

remains highly speculative. Only time and more experimental data will tell.

GRAHAM P. COLLINS

References

1. M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, E. A. Cornell, *Science* **269**, 198 (1995).
2. K. B. Davis, M.-O. Mewes, M. R. An-

- draws, N. J. van Druten, D. S. Durfee, D. M. Kurn, W. Ketterle, *Phys. Rev. Lett.* **75**, 3969 (1995).
3. C. C. Bradley, C. A. Sackett, J. J. Tollett, R. G. Hulet, *Phys. Rev. Lett.* **75**, 1687 (1995).
4. H. T. C. Stoof, *Phys. Rev. A* **49**, 3824 (1994).
5. P. A. Ruprecht, M. J. Holland, K. Burnett, M. Edwards, *Phys. Rev. A* **51**,

- 4704 (1995). Yu. Kagan, G. V. Shlyapnikov, J. T. M. Walraven, "BEC in trapped atomic gases," preprint (1995).
6. G. Baym, C. Pethick, *Phys. Rev. Lett.* **76**, (1996).
7. F. Dalfovo, S. Stringari, "Bosons in anisotropic traps: Ground state and vortices," preprint cond-mat/9510142, University of Trento, Italy, October 1995.

Labs Demonstrate Logic Gates for Quantum Computation

Since the early 1980s, theoreticians of various stripes have been carrying on a lively discussion about how and why one might build a quantum mechanical computer. (See the article by Charles H. Bennett, *PHYSICS TODAY*, October 1995, page 24.) Now the appearance of two back-to-back papers in the 18 December issue of *Physical Review Letters*, reporting the demonstration of experimental "quantum logic gates," has focused the discussion onto the physics laboratory.

In one of these papers,¹ Christopher Monroe, David Wineland and coworkers at the National Institute of Standards and Technology facility in Boulder, Colorado, report the operation of a quantum logic gate that couples the hyperfine splitting of a single trapped ion to its oscillation modes in the ion trap. The device performs the function of a "controlled-NOT" Boolean logic gate on a pair of binary input bits specified by the oscillation mode and the hyperfine state.

In the adjacent paper,² Jeffrey Kimble's quantum optics group at Caltech reports the demonstration of large nonlinear phase shifts for photon pairs coupled by a single atom in a quantum electrodynamic cavity. Such a device would serve as a "quantum phase gate," exhibiting an optical phase shift that depends strongly on the binary input bits embodied by the polarization states of the two incoming photons.

We can label the two states of a binary information bit $|0\rangle$ and $|1\rangle$. A controlled-NOT gate, operating on two input bits (called the control bit and the target bit) will flip the target bit if, and only if, the control bit is $|1\rangle$. Such a two-bit gate, coupled with simple single-bit rotations, could serve as the universal gate for a quantum computer. So could a quantum phase gate, which phase-shifts the input states if, and only if, both are $|1\rangle$. Whether it's more efficient to employ such phase gates in place of controlled-NOT gates depends on the kind of computation one wants to do.

In a quantum computer, the input state can be in any coherent superposition of the basis states. That is its essential distinction from a classical

After years of just writing down Hamiltonians and algorithms, quantum computer enthusiasts have begun creating logic gates in the lab. Where will it end?

computer, and it's what would give a quantum computer unique capabilities for doing massively parallel computations—if the coherence between the superposed states can be adequately preserved. Quantum binary bits have come to be called "qubits" (pronounced like the biblical unit of length).

Lone ion vibrating in a trap

The Boulder group's logic gate starts with a single Be^+ ion sitting in a radio-frequency ion trap and made so cold that its motion can occupy only the first two quantized harmonic-oscillator modes along the trap's axis. These lowest vibrational states, separated by 11 MHz, serve as the $|0_v\rangle$ and $|1_v\rangle$ states of the gate's control qubit. The target qubit is the hyperfine substate of the ion's s-wave ground state. The lower-lying substate $|0_H\rangle$, with the valence electron's spin antiparallel to the spin of the nucleus, is separated from the substate $|1_H\rangle$, with the spins parallel, by an energy that corresponds to 1.250 GHz (called the carrier frequency).

Thus the two-qubit system has four different energy levels: $|0_v\rangle|0_H\rangle$, $|1_v\rangle|0_H\rangle$, $|0_v\rangle|1_H\rangle$ and $|1_v\rangle|1_H\rangle$. Irradiating the trapped ion precisely at the carrier frequency induces transitions between the two hyperfine states *without* changing the trap-oscillation mode. But by shifting the radiation frequency 11 MHz to the red or blue, one can simultaneously flip the hyperfine and oscillation states. While it's being irradiated at a given transition frequency, the atom cycles back and forth coherently between bit states at the so-called Rabi nutation frequency, which depends on the intensity of the perturbing radiation. If one stops irradiating at an arbitrary moment, the qubit ends up in an arbitrary coherent superposition of its two states. To get a complete, clean flip requires a radiation pulse that lasts for precisely

half a Rabi cycle (a " π pulse"), which in this experiment is on the order of a few microseconds.

Even though the gigahertz transition frequencies are in the microwave regime, the Boulder group operates the gate by means of optical fields. A pair of laser beams with a precisely tunable frequency separation induces stimulated Raman transitions when the difference frequency is tuned to the appropriate transition frequency. The strong spatial gradient of the optical field provides the necessary coupling between the ion's internal state and its external motion.

After setting the initial two-qubit state at will to any one of the four energy levels, or any desired coherent superposition of them, the Boulder group operates the controlled-NOT logic gate by applying a sequence of three Raman radiation pulses to the trapped ion:

- (1) a $\pi/2$ pulse with the difference between the two lasers tuned to the carrier frequency,
- (2) a 2π pulse at a difference frequency that would induce a transition between the $|1_v\rangle|1_H\rangle$ state and a convenient "auxiliary" state separated from the ground state by a 3-MHz Zeeman splitting, and finally
- (3) a repeat of the first $\pi/2$ pulse, but this time phase-shifted by π relative to step 1.

A $\pi/2$ carrier pulse lasts precisely $1/4$ of a Rabi cycle. It would convert a pure hyperfine state into a coherent equal superposition of $|0_H\rangle$ and $|1_H\rangle$. But these pulses (steps 1 and 3) have no effect on the ion's vibration mode. Step 2, by way of an excursion to the auxiliary state and back, simply reverses the sign of any component that happens to be in the $|1_v\rangle|1_H\rangle$ state after step 1.

Thus the three-pulse sequence performs the function of a controlled-NOT gate: If the control bit is in the $|0_v\rangle$ state, pulse 2 has no effect whatsoever and pulses 1 and 3 simply cancel each other out to leave the target hyperfine bit in its initial state. But if the control bit is $|1_v\rangle$, step 2 changes the sign of the state's $|1_v\rangle$ component. Thus, instead of canceling each other out, steps

1 and 3 now add constructively to flip the target hyperfine bit.

Quantum interference effects play a central role in this kind of logic gate. It is therefore essential that phases be coherently maintained. That requires minimizing the quantum system's interaction with its environment while at the same time providing strong, nonlinear coupling between the two qubits during the perturbing optical pulses. The duration and shaping of each pulse are also crucial if the pulse is to leave the state with the right superposition and phase.


If the input to the gate is a coherent superposition state rather than the simple Boolean inputs we've been discussing, the gate should operate separately but coherently on the individual components to produce the appropriate output superposition of states. In the general case the output will be an "entangled state"—a state whose wavefunction cannot be written simply as a product of wavefunctions for the individual qubits. It's like the decay state in an Einstein-Podolsky-Rosen (EPR) experiment, where the entangled final-state wavefunction of the two decay products departing in opposite directions cannot be written as a product of local wavefunctions. (Rosen's obituary appears on page 120.)

Applying the experimental logic gate to the four Boolean input states, the Boulder group found that it yielded the appropriate controlled-NOT response about 90% of the time. They attribute this fallibility "primarily to imperfect laser cooling, imperfect state preparation and detection, and decoherence effects." They measured the decoherence time of the system to be about a millisecond, "adequate for a single gate, but certainly not [yet] acceptable for a more extended computation." Peter Zoller and Ignazio Cirac at the University of Innsbruck have suggested that one could make an extended register of logic gates with an array of ions in a single trap, using one of their collective oscillation modes as a transfer qubit.³ "Their proposal was, in fact, the impetus for our experiment," Wineland told us.

Flying qubits

The control and target bits of the Caltech quantum phase gate are two photons of different optical frequency, passing together through a low-loss QED cavity 53 microns long. Kimble calls them "flying qubits." The binary state of each qubit, labeled $|+\rangle$ or $|-\rangle$, is taken to be the photon's circular polarization relative to the spin of a cesium atom that's traversing the cavity at the same time.

Why the cesium atom? The mutual



QUANTUM LOGIC GATE built at the National Institute of Standards and Technology.¹ Through a 3-cm aperture we see an rf ion trap in whose central 250-micron hole (barely visible) a lone beryllium ion oscillates. In operation, two laser beams, perpendicular to each other, thread the tiny trap hole to induce hyperfine transitions in the ion. (Photo courtesy of David Wineland.)

interaction of a photon pair on its own is exceedingly weak. The assistance of an atom coupled to the cavity mediates the strong nonlinear photon-photon interaction needed for the operation of the quantum gate.

The spacing of the cavity's mirrors is precisely tuned to resonate with a particular optical transition between the cesium atom's ground state and an excited p-wave state. If one irradiates the atom in the cavity with circularly polarized photons slightly detuned from the transition frequency, only photons polarized parallel to the atom's spin will (to good approximation) induce the atomic excitation. A $|-\rangle$ photon, with antiparallel circular polarization, has hardly any effect on the atom's ground state; nor is it much affected by the excited state.

The Caltech group's demonstration of the "conditional dynamics" required for a quantum logic gate involved two laser beams at slightly different detuned frequencies traversing the cavity together through its mirrors. The "pump" beam was run with either + or - circular polarization, and its intensity was varied so that the average photon population of the cavity ranged from less than 0.1 to more than 1. The much weaker "probe" beam was linearly polarized, and its intensity corresponded to about 10^{-4} photons in the cavity at any instant. A cesium atomic beam perpendicular to the laser beams crossed the cavity with an intensity such that there was, on average, one atom in the cavity at any time.

Measuring the rotation of the probe

beam's polarization as a function of the pump beam's intensity, the group found a strong coupling effect between the two beams, but only when the pump beam's circular polarization was *positive*. The linearly polarized probe beam is, of course, a superposition of opposite circular polarizations. So one can think of the rotation as resulting from the birefringence of a single atom, which ignores the probe beam's $|-\rangle$ component while phase-shifting its $|+\rangle$ component to a degree that depends on the excitation wrought by the pump beam. It's a kind of single-atom Kerr effect, with the pump-beam intensity having a nonlinear influence on the effective refractive index of the atom at the level of a single photon.

The fact that this effect is manifest at such low laser-beam intensities demonstrates that the conditional nonlinear phase shifts result from the interaction of individual photons. Extrapolating the data to a cavity population of precisely two circularly polarized photons one gets an impressive nonlinear phase shift of about 16° if both photons are in the state $|+\rangle$, and *nothing* if either photon is in the state $|-\rangle$.

"That's the kind of two-photon logic gate we'd like to have," Kimble told us. "But because we don't yet know how to make single-photon sources, we had to use laser beams. And it would also be better to have a single trapped cesium atom rather than an atomic beam."

A conditional phase shift of 16° per photon pair is enormous by the usual standards of nonlinear optics. Much of the credit for this achievement goes

to the cavity, which greatly enhanced the coupling of the laser light to the cesium atoms by concentrating their radiative interactions along the cavity axis. "The cavity gave us, in effect, one-dimensional atoms," Kimble told us.

Serge Haroche, Jean-Michel Raimond and colleagues at the Ecole Normale Supérieure, in Paris, are doing related work in the microwave regime. They're experimenting with single Rydberg atoms interacting with very weak microwave fields in superconducting cavities. In 1994 they reported the observation of atomic wavefunction phase shifts due to microwave fields with, on average, much less than one photon in the cavity.⁴ "That dispersive-cavity QED experiment was already a demonstration of conditional dynamics at the single quantum level," Haroche contends, "at a time when logic gates were not yet fashionable in quantum optics. Whether such experiments should be heralded as first steps toward a large-scale quantum computer, which I consider an impossibility, is a matter of taste."

Caveat emptor

Rolf Landauer at IBM has long been the field's leading voice of caution.⁵ In addition to the much-discussed problem of decoherence, he stresses problems that arise from the fact that quantum computing is, in a sense, a return to old-fashioned analog computing. The coherent superposition states of a binary qubit form a continuum.

Whereas in digital computing you can always restore a slightly degraded bit to a pristine 0 or 1, the degradation of analog information is dangerously cumulative. It can quickly run an extended calculation off the rails. In quantum computing, Landauer points out, small but insidious errors arise particularly from the stringent requirements on the timing, amplitudes and classical phases of the perturbing radiative pulses that are supposed to flip the qubits.

Various error-correcting schemes have recently been put forward. But even if error accumulation proves to be a crippling limit on the number of sequential steps, the investigation of quantum computing still promises numerous rewards short of a general-purpose computer. Very few sequential steps are required, for example, in quantum communication schemes for cryptography or for EPR-type tests of the quantum theory itself. The Boulder group's work began with the idea of creating entangled quantum states to circumvent the shot-noise limit in spectroscopy.

The very notion of quantum computing makes a liberating contribution

to computer science. All attempts to state or prove theorems on the intrinsic computational complexity of a class of problems must now take into account what an idealized quantum computer could do.

BERTRAM SCHWARZSCHILD


References

1. C. Monroe, D. Meekhof, B. King, W. Itano, D. Wineland, Phys. Rev. Lett. **75**,

4714 (1995).

2. Q. Turchette, C. Hood, W. Lange, H. Mabushi, H. J. Kimble, Phys. Rev. Lett. **75**, 4710 (1995).
3. J. I. Cirac, P. Zoller, Phys. Rev. Lett. **72**, 4091 (1994)
4. M. Brune, P. Nussenzveig, F. Schmidt-Kaler, F. Bernadot, A. Maali, J. Raimond, S. Haroche, Phys. Rev. Lett. **72**, 3339 (1994)
5. R. Landauer, Philos. Trans. R. Soc. London, Ser. A **353**, 367 (1995).

Quantum Interference Used to Eliminate Optical Problem



The series of images shown above depicts the onset of a problem in optical imaging (first two panels) and a solution to it (third panel). The panel on the left displays the undistorted image of a low-intensity laser beam that has passed through a circular aperture followed by a refracting medium.

In the middle panel the beam intensity is increased and the image breaks into many random spots. The culprit is the self-focusing that occurs within the refractive medium: As the laser intensity grows, the index of refraction increases quadratically, causing any regions of the beam that are more intense than others to develop refractive index gradients. Thus the local intensity maxima grow, and the beam becomes focused into narrow pencils, or filaments, within the original beam envelope. (A comparable phenomenon involving the defocusing of the beam can occur if the index of refraction decreases rather than increases with growing intensity.)

In the panel on the right, the problem has been corrected, even though the beam is just as intense as in the middle panel. The idea, reported recently by Maneesh Jain, Andrew Merriam, Athos Kasapi, Guang-Yu Yin and Stephen Harris at Stanford University, is to couple a second laser beam to the first within the refractive medium.¹ The coupling laser is then separated out and the probe beam is imaged.

The technique is simple enough, but its explanation is more complex, tied as it is to the quantum interaction among laser beams and atoms. Essentially, the combination of the two laser beams prevents the formation of an atomic dipole moment, which otherwise would contribute to the index of refraction of the material. The two laser beams trap a fixed portion of the atoms in a particular energy state and thereby freeze the dipole moment.

In the Stanford demonstration the refractive medium was a gas of three-level atoms (²⁰⁸Pb). A key requirement was that the frequency difference between the probe laser (at 283 nm) and the coupling laser (at 406 nm) must equal a Raman resonance of the atoms.

Previous work by Harris and his colleagues² has shown that quantum interference can make an optically thick medium become transparent to a certain wavelength of light (See PHYSICS TODAY, May 1992, page 17). Now, notes Jain, the Stanford group has shown that one can also use such interference effects to render a *refractively* thick medium transparent.

Harris warns that this technique is strictly for laboratory-scale applications, because the energy needed is proportional to the number of atoms along the beam path. Except in special cases, the linewidths in solids or liquids are too broad for this technique to work. But within these constraints, Harris sees many potential applications to devices such as nonlinear converters.

BARBARA GOSS LEVI

References

1. M. Jain, A. J. Merriam, A. Kasapi, G. Y. Yin, S. E. Harris, Phys. Rev. Lett. **75**, 4385 (1995).
2. K. J. Boller, A. Imamoglu, S. E. Harris, Phys. Rev. Lett. **66**, 2593 (1991). A. Kasapi, M. Jain, G. Y. Yin, S. E. Harris, Phys. Rev. Lett. **74**, 2447 (1995).