the membrane. The membrane position at each moment was a measure of the pressure drop ΔP. At the same time, the vibrations of the membrane reflected the mass current oscillations. At first, Packard, Davis and company could not distinguish the oscillations from the background noise just by looking at the oscilloscope trace from their detector. But when they connected the output to audio headphones, their ears were able to sort out the signal. Davis says they were ecstatic when they first heard the tone. They hadn't expected that the sound would be so clear.

With the confidence that the desired signal was there, the Berkeley collaboration was then able to extract the graph of the frequency of the oscillations as a function of pressure, as shown in the figure on page 18. All the data from five temperatures fall on a nice straight line, whose slope is close to the expected value of m/ph.

Each of these frequencies was determined by averaging over very short time intervals because the pressure did not remain constant at one value for long; the researchers applied a pressure pulse and listened to the frequencies drop down the scale as the pressure decayed.

Although the published data do not determine the Josephson current–phase relationship embodied in equation 1, Packard, speaking at the Paris symposium, discussed more recent work in which the Berkeley team had made a direct measurement of this relationship.

An intriguing—and rewarding—aspect of their results is the demonstration that the separate flows through the thousands of apertures in the membrane apparently acted coherently: if they hadn't, the various oscillations would have cancelled one another out. The Berkeley researchers had gambled on their expectation that the array would act as a single coherent weak link, and that gamble paid off. It enabled them to effectively magnify the extremely faint signal one would hear through a single opening.

Barbara Goss Levi

References

1. S. V. Pereverzev, A. Loshak, S. Backhaus, J. C. Davis, R. E. Packard, Nature **388**, 449 (1997).
2. O. Avenel, E. Varoquaux, Jpn. J. Appl. Phys. **26**, Supplement 26-3, 1798 (1987); Phys. Rev. Lett. **60**, 416 (1988).
3. See, for example, P. W. Anderson, Rev. Mod. Phys. **38**, 298 (1966).
4. B. S. Deaver, J. M. Pierce, Phys. Lett. **38A**, 81 (1972).
5. H. J. Paik, J. Appl. Phys. **47**, 1168 (1976).

Exhaustive Searching Is Less Tiring with a Bit of Quantum Magic

Quantum computers have been shown to provide a dramatic speedup over classical computers in solving problems by exhaustive searching. For example, the widely used 56-bit Data Encryption Standard could be cracked with a mere 200 million or so computations instead of about 35 quadrillion.

The elementary particle of information used by modern digital computers is the bit—a register or memory element that can be in one of two distinct states, 0 or 1. But we live in a quantum world, and one can design computers in which each elementary unit of information is a quantum bit, or qubit, which can be in any superposition of two quantum states, $|0\rangle$ and $|1\rangle$. A quantum computer built with n such components could itself be in a superposition of 2^n distinct states, each splinter of the superposition performing its own computation in parallel with all the rest.

What computational magic could be performed on such a device? Three years ago, much interest in quantum computation was sparked when Peter Shor of AT&T Laboratories devised a quantum algorithm that could solve the factorization problem much faster than any known classical algorithm. Now, Lov K. Grover of Bell Laboratories, Lucent Technologies, has devised a fast quantum algorithm to search for an entry in an unordered database.[1] (See figure at right.)

"If quantum computers are being used a hundred years from now," said John Preskill of Caltech, "I would guess that they will be used to run Grover's algorithm or something like it." He calls Grover's algorithm "the simplest example of an interesting problem for which a quantum computer has a clear advantage (in principle) over a classical computer."

Furthermore, Preskill said, "the Grover algorithm, much more so than the Shor algorithm, can be adapted to many different computationally hard problems. In principle, the unsorted database search can be used to solve any NP problem—a problem for which the solution may be hard to find but is easy to verify. The database is all the trial solutions; we can invoke quantum parallelism to try them all at once and search for the one that works." If there is only one correct solution among N possibilities, an exhaustive search like this will typically take $N/2$ trials before the answer is found. By contrast, Grover's quantum algorithm almost certainly finds the correct answer in about $\sqrt{N}$ trials.

To get an idea of the significance of this, consider an example cited[2] by Gilles Brassard (University of Montreal): The widely used Data Encryption Standard relies on a 56-bit key. In "a classic scenario in secret intelligence," to crack the code one must try out keys from the $2^{56} = 7 \times 10^{16}$ possible keys. Classical methods will take, on average, about 3.5×10^{16} trials; Grover's algorithm will need only about 200 million. At a million trials per second, that's more than 1000 years versus less than 4 minutes.

The advantage of Grover's algorithm is known with certainty: The $N/2$ time needed on average by a classical algorithm cannot be improved by the discovery of some unexpectedly efficient algorithm. Furthermore, ear-

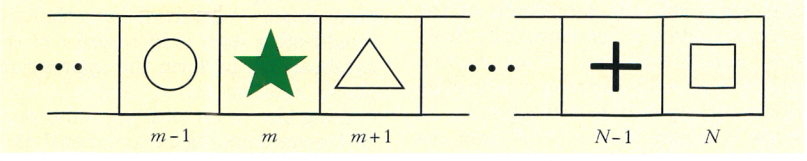

SEARCHING AN UNORDERED DATABASE OF N RECORDS for a unique item (represented by the green star in record m of the database) will take, classically, $N/2$ steps to have even a 50% probability of success. A quantum computer programmed with Grover's algorithm, however, achieves essentially 100% success in only $\pi\sqrt{N}/4$ steps, a dramatic speedup for large N. The algorithm can be used to achieve a comparable speedup in solving many other problems.

lier work by Brassard, Charles Bennett (IBM Thomas J. Watson Research Center), Ethan Bernstein (Microsoft Corp) and Umesh Vazirani (University of California, Berkeley) showed that a $\sqrt{N}$ speedup is essentially as fast as can be achieved with a quantum computer for this type of problem.[3] As Brassard told us, "We proved that quantum magic cannot be used to go faster than the square root of N. We never believed that you could do the square root of N, but to everybody's amazement, Grover found an algorithm that does exactly that." The speed limit found in the early work included an overall arbitrary constant factor, but more recent analysis by Brassard, Michel Boyer and Alain Tapp (both University of Montreal) and Peter Høyer (Odense University, Denmark) shows that Grover's algorithm is essentially within a few percent of the fastest allowed.[4]

Quantum diffusion

Grover told PHYSICS TODAY that the idea for his algorithm came from the nature of Schrödinger's equation itself. "This equation is like a diffusion equation," he said, "except if certain states have a potential, their phase gets rotated." One can think of the quantum computer as a particle, and the solution of the search problem is a particular state that we want the particle to end up in. "If we want to get a particle into a certain state," Grover said, "the way is to lower that state's potential—that is, rotate its phase, and then do the diffusion."

The computer begins in a superposition of all N possible solutions, and the combination of phase rotations and quantum diffusion works to transform this gradually into a state very close to the pure solution state, somewhat like a recurrence. The diffusion transfers amplitude among the various states in the superposition; the phase rotations single out the state that corresponds to the desired solution, ensuring a steady accumulation of amplitude in that state. The box on this page describes these iterations.

The precise number of iterations that optimizes the amplitude of the solution state was determined by Brassard and coworkers.[4] They demonstrated that each iteration of the algorithm corresponds to a rotation by an angle θ, where $\sin\theta = 1/\sqrt{N}$ and the angle $\pi/4$ corresponds to the desired solution. For large N, after $\pi\sqrt{N}/4$ iterations the probability of obtaining the wrong result is less than $1/N$. If one carries out more iterations, however, the probability of success starts declining!

The Steps of Grover's Search Algorithm

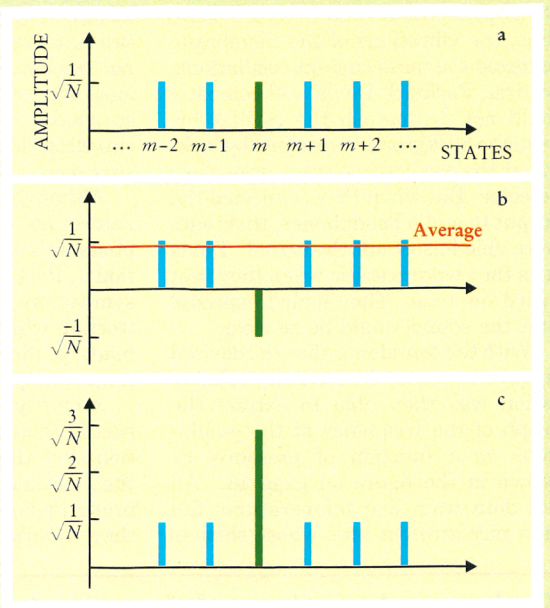

a: Initialization. At the start of the algorithm, the computer is put in a state corresponding to an equal superposition of all possible answers from 1 to N. This is "easy" to set up on a quantum computer: The n quantum bits, or qubits, representing the $N = 2^n$ possible solutions are each independently rotated to the superposition

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

The resulting product state is the one shown.

b: Phase rotation. For each state $|r\rangle$, the computer checks if the sought item is in record r of the database. If so, the phase of the state is rotated by π. That is, the correct state, $|m\rangle$, is multiplied by −1. The computer can do this for all N states in the superposition in a single step.

c: Diffusion. This step amounts to an "inversion about the average": A state with amplitude w is mapped to one with amplitude $2A - w$, where A is the average amplitude (including phase) of the N states. For the illustrated first iteration, the average amplitude is slightly less than $1/\sqrt{N}$. The inversion about average thus flips the amplitude of state m up to slightly less than $3/\sqrt{N}$, and the amplitudes of incorrect states are slightly reduced. The inversion about average can be represented by a unitary $N \times N$ matrix and can be implemented as a product of three unitary local transition matrices. In other words, the diffusion step can be physically implemented.

Iteration. Each successive iteration of phase rotation and diffusion increases the amplitude in the solution state until, after about $\pi\sqrt{N}/4$ iterations, the solution state has amplitude very close to 1. Measuring the computer's quantum state at this stage will, with near certainty, yield the correct solution, the location m of the desired database entry. At the optimum integer number of iterations, the probability of error is less than $1/N$.

Decoherence must be avoided

As with any quantum computation, loss of coherence must be avoided until the quantum algorithm has been completed. Consider what this means when the computer performs the phase rotation step on a state $|r\rangle$: The computer must (1) determine whether the sought-after item is in record r of the database, (2) perform the appropriate phase rotation on $|r\rangle$ and (3) leave nothing else in its total quantum state (or the surrounding environment!) that depends on the outcome of steps 1 and 2. If step 3 is unsuccessful, when the computer acts on a superposition of states the amplitudes of the different states will not combine coherently in the subsequent diffusion steps. Avoiding decoherence of this sort is what makes the construction of quantum computers a tremendous technological challenge.[5] (See the article, "Quantum Computing: Dream or Nightmare?" by Serge Haroche and Jean-Michel Raimond, in PHYSICS TODAY, August 1996, page 51, and the subsequent letters to the editor in the November 1996 issue, page 107.)

Coherence provides an intuitive explanation of how a quantum computer can achieve a $\sqrt{N}$ speedup. Consider a light source with N elements emitting light. If light from these N elements all combines with the same phase, the constructive interference results in an amplitude $\sqrt{N}$ times greater than if they combine with random phases. "Equivalently," Grover told us, "if we think of the algorithm as a Feynman

path integral, summing over all possible paths, the paths leading to the desired result interfere constructively and add up. Those leading to undesired states cancel out. All this is achieved by adjusting phases."

Extensions

A similar process can be carried out in a continuous, or analog, fashion. Edward Farhi (MIT) and Sam Gutmann (Northeastern University) describe this "analog analogue" of Grover's algorithm.[6] In their model, they consider an N-dimensional Hilbert space with a Hamiltonian that has only one nonzero eigenvalue, E. The task is to find the corresponding eigenvector, $|E\rangle$. They show that if one starts with an arbitrary normalized vector and adds an appropriate driving term to the Hamiltonian, after a time of order $\hbar\sqrt{N}/E$ the system will typically have a substantial amplitude in state $|E\rangle$.

Other extensions of Grover's algorithm apply it to more general problems. Boyer, Brassard, Høyer and Tapp consider the case in which the database holds c items that satisfy the desired condition.[4] If the number c is known in advance, one of the c items can be found with near certainty after $\pi\sqrt{N}/4\sqrt{c}$ iterations. (It is assumed that c is small compared to N.) For the case in which c is not known in advance, they provide a modified algorithm that achieves success, on average, after about $4\sqrt{N/c}$ iterations.

Brassard and coworkers[4] have also developed a quantum algorithm that can estimate the number of items c, using a combination of Grover's algorithm and techniques borrowed from Shor's factorization algorithm. The number of steps required is proportional to $\sqrt{N}$, with the details depending on the accuracy required and the value of c itself.

Christoph Dürr (University of Paris—South) and Høyer have devised a fast quantum algorithm for finding the smallest number in an unsorted table of numbers.[7] Brassard, Høyer and Tapp have devised a quantum algorithm that solves the "collision problem" faster than any possible classical algorithm.[8] The collision problem for a function F is to find two distinct integers m_0 and m_1 such that $F(m_0) = F(m_1)$. The speedup, for F acting on integers up to N, is from $N^{1/2}$ to $N^{1/3}$. This may seem like a small advance for an arcane problem, but the security of cryptographic protocols that use hashing functions depends on the difficulty of solving the collision problem.

Quantum telecomputation

Grover has used techniques similar to those of his search algorithm to develop fast algorithms for estimating the median and mean of populations of numbers to a specified accuracy. He has also shown how to further speed up the mean calculation by a constant factor by "quantum telecomputation," in which a number of quantum processors are run in parallel and linked by quantum entanglement (as occurs in Einstein–Podolsky–Rosen pairs of particles).[9]

Grover and—independently—Barbara M. Terhal (University of Amsterdam) and John A. Smolin (IBM Thomas J. Watson Research Center) have devised fast quantum algorithms for retrieving information from a database in a single quantum query of the database.[10] (The result was also implicit, but not emphasized, in a 1993 paper by Bernstein and Vazirani.) In Grover's search algorithm, each database query returns one bit of information relating to whether the sought item is present in a particular location. The single-query algorithms return one bit of information that answers a complicated question about many of the database's records. (For example: "Is the sought item present in the specified sequence of records?") These algorithms are of interest because classical algorithms must take at least $\log N$ database queries to perform what can be done in a single quantum query. The quantum algorithms do not provide faster searches because about N or more other quantum steps are needed before or after the database query, and large numbers of registers are needed. Nevertheless, they are of importance for showing another way in which quantum operations can outperform classical systems.

The variety of these applications suggests that, in Grover's search algorithm, quantum computer programmers may have the "killer app" that will drive demand for working devices.

GRAHAM P. COLLINS

References

1. L. K. Grover, Phys. Rev. Lett. **79**, 325 (1997).
2. G. Brassard, Science **275**, 627 (1997).
3. C. H. Bennett, E. Bernstein, G. Brassard, U. Vazirani, to appear in SIAM J. Computing, October 1997; also available on http://xxx.lanl.gov/ as quant-ph/9701001.
4. M. Boyer, G. Brassard, P. Høyer, A. Tapp, in *PhysComp 96: Proc. 4th Workshop on Physics and Computation*, T. Toffoli, M. Biafore, J. Leão, eds., New England Complex Systems Institute, Cambridge, Mass. (1996), p. 36.
5. J. Preskill, presentations given at the ITP Conference on Quantum Coherence and Decoherence, University of California, Santa Barbara, 15–18 December 1996. Available as quant-ph/9705031 and quant-ph/9705032.
6. E. Farhi, S. Gutmann, quant-ph/9612026.
7. C. Dürr, P. Høyer, quant-ph/9607014.
8. G. Brassard, P. Høyer, A. Tapp, quant-ph/9705002.
9. L. K. Grover, quant-ph/9607024; quant-ph/9704012.
10. L. K. Grover, quant-ph/9706005. B. M. Terhal, J. A. Smolin, quant-ph/9705041.

From Ethane to Benzene through a Supersonic Nozzle

Organic chemists have long been challenged to find an easy way to break apart, or activate, the single bonds in methane so that they could use this very abundant compound as a building block to synthesize organic compounds. But methane is stubbornly unreactive. Recently, Dudley Herschbach and his colleagues at Harvard University found that they could use a catalytic supersonic nozzle to produce high yields of more reactive intermediaries starting with ethane—a cousin of methane.[1] The new method is tantalizing, although it has not yet worked with methane—and even if it were to, it would have to be scaled up from producing thimblefuls to churning out millions of barrels of hydrocarbons before it would be of interest to the petroleum industry. Short of that, the catalytic supersonic nozzle is of interest as a way to do nonequilibrium chemistry.

By exploiting a nonequilibrium reaction, Harvard chemists have found they can readily produce a high yield of unsaturated hydrocarbons starting with a hydrocarbon that is normally rather unreactive.

Methane (CH_4) and ethane (C_2H_6) are members of the alkane family (C_nH_{2n+2}). All family members feature single bonds between the carbon and hydrogen atoms that are relatively hard to break. The Harvard chemists start with a chamber of ethane gas at pressures around 85 torr and temperatures of about 1000 °C. The gas flows out of the chamber through a supersonic nozzle (essentially a pinhole) made from either nickel or molybdenum, elements known to catalyze the