Proc. Natl. Acad. Sci. USA
Vol. 95, pp. 11032–11033, September 1998
From the Academy

This paper is a summary of a session presented at the Ninth Annual Frontiers of Science Symposium, held November 7–9, 1997, at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, CA.

Quantum computing

GILLES BRASSARD*†, ISAAC CHUANG‡, SETH LLOYD§, AND CHRISTOPHER MONROE¶

*Département d'informatique et de recherche opérationnelle, Université de Montréal, C.P. 6128, succursale centre-ville, Montréal, QC H3C 3J7 Canada; ‡IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120; §Massachusetts Institute of Technology, Department of Mechanical Engineering, MIT 3-160, Cambridge, MA 02139; and ¶National Institute of Standards and Technology, Time and Frequency Division, 325 Broadway, Boulder, CO 80303

Quantum computation is the extension of classical computation to the processing of quantum information, using quantum systems such as individual atoms, molecules, or photons. It has the potential to bring about a spectacular revolution in computer science. Current-day electronic computers are not fundamentally different from purely mechanical computers: the operation of either can be described completely in terms of classical physics. By contrast, computers could in principle be built to profit from genuine quantum phenomena that have no classical analogue, such as entanglement and interference, sometimes providing exponential speed-up compared with classical computers.

Quantum Information. All computers manipulate information, and the unit of quantum information is the quantum bit, or *qubit*. Classical bits can take either value 0 or 1, but qubits can be in a linear *superposition* of the two classical states. If we denote the classical bits by $|0\rangle$ and $|1\rangle$, a quantum bit can be in any state $\alpha|0\rangle + \beta|1\rangle$, where α and β are complex numbers called *amplitudes* subject to $|\alpha|^2 + |\beta|^2 = 1$. Any attempt at measuring qubits induces an irreversible disturbance. For example, the most direct measurement on $\alpha|0\rangle + \beta|1\rangle$ results in the qubit making a probabilistic decision: with probability $|\alpha|^2$, it becomes $|0\rangle$ and with complementary probability $|\beta|^2$, it becomes $|1\rangle$; in either case the measurement apparatus tells us which choice has been taken, but all previous memory of the original amplitudes α and β is lost.

Unlike classical bits, where a single string of n zeros and ones suffices to describe the state of n bits, a physical system of n qubits requires 2^n complex numbers to describe its state. For example, two qubits can be in the state $\alpha|00\rangle + \beta|01\rangle + \gamma|10\rangle + \delta|11\rangle$ for arbitrary complex numbers α, β, γ, and δ subject only to the constraint $|\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 = 1$.

Another feature of qubits is the property of *entanglement*. Consider the two-qubit state $(|00\rangle - |01\rangle - |10\rangle + |11\rangle)/2$. This state is less complicated than it actually looks, because it can be factored into the product of two one-qubit states, each of which is $(|0\rangle - |1\rangle)/\sqrt{2}$. Similarly, many n-qubit states can be written in factored form and thus require only $2n$ numbers for their description, which is much less than the 2^n numbers generally required. However, some special states such as $(|01\rangle - |10\rangle)/\sqrt{2}$ cannot be factored. When these two qubits are measured, they yield either 0 and 1 or 1 and 0, with equal probability $(1/\sqrt{2})^2 = 1/2$, but which of these two outcomes will occur is not determined until the measurement is actually performed. This has no classical analogue.

Quantum Computing. Computers that thrive on entangled quantum information could run exponentially faster than classical computers because n qubits require 2^n numbers for their description. A few simple operations on these qubits can affect all 2^n numbers through the use of quantum parallelism and quantum interference.

Quantum parallelism arises because a quantum operation acting on a superposition of inputs produces a superposition of outputs. For example, consider some function f and a quantum logic circuit U that computes it by mapping quantum register $|x, 0\rangle$ to output $|x, f(x)\rangle$. Let x and y be two distinct inputs, and prepare the superposition $(|x, 0\rangle + |y, 0\rangle)/\sqrt{2}$. Applying U produces $(|x, f(x)\rangle + |y, f(y)\rangle)/\sqrt{2}$: the value of function f is computed on both inputs x and y even though circuit U is used once only. This works for even larger superpositions: applying U to $\Sigma_x c|x, 0\rangle$, where $c = 2^{-n/2}$ is a normalization factor, gives $\Sigma_x c|x, f(x)\rangle$, an equal superposition of all input–output pairs. *An exponential amount of computation has been achieved in the time it takes to compute the function on a single input.* Unfortunately, if this exponentially rich state is measured, the entire state collapses into a single randomly chosen input–output pair $|x, f(x)\rangle$. Indeed, it would have been easier to choose x at random before computing $f(x)$ by classical means! To make good use of quantum parallelism, we also must use the notion of *quantum interference*.

Imagine that we have an unknown function f that has a one-bit input and a one-bit output, and that we are interested neither in the value of $f(0)$ nor that of $f(1)$, but rather in whether or not those two values are equal. If we have no other access to f than a classical circuit to compute it, nothing better can be done than to run the circuit twice, to obtain and then compare $f(0)$ and $f(1)$. However, if we are given a quantum circuit U that transforms input $|x, b\rangle$ into output $|x, b \oplus f(x)\rangle$, where $\oplus$ denotes addition modulo two, we can do better.

Prepare the first input in state $(|0\rangle + |1\rangle)/\sqrt{2}$ and the second input in state $(|0\rangle - |1\rangle)/\sqrt{2}$. Put together, the input state is $(|00\rangle - |01\rangle + |10\rangle - |11\rangle)/2$. Applying U, the output is $(|0, f(0)\rangle - |0, \overline{f(0)}\rangle + |1, f(1)\rangle - |1, \overline{f(1)}\rangle)/2$, where $\bar{x}$ denotes the complement of bit x. Now apply to the first qubit the quantum transformation that maps $|0\rangle$ to $(|0\rangle + |1\rangle)/\sqrt{2}$ and $|1\rangle$ to $(|0\rangle - |1\rangle)/\sqrt{2}$. The reader is encouraged to verify that the first qubit ends up in the state $|f(0) \oplus f(1)\rangle$: measuring this qubit now produces the desired answer *even though the quantum circuit* U *to compute function f has been invoked once only.*

Intuitively, what happens is that interference occurs between computation paths. For example, there are two logical paths for the case $f(0) = f(1) = 0$ from the initial input state to output $|10\rangle$: one via $f(0) = 0$, whose amplitude is $1/2\sqrt{2}$, and one via $f(1) = 0$, whose amplitude is $-1/2\sqrt{2}$. Both paths appear to have nonzero probability of being observed, but actually they interfere *destructively* in a way that output $|10\rangle$ is in fact never observed. Interference also can be *constructive*, as in the case for output $|00\rangle$, which is twice more likely to be observed than the sum of the probabilities associated with the two paths leading to it.

The art of quantum programming is the clever exploitation of quantum interference to solve interesting problems by reinforcing the probability of obtaining desired results while at

†To whom reprint requests should be addressed. email: brassard@iro.umontreal.ca.

the same time reducing or even annihilating the probability of obtaining unwanted results. It took nearly ten years after David Deutsch first suggested that such techniques could be useful (following earlier work by Paul Benioff and Richard Feynman) (1–3) that Peter Shor discovered in 1994 a quantum algorithm for efficiently factoring large numbers (4). This attracted considerable interest, not only for its tremendous cryptographic significance, but also as a first indication that quantum computers could be genuinely faster than classical computers for solving natural problems. Subsequently, another application was discovered: imagine searching for the name of a stranger whose phone number you know, given an ordinary phone directory ordered alphabetically by names. Classically, this would require one to sift through one-half of the directory on average. Lov Grover's quantum search algorithm (5) solves this problem in about square-root the amount of time required classically. This algorithm has been extended and applied to database searches and cryptography (6).

All this theory of quantum computation is very nice, but is it only that: a theory? For the moment, no large-scale quantum computation has been achieved in the laboratory, and there is no conclusive evidence than one will ever be. Nevertheless, several teams around the globe are working at small-scale prototypes. Next, we give a glimpse into that experimental world.

Implementation. Implementation of quantum computers presents a profound experimental challenge. Quantum computer hardware must satisfy fundamental constraints: (*i*) the qubits must interact very weakly with the environment to preserve their superpositions, (*ii*) the qubits must interact very *strongly* with one another to make logic gates and transfer information, and (*iii*) the states of the qubits must be able to be initialized and read-out with high efficiency. Although few physical systems can satisfy these seemingly conflicting requirements, a notable exception is a collection of charged atoms (ions) held in an electromagnetic trap (7, 8). Here, each atom stores a qubit of information in a pair of internal electronic levels. Each atom's levels are well protected from environmental influences—indeed, this is exactly why such energy levels also are used for atomic clocks. Scaling to larger numbers of qubits is simply a matter of adding more atoms to the collection. An image of ≈35 trapped mercury atomic ions is shown in Fig. 1. When appropriate laser radiation is applied to the atoms, only one of the two internal states fluoresces (in Fig. 1, all atoms were prepared in this state). This allows near-perfect detection of the state of each qubit. The atoms are coupled by virtue of their mutual Coulomb repulsion. A particular atom's internal state can be mapped onto the collective motion of the atoms, which can subsequently be transferred to another atom's internal levels. In this way, the quantum motion of the atom collection acts as a "data bus," which allows any quantum computation to proceed. Experimental activity in trapped ion quantum computation is still in its infancy because only single-ion and two-ion quantum logic has been demonstrated to date. Extensions to larger numbers of trapped ions is hampered by a variety of technical difficulties, but there appear to be no fundamental limits to the scaling.

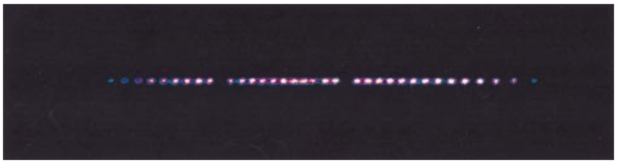

FIG. 1. Image of fluorescence from several trapped mercury (^{199}Hg$^+$) atomic ions. The ions are spaced by ≈15 μm, and the two apparent gaps are different isotopes of mercury, which do not respond to the probe laser. (Courtesy of J. Miller and J. Bergquist, NIST)

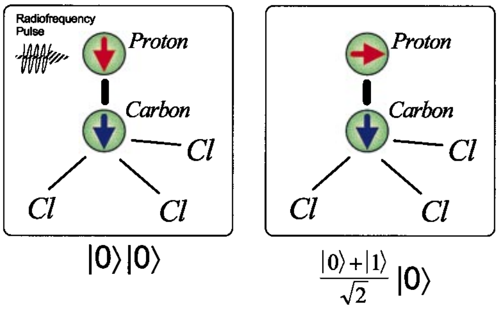

FIG. 2. Chloroform molecule as a two-bit quantum computer.

Another nearly ideal physical system that can be used as quantum computer is a single molecule, in which nuclear spins of individual atoms represent qubits (9). Using NMR techniques, invented in the 1940's and widely used in chemistry and medicine today, these spins can be manipulated, initialized, and measured. Most NMR applications treat spins as little "bar magnets," whereas in reality, the naturally well isolated nuclei are nonclassical objects. The spins' quantum behavior can be exploited to perform quantum computation; for example, the carbon and hydrogen nuclei in a chloroform molecule (Fig. 2) represent two qubits. Applying a radio-frequency pulse to the hydrogen nucleus addresses that qubit and causes it to rotate from a $|0\rangle$ to a superposition $(|0\rangle + |1\rangle)/\sqrt{2}$ state. Interactions through chemical bonds allow multiple-qubit logic to be performed. In this manner, applying newly developed techniques to allow bulk samples with many molecules to be used, small-scale quantum algorithms have been experimentally demonstrated with molecules such as Alanine, an amino acid. This includes the algorithm described above (10), to test if $f(0) = f(1)$, as well as Grover's algorithm (11) on a database with four entries. With current schemes, the difficulty of creating and maintaining quantum states grows exponentially with the number of qubits, i.e., the size of the molecules. Quantum computation is an exciting challenge, and it is expected that future experimental developments will lead to a better understanding of the practical reality of quantum computers.

We express our gratitude to Jennifer Chayes and Steven Chu for suggesting and organizing this session for the 1997 Frontiers of Science Symposium. I.C. acknowledges a collaboration with Neil Gershenfeld and Alex Pines. C.M. acknowledges a collaboration with David Wineland and his group at the National Institute of Standards and Technology. G.B. was supported in part by Canada's Natural Sciences and Engineering Research Council, Québec's Fonds pour la Formation de Chercheurs et l'Aide à la Recherche and the Canada Council. I.C. was supported in part by the Defense Advanced Research Projects Agency. S.L. was supported by the Office of Naval Research, Defense Advanced Research Projects Agency, and Army Research Office under the Quantum Information and Computation Initiative. C.M. was supported by the National Security Agency, Office of Naval Research, and Army Research Office.

1. Deutsch, D. (1985) *Proc. R. Soc. London Ser. A* **400,** 97–117.
2. Lloyd, S. (1995) *Sci. Am.* **273,** (10) 44–50.
3. Barenco, A. (1996) *Contemp. Physics* **38,** 357–389.
4. Shor, P. W. (1997) *SIAM J. Comput.* **26,** 1484–1509.
5. Grover, L. K. (1997) *Phys. Rev. Lett.* **79,** 325–328.
6. Brassard, G. (1997) *Science* **275,** 627–628.
7. Cirac, J. I. & Zoller, P. (1995) *Phys. Rev. Lett.* **74,** 4091–4094.
8. Monroe, C., Meekhof, D. M., King, B. E., Itano, W. M. & Wineland, D. J. (1995) *Phys. Rev. Lett.* **75,** 4714–4717.
9. Gershenfeld, N. & Chuang, I. L. (1997) *Science* **275,** 350–356.
10. Chuang, I. L., Vandersypen, L. M. K., Zhou, X., Leung, D. W. & Lloyd, S. (1998) *Nature (London)* **393,** 143–145.
11. Chuang, I. L., Gershenfeld, N. & Kubinec, M. (1998) *Phys. Rev. Lett.* **80,** 3408–3411.