

PHYSICS

Quantum computing

Emanuel Knill

The race is on to build a computer that exploits quantum mechanics. Such a machine could solve problems in physics, mathematics and cryptography that were once thought intractable, revolutionizing information technology and illuminating the foundations of physics. But when?

What is a quantum computer?

We can think of it as a traditional computer that has access to, and can manipulate, quantum information. Information in traditional, or classical, computers is represented by strings of '0's and '1's called bits. Quantum information uses quantum bits, or qubits, instead. A qubit is a quantum system that has two distinguishable configurations corresponding to the bit values 0 and 1. But being a quantum system, the superposition principle applies: the qubit's state can be any combination, or superposition, of the two configurations. Likewise, a string of qubits can be in any superposition of its bit-string configurations. This allows quantum interference to be exploited and greatly enriches the kind of information that can be represented.

What are qubits made of?

One of the first proposals for making qubits was based on photons (quanta of light). In this case, one can associate the two distinguishable configurations with whether or not a photon passes through a polarizer, such as a lens from a pair of polarizing sunglasses. Many other quantum systems can be used. Examples include trapped atoms or ions (Fig. 1) and the collective behaviour of electrons in superconducting circuits. For atoms and ions, the two configurations arise from different electron and nuclear arrangements; for superconducting circuits, the configurations can be characterized by the system's charge (for example, charged/uncharged) or flux (for example, as induced by clockwise/anticlockwise current).

How is quantum information manipulated?

Quantum computing requires the physical manipulation — initialization, control and measurement — of the state of many qubits (Box 1, overleaf). Control is accomplished




Figure 1 | Artist's impression of an ion-trap quantum computing laboratory. In general, today's physical quantum bits (qubits) are microscopic in size and confined to superconducting circuits, nanoscale semiconductor islands, impurities in solids, nuclei of molecules, or electromagnetically trapped atoms or ions. But the devices needed for confinement substantially increase the bulk of the 'quantum core'. Depicted here is an experimental set-up for a two-qubit system that consists of two trapped ions (blue spheres). The trapping requires components (shown in the series of blow-ups) such as an array of electrodes, a copper case and vacuum housing (glass structure). Beyond the core, the steps that are necessary to manipulate the quantum information encoded in the qubits — initialization, control and measurement of the qubits' state — require a roomful of gear: lasers, optical devices and electronic equipment to apply pulsed electromagnetic fields to the system, and, of course, a classical computer and a scientist to run it all. (Table set-up based on a sketch by D. Leibfried.)

by applying quantum gates, the quantum analogues of classical logic gates.


Are quantum computers faster than classical computers?

Some problems can be solved much faster on quantum computers. Because they are an extension of classical computers, they are always at least as fast. But the cost of qubit manipulations and the nature of many applications such as word processing imply that

there is no general speed advantage in 'going quantum'.

When was the potential of quantum computing first recognized?

During the second half of the twentieth century, there was a growing appreciation of the practical potential of quantum systems. Researchers realized that, by manipulating quantum systems directly, they could circumvent the hurdles of simulating quantum

Box 1 | A quantum computation

a, The computation shown requires two qubits, depicted as particles (blue spheres in the copper case) of spin $\frac{1}{2}$. The particle configurations, spinning left or right, correspond to the bit values 0 and 1, respectively. A strong interaction (not shown) initializes the qubits to their left-spinning states (arrowed), which are then isolated from decoherence processes in the copper case. We can visualize the two-qubit state as a wavefunction on the space of four possible configurations — 00, 01, 10 and 11. Each configuration contributes an amplitude to the wavefunction. The initial state has all of its amplitude at the configuration 00 (red circle).

b, The computation starts with an operation (shown here as light

passing through a lens and shining on the particles) that changes the state to a uniform superposition of the configurations. The new wavefunction has equal amplitudes at each configuration. If we measure the state (which we must not yet do), the probability of observing a certain configuration is given by its amplitude squared ($0.5^2 = 0.25$ in this case). The shade and colour of the circles associated with each configuration represent the magnitude and phase (red for positive phase, blue for negative) of the amplitudes. Although a wavefunction visualization is possible for any number of qubits, it becomes extremely unwieldy owing to the exponential growth of the number of possible configurations.

c, We apply a ‘black box’, or oracle, computation whose behaviour is difficult to predict except by using it. Because the black box acts in one step on a state that involves all configurations simultaneously, it makes use of ‘quantum parallelism’ (not to be confused with classical parallelism). In this case, the behaviour of the black box is to ‘mark’ one configuration by changing the phase of its amplitude. But which one did it change? At this point, a configuration measurement would reveal nothing, as the probabilities of the measurement outcomes depend only on the magnitudes.

d, We next exploit interference to concentrate the amplitudes on the marked configuration.

In terms of wavefunctions, this involves recombining amplitudes as if by sending them through an interferometer. Here, another type of light (blue) does the trick.

e, The qubits’ state is measured by opening the isolating case and checking the orientations of the spins. Because only one amplitude is non-zero, the outcome is deterministic and reveals the location of the mark. In this case, it turns out to be at 10. For many algorithms, there are many non-zero amplitudes, and the measurement outcomes are probabilistic. This quantum computation required a single application of the black box. Analogous classical computations would need up to three applications.

E.K.

phenomena using classical computers. This led to the idea of quantum computing, which exploits quantum systems to perform more efficient calculations. But the subject still seemed primarily of academic interest.

What kick-started the field?

It was a remarkably efficient quantum algorithm designed by Peter Shor in 1994 for factoring large numbers; this followed increasingly impressive examples of quantum solutions to ‘oracle’ problems (those involving a black box whose internal workings are inaccessible). Shor’s algorithm could break cryptographic codes commonly used for Internet communication and commerce. It became clear that quantum computers had tangible advantages, and that it was necessary and worthwhile to assess the practicality of building them.

What else are quantum computers good for?

There are other, similarly well-structured mathematical problems for which quantum algorithms can yield solutions dramatically faster than can classical computers. Significantly,

exponential speed-ups are found in simulations of quantum phenomena — one of the original motivations for considering quantum computers. In particular, they can provide a virtual laboratory, realizing quantum models of one’s choice. Less dramatic but still substantial speed-ups exist for optimization and integration problems. There are also problems involving the reduction of the amount of information being communicated between different parties, for which quantum algorithms possess theoretically proven exponential advantages.

Is that all?

That is a lot! What we have is a number of fundamental quantum-algorithm tools that can be widely applied, and we are seeking more such tools. We don’t know how to recognize or classify problems that can be solved more efficiently with a quantum computer. I have no doubt that there are applications that we cannot yet imagine.

What do we require to build a quantum computer?

We need a physical medium that can support

quantum systems possessing two distinguishable configurations for realizing the qubits, and the set-up must be such that the qubits are adequately isolated. In particular, their surroundings should not be able to ‘sense’ their configuration; otherwise, features of the analog configuration amplitudes (Box 1) required to exploit interference will become unstable and ‘decohere’. Such decoherence is one of the reasons that the effects of superpositions are not normally observed in the macroscopic world. To manipulate the qubits’ state, we need to break the isolation in controlled ways that do not introduce unwanted computational errors or decoherence. This typically means that we must be able to access and modify the state of individual qubits and qubit pairs without affecting the other qubits — in much the same way as bits in classical computers are manipulated by applying circuit elements that act on one or two bits.

Don’t we need ‘entanglement’?

Entanglement is a property of certain joint states of two or more spatially separated quantum systems. It links the systems in a way that

cannot be explained by a classical combination of the states of each of the constituent systems. Because of this, entanglement has a key role as a resource for unlocking the power of quantum communication between different parties or quantum computers. Similarly, sufficiently pure entanglement can be used to establish links between distant qubits for the purpose of applying two-qubit gates. Although quantum computing does not intrinsically require spatial separation of qubits, once one has identified the qubits in a physical medium, one can consider the entanglement between them; many researchers look for its signatures as a first step. However, we know that entanglement is not sufficient for making quantum computers powerful. There are notable examples of qubits, such as photon qubits, that are easy to entangle but hard to usefully quantum compute with.

Does the analog nature of configuration amplitudes cause problems?

The short answer is no. But it took much research to understand why this is. Originally, most of us felt that the difficulties encountered in building classical analog (as opposed to digital) computers would also apply to quantum computers. In particular, it seemed that, for the output of a computation to be sufficiently close to the desired answer, the quantum gates would have to be increasingly accurate as the number of gates grows. Because the physical interactions that underlie the computation's gates depend on parameters that take a continuous range of values, it is generally believed that it is impractical to achieve high accuracy directly. As it turns out, it is possible to digitize quantum computations arbitrarily accurately, using relatively limited resources, by applying quantum-error-correction strategies developed for this purpose.

Will quantum error correction work?

It can remove the effects of physically reasonable computational errors and decoherence processes, provided that enough of the requirements for building a quantum computer can be met. In particular, the physical manipulations of qubits must be performed sufficiently accurately. But it cannot correct every conceivable error. The ultimate test of quantum error correction will be the demonstration of a quantum computer with no apparent limit (except for cost) to the number of qubits and gates — that is, a ‘scalable’ quantum computer. Because there are no known fundamental obstacles to such scalability, it has been suggested that failure to achieve it would reveal new physics.

What is the required accuracy of physical qubit manipulations?

The accuracy of initialization, control and measurement of qubit states that is required depends on a number of factors, such as the dominant types of error (of which there are

many more than in classical computing), constraints of the physical medium supporting the qubits, and the chosen error-correction strategies. All complete error-correction strategies require many additional qubits and gates. Just how many depends on the accuracy of the basic physical manipulations, and the number required increases indefinitely as the accuracy decreases towards a limit. Although this limit is case-dependent, there is nevertheless a consensus that, for practical scalability, the probability of error introduced by the application of a quantum gate must be less than 0.0001. (Probability of error is a somewhat loose way of quantifying the inaccuracy of quantum processes.) Requirements for qubit-state initialization and measurement are more benign, and I have proposed that a probability of error below 0.01 is a reasonable goal in these cases. There are indications that scalability is possible at error probabilities as high as a few per cent for all basic physical manipulations. Whether this is practical remains to be seen.

Have sufficiently accurate quantum gates been demonstrated?

No. And it is one of the main as-yet-unmet challenges.

What can be done in the absence of such accurate gates?

Interesting quantum states of a few qubits have already been prepared. Similarly, experiments that implement and characterize the behaviour of steps of small quantum computations or error-correction protocols have been performed. Such experiments are useful even if their accuracy is insufficient for use in long computations. It has been proposed, and experimental efforts are under way, to use large numbers of moderate-quality qubits to simulate and make measurements of quantum models, such as those that have been suggested to explain different types of superconductivity.

How many qubits can be used for quantum computing today?

Experiments have demonstrated quantum states and processes useful in quantum computations with up to eight trapped-ion qubits. However, existing experimental systems are best thought of as incipient ‘quantum registers’ of data, because the number of qubits available is fixed and small. Currently, there are no generally accepted criteria for when a device can be called a one-, two- or n -qubit register. In my opinion, the following are requirements for such a quantum register: one can use it to run an arbitrary computation without manually reconfiguring its hardware; the computation should be sufficiently accurate to exploit quantum features better than would a smaller register; and the register should be realized in a system for which there is an apparently practical way to scale to larger numbers of qubits and eventually breach the scal-

ability threshold. Bearing such requirements in mind, one-qubit registers have been demonstrated, and two-qubit registers are close at hand.

What are the leading technologies for quantum computing?

Many experiments have been performed involving nuclear magnetic resonance (qubits associated with nuclear spins in molecules) and optics (qubits carried by photons). In their current form, these two technologies are close to their practical scalability limit of about ten qubits. Among technologies for which a practical path to scalability is known, that involving qubits carried by trapped ions is currently the most advanced, and will probably lead to the realization of the first quantum registers that have tens of qubits. Because qubit configurations in superconducting circuits involve the collective behaviour of many electrons, I was surprised that this technology has advanced so rapidly and can be considered the current runner-up. Several other approaches are being pursued and are waiting in the wings. Large arrays of trapped atoms or ions may soon be used for specially devised quantum simulations.

Which will be the best in the long term?

I hope that eventually we will be able to use qubits whose effective gate speeds approach those of classical gates. Ideally, all computation would be based on elementary devices whose quantum features could be exploited with few additional resources over those needed to make them compute classically. The technologies required to achieve this ideal are unknown.

When will quantum computers outperform classical computers?

Many of us are reluctant to make predictions about when quantum computers will start to be used productively. I am optimistic that I will be able to perform interesting computations on quantum devices in my lifetime, but would not be disappointed if we encounter unexpected obstacles along the way.

Emanuel Knill is at the Mathematical and Computational Sciences Division, National Institute of Standards and Technology, Boulder, Colorado 80305, USA.

e-mail: emanuel.knill@nist.gov

FURTHER READING

Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2001).

Preskill, J. *Quantum Computation* (lecture notes) www.theory.caltech.edu/people/preskill/ph219/#lecture

Spec. Issue on Implementations of Quantum Computers *Fortsch. Phys.* **48**, nos 9–11 (2000).

Nakahara, M. & Ohmi, T. *Quantum Computing: From Linear Algebra to Physical Realizations* (CRC Press, 2008).

Ladd, T. D. et al. Quantum computers *Nature* doi:10.1038/nature08812 (in the press).