available.) In quantum key distribution, an eavesdropper ('Eve') is allowed to interact with the quantum information carriers (for example, photons) *en route* from Alice to Bob—at the risk of disturbing them—and can also passively listen to all classical communication between Alice and Bob, but she cannot alter or suppress the classical messages. Sometimes (for example, if Eve jams or interacts strongly with the quantum signals) Alice and Bob will detect the excessive eavesdropping and abort the protocol; but, for every eavesdropping strategy, Eve's probability of remaining undetected and obtaining significant information on the key should be negligible.

The practical implementation of quantum key distribution is much farther advanced than other kinds of quantum information processing, owing to the fact that the standard quantum key distribution protocols require no two-qubit interactions, only preparation and measurement of simple quantum states, along with classical communication and computations. Optical prototypes working over tens of kilometres of fibre, or even through a kilometre of open air at night, have been built and tested. In principle, however, a quantum key distribution protocol could involve quantum computations by Alice and Bob; and to be sure of its security, one ought to allow Eve the full power of a quantum computer, even though Alice and Bob do not need one for the standard protocols.

Various proofs of security of quantum key distribution protocols, especially the four-state "BB84" protocol of ref. 49, have been offered. A complete security proof should encompass all attacks

allowed by the laws of quantum mechanics, and should also be able to cope with noise under the realistic assumption that it arises not only from eavesdropping but also from noisy channels and detectors. Finally, it should provide a way of calculating a safe rate of key generation as a function of the noise level observed by Alice and Bob. Recent proofs^{50,51} building on a long history of previous security proofs against more limited attacks^{48,49,52,53}, have largely met these criteria, the main remaining problems being to simplify the proofs, improve the error thresholds, and to extend them to cover realistic sources, which do not emit exact single-photon states or exact EPR pairs, and in extreme cases may even have been sabotaged by Eve.

Given the success of quantum key distribution, there was high hope that quantum techniques could help with another task, two-party oblivious function evaluation, a better name for which might be "discreet decision-making". This is the task, which arises frequently in commerce and diplomacy, of enabling two mutually distrustful parties to cooperate in evaluating a publicly agreed function of private data held separately by each party, without compromising the private data any more than it would have been compromised had they assigned the job of evaluating the function to a trusted intermediary. Initially Alice knows data x and Bob knows data y ; when the protocol is finished, Alice and Bob should each also know $f(x,y)$, but neither party should know any more about the other party's private input than can logically be inferred from a knowledge of their own data and the common function value $f(x,y)$. Classical protocols for oblivious function evaluation exist, but, like classical key agreement protocols, they are not informationally secure and could be broken by a quantum computer. Hopes for finding a quantum basis for absolutely secure oblivious function evaluation were dashed by the discovery that a fundamental building block of all known oblivious function evaluation protocols, called bit commitment, is insecure in principle against quantum attacks^{54,55}. Bit commitment is the idealization of a protocol in which Alice sends Bob a locked box containing a bit 0 or 1 of her choosing, written on a piece of paper, then later, at a time of her choosing, sends him the key so that he can open the box and then read the bit. Quantum bit commitment is insecure because of a fundamental property of entangled states, namely that if two pure states of the Alice–Bob system are indistinguishable to Bob, they must be interconvertible by a local action of Alice; thus there is in principle no way of implementing a locked box containing a bit value that is both unmodifiable by Alice and unobservable by Bob.

The similarities and differences between classical and quantum information are summarized in Table 1.

Experimental studies of quantum information

The continuing maturation of the theory of quantum information and quantum computation has stimulated experimental work in a great variety of disciplines, in optics and quantum optics, in single-atom and single-ion research, and in several areas of precision spectroscopy. We will touch on some of the progress along these lines here. We will not mention here the very interesting prospects of using solid-state quantum technology—quantum dots, semiconductor microcavities, ultrasmall Josephson junctions, and so forth—to achieve quantum gate operation, which apparently lie several years further into the future.

'Flying qubits' will be needed to implement many of the quantum processing protocols described above. Because of developments in quantum cryptography, high-quality flying qubits in the form of photons travelling on optical fibres are now produced routinely in several laboratories. An important innovation, introduced by the Gisin group at the University of Geneva⁵⁶, helps to make reliable photon transmission through unreliable fibres a possibility. It involves the use of a Faraday mirror, which reflects any light that strikes it into an orthogonal polarization. In their scheme a strong coherent-state double pulse of light is sent from Bob to Alice on an

optical fibre; Alice attenuates it to one-photon intensity, sets the relative phase of the two pulses to obtain one of four quantum states of the photon, and finally Faraday-reflects this photon back to Bob. The Faraday reflection ensures that any distortions or variations of the propagating light mode due to birefringence (anisotropy of the index of refraction) in the photon transmission from Bob to Alice are undone in the return transmission. With this invention a remarkable interferometric stability is attained: the fringe visibility of their 23-km transmission system used as an interferometer has reached 99.98%, implying that the phase of the photon is reliable to 0.03 radians. This means that high-purity quantum states of light are being successfully transmitted in this system.

Their ability to store and process qubits with 'standing qubits' can greatly augment the capabilities of 'flying qubits' for the processing of quantum information. For example, the ability to do quantum computation in conjunction with quantum communication would qualitatively enhance the ability to do quantum cryptography, permitting the use of quantum repeater elements and opening up new techniques for defeating eavesdropping, as well as permitting cryptography over indefinitely long distances^{57,58}. With this in mind, workers have proposed a marriage of techniques from photon-fibre systems and trapped-atom (or ion) systems. In these schemes^{60,61}, a 'standing qubit' encoded in a state of the atom is mapped by an appropriate laser pulse⁵⁹ into the same qubit state of the photon state of a surrounding electromagnetic cavity, and can from there become a 'flying qubit' by leaking out into a propagating mode in free space or in an optical fibre. The unexpected feature of this procedure is the next step, in which the propagating photon then impinges on a replica of the sending system. If the photon wave packet has been tailored properly, this cavity-atom system can be made to recapture the photon into the atomic state by a suitable time-reverse of the sending procedure.

Although extensions of the original two-bit gate demonstration in cavity quantum electrodynamics⁶² have brought us closer to this goal of marrying flying and standing qubits we do not yet have an

elementary functioning prototype. Optical quantum electrodynamics (QED) experiments have not succeeded in entangling the states of two 'standing qubits', but such entanglement has been achieved in experiments in related areas, in microwave cavity QED (ref. 63) and in ion-trap studies⁶⁴.

Unfortunately, the controllable creation of entanglement with two-qubit quantum gates is only one of a formidable checklist of ingredients that a physical experiment must have if it is to realize a quantum computer. There are at least four other milestones which must be achieved⁶⁶: (1) The system should be extendible to a large number of qubits. (2) It must be possible to place the qubits reliably in the "0" or cleared state at the outset. (3) The decoherence rate must, as explained above, be very low (that is, below some suitable threshold). (4) It must be possible to do single-quantum sensitivity measurements (if only one copy of the quantum computer is available) or an accurate ensemble measurement in a qubit-specific fashion (if many copies of the quantum computer are available).

A full-scale experiment of any type to realize all of these criteria simultaneously is still a long way off. In the area of ion-trap research, concerted efforts are being undertaken by a number of experimental groups to realize the original Cirac and Zoller proposal⁶⁷ for ion-trap quantum computing, which created great excitement and interest five years ago. The proposal of these authors was nothing less than a scheme for realizing all of the requirements for quantum computation mentioned above: qubits are to be represented by the internal (spin) states of individual ions held in the electromagnetic trap; extending the number of qubits is to be achieved by adding more atoms to the trap. The techniques of laser cooling would serve to put the system in the "0" state. Coupling to the environment in the ion trap is low, and thus qubits with acceptable decoherence properties are known. The technique known as quantum-jump spectroscopy provides for the possibility of virtually single-quantum measurements of almost 100% efficiency. The heart of the proposal is a detailed scheme for the realization of two-qubit quantum operations: their procedure involves a coupling of the internal ion state with the quantum state of vibration of the ions in the trap. Because these oscillations involve collective modes of all the ions, entanglement of the internal ion states becomes possible. Unfortunately, one feature of the Cirac-Zoller computer—cooling to the ground state of motion of the trap—has proved to be very difficult to achieve, and this essential step has only been accomplished reliably by one group for one⁶⁵ or two⁶⁴ ions.

While it is clear that the trap ideas are on a steady track of progress, naturally workers in other fields hope that their techniques will enable them to leapfrog the atomic physicists and get ahead in the 'quantum computer race'. The proponents of nuclear magnetic resonance spectroscopy (NMR), as practised in organic chemistry, have made a bold move in this direction. NMR spectroscopy has many useful features for application to quantum computation: in a well-understood limit of rapidly tumbling molecules in solution, the hamiltonian of the nuclear spins of the molecule assumes a very simple form:

$$H = \sum_i \omega_i \sigma_z^i + \sum_{i,j} J_{ij} \sigma_z^i \sigma_z^j \quad (3)$$

(Here σ_z is the angular momentum operator of the spins, ω is the Zeeman splitting, and J is the exchange interaction parameter.) This depends only on the z -component of the nuclear spins (labelled i and j here). A system with this hamiltonian is well adapted to quantum computing^{68,69}: as it commutes with all σ_z^i operators, every computational basis state is an eigenstate. Thus, the state of the system only changes when a resonance pulse is applied, so that the dynamics of the system is entirely under external control. By appropriate frequency and time selectivity, an external pulse can perform a very finely tailored operation, for example, flipping one particular spin i if another specific spin j is up; this is the essence of the fundamental two-qubit XOR gate for quantum computing. In




Figure 4 Fault-tolerant computation. **a**, Fault-tolerant error-correction circuit with cold ancillas coming in and corrupted ones being discarded. **b**, Performing the XOR operation on encoded data without decoding it.

addition, the pulse operations can be done much faster than the decoherence time of 1–10 seconds in a well-chosen molecule. Finally, in a situation in which many identical copies of the quantum computer are available (the many identical molecules in solution), the readout of the final result can be accomplished by an ensemble measurement of the transverse magnetization, a standard operation in NMR.

After the initial proposals of refs 70 and 71, there has been a flood of work on few-qubit systems, so numerous that we will just enumerate the accomplishments in this area briefly here: the action of two- and three-bit quantum gates has been demonstrated in the protons in 2,3-dibromothiophene and in 1-chloro-2-nitrobenzene⁷², the Deutsch–Jozsa algorithm⁷³ and the Grover algorithm⁷⁴ have been demonstrated using the H and ¹³C spins in chloroform, the three H and C spins in trichlorethylene have been used to simulate the synthesis of Greenberger–Horne–Zeilinger states⁷⁵, to perform teleportation³⁷ (see above), and to simulate the action of the three-qubit quantum error-correcting code⁷⁶ (the three C spins in alanine were also used in this last study), protons in cytosine have been used to implement the original Deutsch algorithm⁷⁷ as well as the quantum counting algorithm⁷⁸, and 2,3-dibromopropanoic acid has been used for some simple three-qubit gate arrays⁷⁹.

The NMR practitioners are pressing on to implement more quantum information processing in molecules with larger numbers of spins. However, there are a couple of large obstacles to immediately going on to large-scale quantum computation; probably these are not insuperable, but they may serve to make the progress of NMR quantum computing no faster than that in atomic physics or in other areas. One problem is just that the frequency-domain addressing which is used, in which each qubit has a distinct chemical shift ω_i , becomes difficult when the number of qubits grows large. A second problem (this will probably be the more immediate reason that the NMR technique will have to be radically modified to do quantum computation at a scale much greater than about 10 qubits) has to do with state preparation: the spin states of the molecules in the solution at room temperature are almost perfectly random, with a slight bias ϵ for the zero qubit state (typically ϵ , proportional to $k_B T$ divided by the nuclear Zeeman energy, is of the order of 10^{-6}). The number of molecules in the solution starting in the correct state, rather than the completely random state, scales with $\epsilon 2^{-n}$, where n is the number of spins in the molecule. The signal strength thus becomes exponentially small in the number of qubits, and all advantage gained from doing quantum computation is lost. This problem can be solved if ϵ can be increased to nearly one; there are innovative techniques in optical pumping which hold out the hope of doing this. While there is reason for optimism in these areas, we think that it will require many years of concerted effort.

We recall⁵⁹ the incident at a quantum computation meeting in Torino in 1995, when Shor offered a bet that the first factoring of a 500-digit number would be accomplished by a quantum and not a classical computer. There were no takers on the other side, but some commented that they would prefer to bet on a third possibility, that the Sun would burn up first. Although these sceptics have not been entirely silenced, on balance we are more in agreement with Shor than we were then. We think the odds in favour of the quantum computer have improved and will continue to increase slowly as more years of steady progress are made. □

1. Barenco, A. *et al.* Elementary gates for quantum computation. *Phys. Rev. A* **52**, 3457–3467 (1995).
2. Zurek, W. Decoherence and the transition from quantum to classical. *Phys. Today* **44**, 36–44 (1991).
3. Preskill, J. Reliable quantum computers. *Proc. R. Soc. Lond. A* **454**, 385–410 (1998).
4. Shor, P. W. in *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science* 124–133 (IEEE Computer Society Press, Los Alamitos, California, 1994).
5. Ekert, A. & Jozsa, R. Shor's quantum algorithm for factorising numbers. *Rev. Mod. Phys.* **68**, 733–753 (1996).
6. Wiesner, S. Simulations of many-body quantum systems by a quantum computer. Preprint quant-ph/9603028 at (<http://xxx.lanl.gov>) (1996).

7. Abrams, D. S. & Lloyd, S. Simulations of many-body Fermi systems on a universal quantum computer. *Phys. Rev. Lett.* **79**, 2586–2589 (1997).
8. Grover, L. K. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.* **79**, 325–328 (1997).
9. Boyer, M., Brassard, G., Hoyer, P. & Tapp, A. Tight bounds on quantum searching. *Fortschr. Phys.* **46**, 493–506 (1998).
10. Ozhigov, Y. Quantum computers cannot speed up iterated applications of a black box. Preprint quant-ph/9712051 at (<http://xxx.lanl.gov>) (1997).
11. Terhal, B. M. *Quantum Algorithms and Quantum Entanglement*. Thesis, Univ. Amsterdam (1999).
12. Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. A limit on the speed of quantum computation in determining parity. *Phys. Rev. Lett.* **81**, 5442–5444 (1998).
13. Beals, R., Buhrman, H., Cleve, R., Mosca, M. & de Wolf, R. in *Proceedings of the 39th Annual Symposium on the Foundations of Computer Science* 352–361 (IEEE Computer Society Press, Los Alamitos, California, 1998).
14. Jozsa, R. & Schumacher, B. A new proof of the quantum noiseless coding theorem. *J. Mod. Opt.* **41**, 2343–2349 (1994).
15. Shor, P. W. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A* **52**, 2493–2496 (1995).
16. Calderbank, A. R. & Shor, P. W. Good quantum error correcting codes exist. *Phys. Rev. A* **54**, 1098–1105 (1996).
17. Steane, A. Multiple particle interference and quantum error correction. *Proc. R. Soc. Lond. A* **452**, 2551–2577 (1996).
18. Knill, E. & Laflamme, R. Theory of quantum error correcting codes. *Phys. Rev. A* **55**, 900–911 (1997).
19. Gottesman, D. A class of quantum error-correcting codes saturating the Hamming bound. *Phys. Rev. A* **54**, 1862–1868 (1996).
20. Bennett, C. H., DiVincenzo, D. P., Smolin, J. & Wootters, W. K. Mixed state entanglement and quantum error correction. *Phys. Rev. A* **54**, 3824–3851 (1996).
21. Helstrom, C. W. *Quantum Detection and Estimation Theory* (Academic, New York, 1976).
22. Kholevo, A. S. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii* **9**, 3–11 (1973); translated in *Problems Inf. Transmiss.* **9**, 177–183 (1973).
23. Holevo, A. S. Problems in the mathematical theory of quantum communication channels. *Rep. Math. Phys.* **12**, 273–278 (1977).
24. Schumacher, B., Westmoreland, M. & Wootters, W. K. Limitation on the amount of accessible information in a quantum channel. *Phys. Rev. Lett.* **76**, 3452–3455 (1997).
25. Holevo, A. S. The capacity of the quantum channel with general signal states. *IEEE Trans. Inf. Theory* **44**, 269–273 (1998).
26. Kholevo, A. S. Capacity of a quantum communications channel. *Problemy Peredachi Informatsii* **15**, 3–11 (1979); translated in *Problems Inf. Transmiss.* **15**, 247–253 (1979).
27. Sassaki, M., Kato, K., Izutsu, M. & Hirota, O. Quantum channels showing superadditivity in channel capacity. *Phys. Rev. A* **58**, 146–158 (1998).
28. Fuchs, C. A. Nonorthogonal quantum states maximize classical information capacity. *Phys. Rev. Lett.* **79**, 1162–1165 (1997).
29. Bennett, C. H. *et al.* Teleporting an unknown quantum state via dual classical and Einstein–Podolsky–Rosen channels. *Phys. Rev. Lett.* **70**, 1895–1898 (1993).
30. Bennett, C. H. & Wiesner, S. J. Communication via one- and two-particle operators on Einstein–Podolsky–Rosen states. *Phys. Rev. Lett.* **69**, 2881–2884 (1992).
31. Bennett, C. H., Shor, P. W., Smolin, J. A. & Thapliyal, A. V. Entanglement enhanced classical capacity of noisy quantum channels. *Phys. Rev. Lett.* **83**, 3081–3084 (1999).
32. Mattle, K., Weinfurter, H., Kwiat, P. G. & Zeilinger, A. Dense coding in experimental quantum communication. *Phys. Rev. Lett.* **76**, 4656–4659 (1996).
33. Bouwmeester, D. *et al.* Experimental quantum teleportation. *Nature* **390**, 575–579 (1997).
34. Boschi, D., Branca, S., De Martini, F., Hardy, L. & Popescu, S. Experimental realization of teleporting an unknown pure quantum state via dual classical and Einstein–Podolsky–Rosen channels. *Phys. Rev. Lett.* **80**, 1121–1124 (1998).
35. Furusawa, A. *et al.* Unconditional quantum teleportation. *Science* **282**, 706–709 (1998).
36. Vaidman, L. Teleportation of quantum states. *Phys. Rev. A* **49**, 1473–1476 (1994).
37. Nielsen, M. A., Knill, E. & Laflamme, R. Complete quantum teleportation using nuclear magnetic resonance. *Nature* **396**, 52–55 (1998).
38. Cleve, R. & Buhrman, H. J. Substituting quantum entanglement for communication. *Phys. Rev. A* **56**, 1201–1204 (1997).
39. Buhrman, H., Cleve, R. & Wigderson, A. in *Proceedings of the 39th Annual ACM Symposium on the Theory of Computing* 63–68 (ACM Press, New York, 1998).
40. Bennett, C. H., Bernstein, H. J., Popescu, S. & Schumacher, B. Concentrating partial entanglement by local operators. *Phys. Rev. A* **53**, 2046–2052 (1996).
41. Lo, H.-K. & Popescu, S. Concentrating entanglement by local actions—beyond mean values. Preprint quant-ph/9707038 at (<http://xxx.lanl.gov>) (1997).
42. Vidal, G. Entanglement monotones. Preprint quant-ph/9807077 at (<http://xxx.lanl.gov>) (1998).
43. Linden, N., Popescu, S. & Sudbury, A. Non-local properties of multi-partite density matrices. *Phys. Rev. Lett.* **83**, 243–247 (1999).
44. Thapliyal, A. V. On multipartite pure-state entanglement. *Phys. Rev. A* **59**, 3336–3342 (1999).
45. Kempe, J. On multi-particle entanglement and its application to cryptography. *Phys. Rev. A* **60**, 910–916 (1999).
46. Horodecki, M., Horodecki, P. & Horodecki, R. Mixed state entanglement and distillation: is there a 'bound' entanglement in nature? *Phys. Rev. Lett.* **80**, 5239–5242 (1998).
47. Bennett, C. H. *et al.* Purification of noisy entanglement, and faithful teleportation via noisy channels. *Phys. Rev. Lett.* **76**, 722–725 (1996).
48. Deutsch, D. *et al.* Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Phys. Rev. Lett.* **77**, 2818–2821 (1996); **80**, 2022 (1998) (errata).
49. Bennett, C. H. & Brassard, G. in *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing, Bangalore, India* 175–179 (IEEE, New York, 1984).
50. Mayers, D. Unconditional security in quantum cryptography. Preprint quant-ph/9802025 at (<http://xxx.lanl.gov>) (1998).
51. Lo, H.-K. & Chau, H. F. Unconditional security of quantum key distribution over arbitrarily long distances. *Science* **283**, 2050–2056 (1999).

52. Griffiths, R. B. & Niu, C.-S. Optimal eavesdropping in quantum cryptography. II. Quantum circuit. *Phys. Rev. A* **56**, 1173–1176 (1997).
53. Biham, E., Boyer, M., Brassard, G., van de Graaf, J. & Mor, T. Security of quantum key distribution against all collective attacks. *Phys. Rev. Lett.* **78**, 2256–2259 (1997).
54. Mayers, D. Unconditionally secure quantum bit commitment is impossible. *Phys. Rev. Lett.* **78**, 3414–3417 (1997).
55. Lo, H.-K. & Chau, H. F. Is quantum bit commitment really possible? *Phys. Rev. Lett.* **78**, 3410–3413 (1997).
56. Muller, A. *et al.* 'Plug and Play' systems for quantum cryptography. *Appl. Phys. Lett.* **70**, 793–795 (1997).
57. Briegel, H. J., Dür, W., Cirac, J. I. & Zoller, P. Quantum repeaters for communication. *Phys. Rev. Lett.* **81**, 5932–5935 (1998).
58. Dür, W., Briegel, H. J., Cirac, J. I. & Zoller, P. Quantum repeaters based on entanglement purification. *Phys. Rev. A* **59**, 169–181 (1999).
59. Bennett, C. H. & DiVincenzo, D. P. Quantum computing—towards an engineering era? *Nature* **377**, 389 (1995).
60. van Enk, S. J., Cirac, J. I. & Zoller, P. Ideal quantum communication over noisy channels: a quantum optical implementation. *Phys. Rev. Lett.* **78**, 4293–4296 (1997).
61. van Enk, S. J., Kimble, H. J., Cirac, J. I. & Zoller, P. Quantum communication with dark photons. *Phys. Rev. A* **59**, 2659–2664 (1999).
62. Mabuchi, H., Turchette, Q. A., Chapman, M. S. & Kimble, H. J. Real-time detection of individual atoms falling through a high-finesse optical cavity. *Opt. Lett.* **21**, 1393–1395 (1996).
63. Haroche, S., Brune, M. & Raimond, J. M. Experiments with single atoms in a cavity: entanglement, Schrödinger's cats and decoherence. *Phil. Trans. R. Soc. Lond. A* **355**, 2367–2380 (1997).
64. Turchette, Q. A. *et al.* Deterministic entanglement of two ions. *Phys. Rev. Lett.* **81**, 3631–3634 (1998).
65. Monroe, C., Meekhof, D. M., King, B. E., Itano, W. M. & Wineland, D. J. Demonstration of a fundamental quantum logic gate. *Phys. Rev. Lett.* **75**, 4714–4717 (1995).
66. DiVincenzo, D. P. & Loss, D. Quantum information is physical. *Superlatt. Microstruct.* **23**, 419–432 (1998).
67. Cirac, J. I. & Zoller, P. Quantum computations with cold trapped ions. *Phys. Rev. Lett.* **74**, 4091–4094 (1995).
68. Lloyd, S. A potentially realizable quantum computer. *Science* **261**, 1569–1571 (1993).
69. Lloyd, S. Envisioning a quantum supercomputer. *Science* **263**, 695 (1994).
70. Cory, D. G., Fahmy, A. F. & Havel, T. F. Ensemble quantum computing by nuclear magnetic resonance spectroscopy. *Proc. Natl. Acad. Sci. USA* **94**, 1634–1639 (1997).
71. Gershenfeld, N. A. & Chuang, I. L. Bulk spin resonance quantum computation. *Science* **275**, 350–356 (1997).
72. Cory, D. G., Price, M. D. & Havel, T. F. Nuclear magnetic resonance spectroscopy: an experimentally accessible paradigm for quantum computing. *Physica D* **120**, 82–101 (1998).
73. Chuang, I. L., Vandersypen, L. M. K., Xinlan Zhou, Leung, D. W. & Lloyd, S. Experimental realization of a quantum algorithm. *Nature* **393**, 143–146 (1998).
74. Chuang, I. L., Gershenfeld, N. & Kubinec, M. Experimental implementation of fast quantum searching. *Phys. Rev. Lett.* **80**, 3408–3411 (1998).
75. Laflamme, R., Knille, E., Zurek, W. H., Catasti, P. & Mariappan, S. V. S. NMR Greenberger Horne Zeilinger states. *Phil. Trans. R. Soc. Lond. A* **356**, 1941–1948 (1998).
76. Cory, D. G. *et al.* Experimental quantum error correction. *Phys. Rev. Lett.* **81**, 2152–2155 (1998).
77. Jones, J. A. & Mosca, M. Implementation of a quantum algorithm on a nuclear magnetic resonance quantum computer. *J. Chem. Phys.* **109**, 1648–1653 (1998).
78. Jones, J. A. & Mosca, M. Approximate quantum counting on an NMR ensemble quantum computer. *Phys. Rev. Lett.* **83**, 1050–1053 (1999).
79. Linden, N., Barjat, H. & Freeman, R. An implementation of the Deutsch Jozsa algorithm on a three qubit NMR quantum computer. *Chem. Phys. Lett.* **296**, 61–67 (1998).

Acknowledgements

This work was supported by the US Army Research office.

Correspondence and requests for materials should be addressed to C.H.B. (e-mail: bennetc@watson.ibm.com).