Cavity Quantum Electrodynamics

Atoms and photons in small cavities behave completely unlike those in free space. Their quirks illustrate some of the principles of quantum physics and make possible the development of new sensors

by Serge Haroche and Jean-Michel Raimond

Fleeting, spontaneous transitions are ubiquitous in the quantum world. Once they are under way, they seem as uncontrollable and as irreversible as the explosion of fireworks. Excited atoms, for example, discharge their excess energy in the form of photons that escape to infinity at the speed of light. Yet during the past decade, this inevitability has begun to yield. Atomic physicists have created devices that can slow spontaneous transitions, halt them, accelerate them or even reverse them entirely.

Recent advances in the fabrication of small superconducting cavities and other microscopic structures as well as novel techniques for laser manipulation of atoms make such feats possible. By placing an atom in a small box with reflecting walls that constrain the wavelength of any photons it emits or absorbs—and thus the changes in state that it may undergo—investigators can cause single atoms to emit photons ahead of schedule, stay in an excited state indefinitely or block the passage of a laser beam. With further refinement of this technology, cavity quantum electrodynamic (QED) phenomena may find use in the generation and precise measurement of electromagnetic fields consisting of only a handful of photons. Cavity QED processes engender an intimate correlation between the states of the atom and those of the field, and so their study provides new insights into quantum aspects of the interaction between light and matter.

SERGE HAROCHE and JEAN-MICHEL RAIMOND work in a team of about a dozen researchers and students in the physics department of the École Normale Supérieure (ENS) in Paris. They have been studying the behavior of atoms in cavities for about 10 years. Haroche received his doctorate from ENS in 1971; he has been a professor of physics at Paris VI University since 1975. He has also been teaching and doing research at Yale University since 1984. In 1991 he became a member of the newly created Institut Universitaire de France. Raimond is also an alumnus of ENS; he earned his doctorate in 1984 working in Haroche's research group and is also a professor of physics at Paris VI University.

To understand the interaction between an excited atom and a cavity, one must keep in mind two kinds of physics: the classical and the quantum. The emission of light by an atom bridges both worlds. Light waves are moving oscillations of electric and magnetic fields. In this respect, they represent a classical event. But light can also be described in terms of photons, discretely emitted quanta of energy. Sometimes the classical model is best, and sometimes the quantum one offers more understanding.

When an electron in an atom jumps from a high energy level to a lower one, the atom emits a photon that carries away the difference in energy between the two levels. This photon typically has a wavelength of a micron or less, corresponding to a frequency of a few hundred terahertz and an energy of about one electron volt. Any given excited state has a natural lifetime—similar to the half-life of a radioactive element—that determines the odds that the excited atom will emit a photon during a given time interval. The probability that an atom will remain excited decreases along an exponential curve: to one half after one tick of the internal clock, one quarter after two ticks, one eighth after three and so on.

In classical terms, the outermost electron in an excited atom is the equivalent of a small antenna, oscillating at frequencies corresponding to the energy of transitions to less excited states, and the photon is simply the antenna's radiated field. When an atom absorbs light and jumps to a higher energy level, it acts as a receiving antenna instead.

If the antenna is inside a reflecting cavity, however, its behavior changes—as anyone knows who has tried to listen to a radio broadcast while driving through a tunnel. As the car and its receiving antenna pass underground, they enter a region where the long wavelengths of the radio waves are cut off. The incident waves interfere destructively with those that bounce off the steel-reinforced concrete walls of the tunnel. In fact, the radio waves cannot propagate unless the tunnel walls are separated by more than half a wavelength. This is the minimal width that permits a standing wave with at least one crest, or field maximum, to build up—just as the vibration of a violin string reaches a maximum at the middle of the string and vanishes at the ends. What is true for reception also holds for emission: a confined antenna cannot broadcast at long wavelengths.

An excited atom in a small cavity is precisely such an antenna, albeit a microscopic one. If the cavity is small enough, the atom will be unable to radiate because the wavelength of the oscillating field it would "like" to produce

CAVITY QED apparatus in the authors' laboratory contains an excitation zone for preparing a beam of atoms in highly excited states (*left*) and a housing surrounding a superconducting niobium cavity (*center*). Ionization detectors (*right*) sense the state of atoms after they have passed through the cavity. The red laser beam traces the line of the infrared laser used to excite the atoms; the blue beam marks the path of the atoms themselves. When in use, the entire apparatus is enclosed in a liquid-helium cryostat that cools it to less than one kelvin.

cannot fit within the boundaries. As long as the atom cannot emit a photon, it must remain in the same energy level; the excited state acquires an infinite lifetime.

In 1985 research groups at the University of Washington and at the Massachusetts Institute of Technology demonstrated suppressed emission. The group in Seattle inhibited the radiation of a single electron inside an electromagnetic trap, whereas the M.I.T. group studied excited atoms confined between two metallic plates about a quarter of a millimeter apart. The atoms remained in the same state without radiating as long as they were between the plates.

Millimeter-scale structures are much too wide to alter the behavior of conventionally excited atoms emitting micron or submicron radiation; consequently, the M.I.T. experimenters had to work with atoms in special states known as Rydberg states. An atom in a Rydberg state has almost enough energy to lose an electron completely. Because this outermost electron is bound only weakly, it can assume any of a great number of closely spaced energy levels, and the photons it emits while jumping from one to another have wavelengths ranging from a fraction of a millimeter to a few centimeters. Rydberg atoms are prepared by irradiating ground-state atoms with laser light of appropriate wavelengths and are widely used in cavity QED experiments.

The suppression of spontaneous emission at an optical frequency requires much smaller cavities. In 1986 one of us (Haroche), along with other physicists at Yale University, made a micron-wide structure by stacking two optically flat mirrors separated by extremely thin metallic spacers. The workers sent atoms through this passage, thereby preventing them from radiating for as long as 13 times the normal excited-state lifetime. Researchers at the University of Rome used similar micron-wide gaps to inhibit emission by excited dye molecules.

The experiments performed on atoms between two flat mirrors have an interesting twist. Such a structure, with no sidewalls, constrains the wavelength only of photons whose polarization is parallel to the mirrors. As a result, emission is inhibited only if the atomic dipole antenna oscillates along the

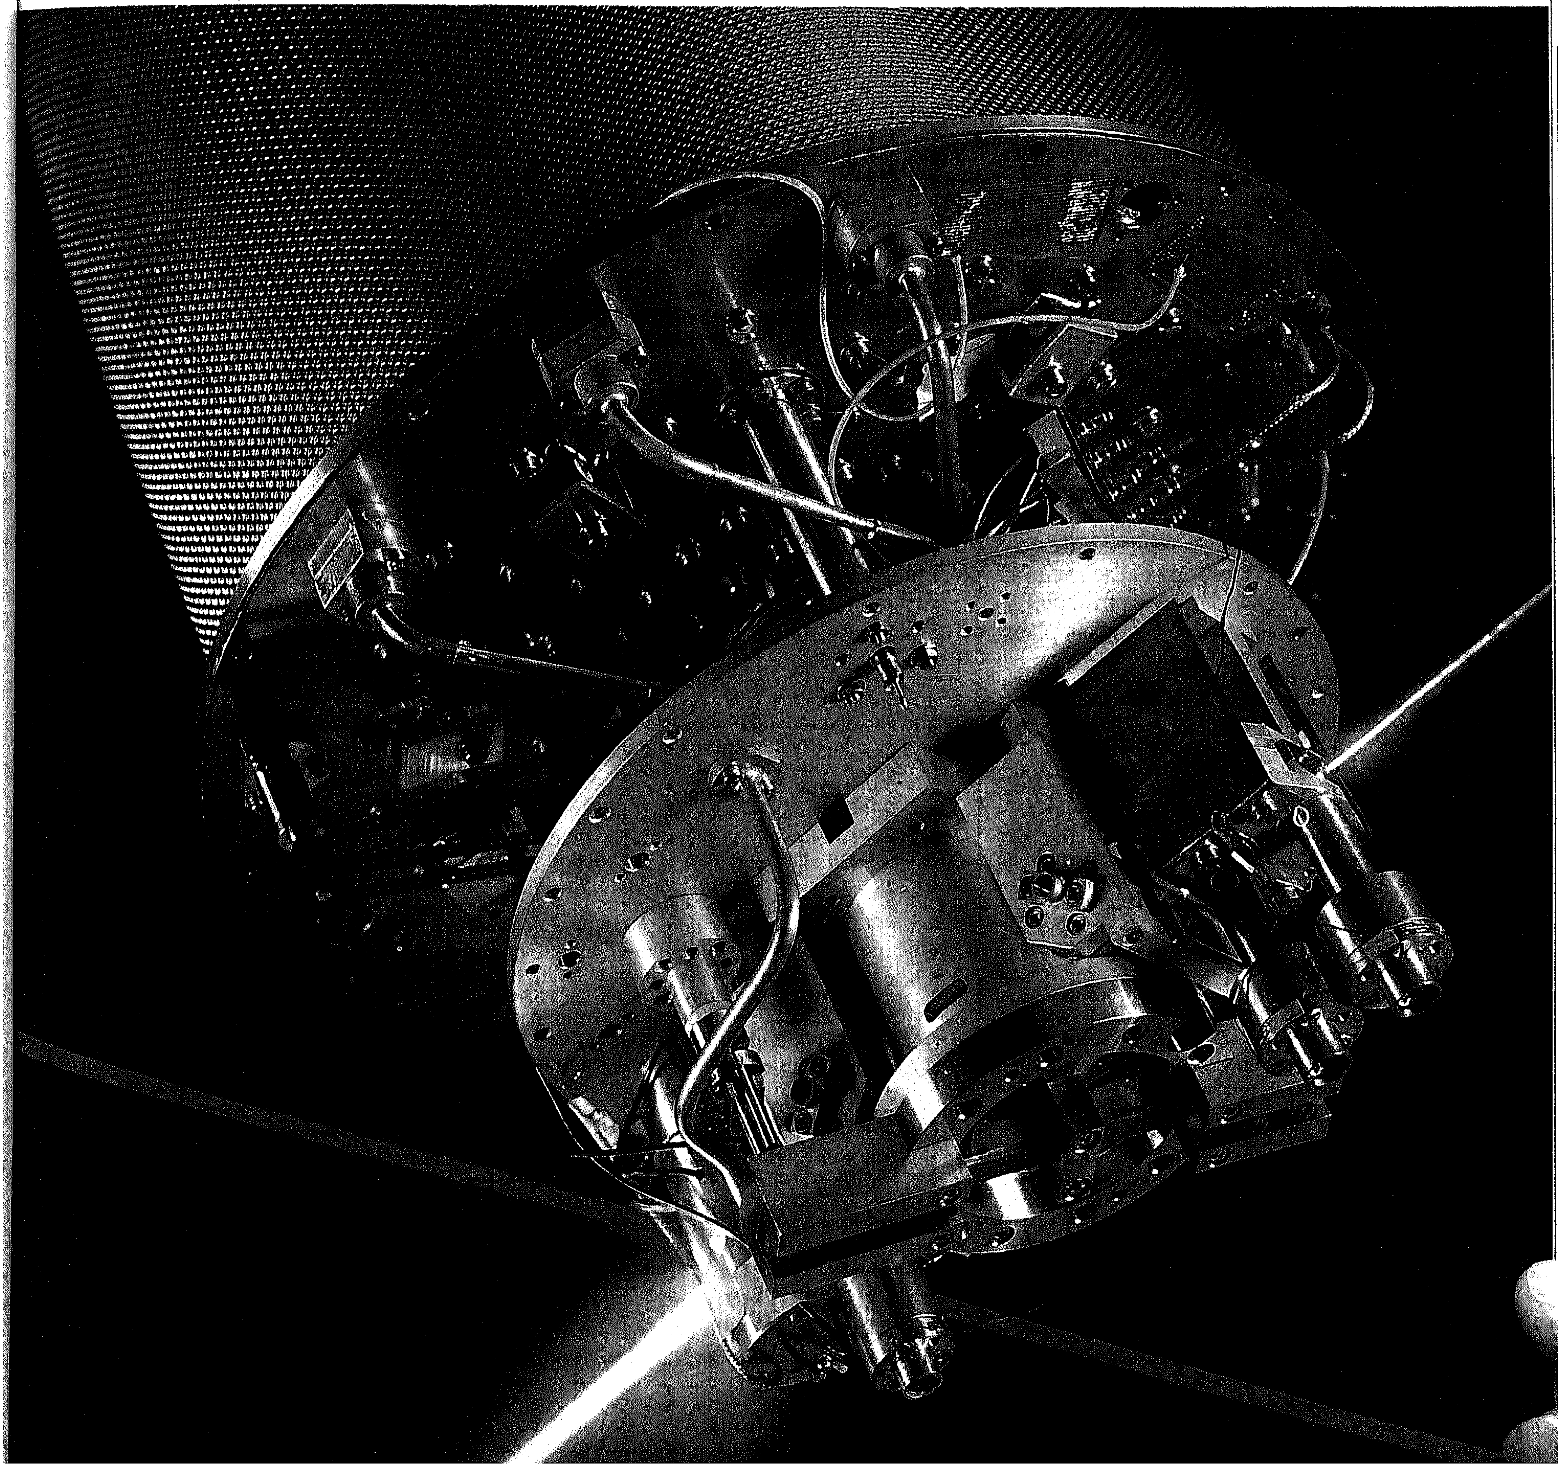

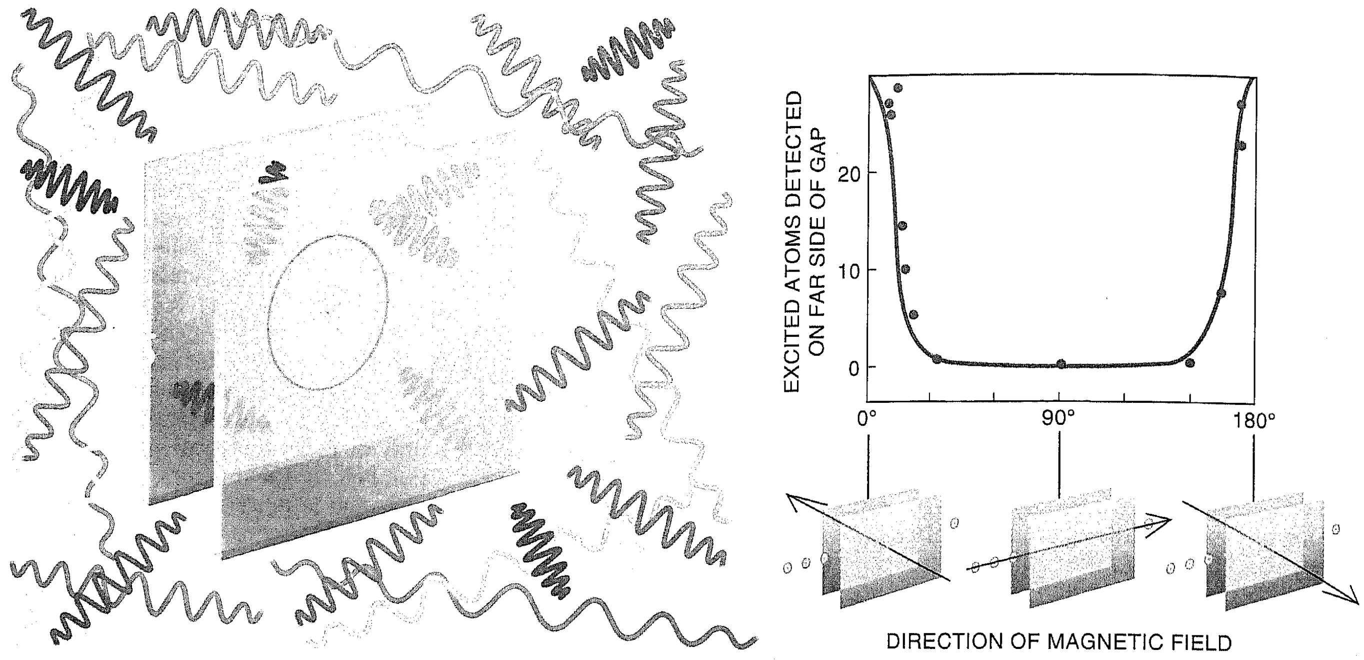

EXCITED ATOM between two mirrors (*left*) cannot emit a photon. The atom is sensitive to long-wavelength vacuum fluctuations whose polarization is parallel to the mirrors, but the narrow cavity prevents such fluctuations. Atoms passing through a micron-wide gap between mirrors have remained in the excited state for 13 natural lifetimes. Subjecting the atoms to a magnetic field causes their dipole axes to precess and changes the transmission of excited atoms through the gap (*right*). When the field is parallel to the mirrors, the atom rotates out of the plane of the mirrors and can quickly lose its excitation.

plane of the mirrors. (It was essential, for example, to prepare the excited atoms with this dipole orientation in the M.I.T. and Yale spontaneous-emission inhibition experiments.) The Yale researchers demonstrated these polarization-dependent effects by rotating the atomic dipole between the mirrors with the help of a magnetic field. When the dipole orientation was tilted with respect to the mirrors' plane, the excited-state lifetime dropped substantially.

Suppressed emission also takes place in solid-state cavities—tiny regions of semiconductor bounded by layers of disparate substances. Solid-state physicists routinely produce structures of submicron dimensions by means of molecular-beam epitaxy, in which materials are built up one atomic layer at a time. Devices built to take advantage of cavity QED phenomena could engender a new generation of light emitters [see "Microlasers," by Jack L. Jewell, James P. Harbison and Axel Scherer; SCIENTIFIC AMERICAN, November 1991].

These experiments indicate a counterintuitive phenomenon that might be called "no-photon interference." In short, the cavity prevents an atom from emitting a photon because that photon would have interfered destructively with itself had it ever existed. But this begs a philosophical question: How can the photon "know," even before being emitted, whether the cavity is the right or wrong size?

Part of the answer lies in yet another odd result of quantum mechanics. A cavity with no photon is in its lowest-energy state, the so-called ground state, but it is not really empty. The Heisenberg uncertainty principle sets a lower limit on the product of the electric and magnetic fields inside the cavity (or anywhere else for that matter) and thus prevents them from simultaneously vanishing. This so-called vacuum field exhibits intrinsic fluctuations at all frequencies, from long radio waves down to visible, ultraviolet and gamma radiation, and is a crucial concept in theoretical physics. Indeed, spontaneous emission of a photon by an excited atom is in a sense induced by vacuum fluctuations.

The no-photon interference effect arises because the fluctuations of the vacuum field, like the oscillations of more actual electromagnetic waves, are constrained by the cavity walls. In a small box, boundary conditions forbid long wavelengths—there can be no vacuum fluctuations at low frequencies. An excited atom that would ordinarily emit a low-frequency photon cannot do so, because there are no vacuum fluctuations to stimulate its emission by oscillating in phase with it.

Small cavities suppress atomic transitions; slightly larger ones, however, can enhance them. When the size of a cavity surrounding an excited atom is increased to the point where it matches the wavelength of the photon that the atom would naturally emit, vacuum-field fluctuations at that wavelength flood the cavity and become stronger than they would be in free space. This state of affairs encourages emission; the lifetime of the excited state becomes much shorter than it would naturally be. We observed this emission enhancement with Rydberg atoms at the École Normale Supérieure (ENS) in Paris in one of the first cavity QED experiments, in 1983.

If the resonant cavity has absorbing walls or allows photons to escape, the emission is not essentially different from spontaneous radiation in free space—it just proceeds much faster. If the cavity walls are very good reflectors and the cavity is closed, however, novel effects occur. These effects, which depend on intimate long-term interactions between the excited atom and the cavity, are the basis for a series of new devices that can make sensitive measurements of quantum phenomena.

Instead of simply emitting a photon and going on its way, an excited atom in such a resonant cavity oscillates back and forth between its excited and unexcited states. The emitted photon remains in the box in the vicinity of the atom and is promptly reabsorbed. The atom-cavity system oscillates between two states, one consisting of an excited atom and no photon, and the other of a de-excited atom and a photon trapped in the cavity. The frequency of this oscillation depends on the transition energy, on the size of the atomic dipole and on the size of the cavity.

This atom-photon exchange has a deep analogue in classical physics. If two identical pendulums are coupled by a weak spring and one of them is set in motion, the other will soon start swinging while the first gradually comes to rest. At this point, the first pendulum starts swinging again, commencing an ideally endless exchange of energy. A state in which one pendulum is excited and the other is at rest is clearly not stationary, because energy moves continuously from one pendulum to the other. The system does have two steady states, however: one in which the pendulums swing in phase with each other, and the other in which they swing alternatively toward and away from each other. The system's oscillation in each of these "eigenmodes" differs because of the additional force imposed by the coupling—the pendulums oscillate slightly slower in phase and slightly faster out of phase. Furthermore, the magnitude of the frequency difference between the two eigenmodes is precisely equal to the rate at which the two pendulums exchange their energy in the nonstationary states.

Researchers at the California Institute of Technology recently observed this "mode splitting" in an atom-cavity system. They transmitted a weak laser beam through a cavity made of two spherical mirrors while a beam of cesium atoms also crossed the cavity. The atomic beam was so tenuous that there was at most one atom at a time in the cavity. Although the cavity was not closed, the rate at which it exchanged photons with each atom exceeded the rate at which the atoms emitted photons that escaped the cavity; consequently, the physics was fundamentally the same as that in a closed resonator.

The spacing between the mirrors was an integral multiple of the wavelength of the transition between the first excited state of cesium and its ground state. Experimenters varied the wavelength (and hence frequency) of the laser and recorded its transmission across the cavity. When the cavity was empty, the transmission reached a sharp maximum at the resonant frequency of the cavity. When the resonator contained one atom on average, however, a symmetrical double peak appeared; its valley matched the position of the previous single peak. The frequency splitting, about six megahertz, marked the rate of energy exchange between the atom and a single photon in the cavity.

This apparatus is extremely sensitive: when the laser is tuned to the cavity's resonant frequency, the passage of a single atom lowers transmission significantly. This phenomenon can be used to count atoms in the same way one currently counts cars or people intercepting an infrared light in front of a photodetector.

Although simple in principle, such an experiment is technically demanding. The cavity must be as small as possible because the frequency splitting is proportional to the vacuum-field amplitude, which is inversely proportional to the square root of the box's volume. At the same time, the mirrors must be very good reflectors so that the photon remains trapped for at least as long as it takes the atom and cavity to exchange a photon. The group at Caltech used mirrors that were coated to achieve 99.996 percent reflectivity, separated by about a millimeter. In such a trap, a photon could bounce back and forth about 100,000 times over the course of a quarter of a microsecond before being transmitted through the mirrors.

Experimenters have been able to achieve even longer storage times—as great as several hundred milliseconds—by means of superconducting niobium cavities cooled to temperatures of about one kelvin or less. These cavities are ideal for trapping the photons emitted by Rydberg atoms, which typically range in wavelength from a few millimeters to a few centimeters (corresponding to frequencies between 10 and 100 gigahertz). In a recent experiment in our laboratory at ENS, we excited rubidium atoms with lasers and sent them across a superconducting cylindrical cavity tuned to a transition connecting the excited state to another Rydberg level 68 gigahertz higher in energy. We observed a mode splitting of about 100 kilohertz when the cavity contained two or three atoms at the same time.

There is a striking similarity between the single atom-cavity system and a laser or a maser. Either device, which emits photons in the optical and microwave domain, respec-

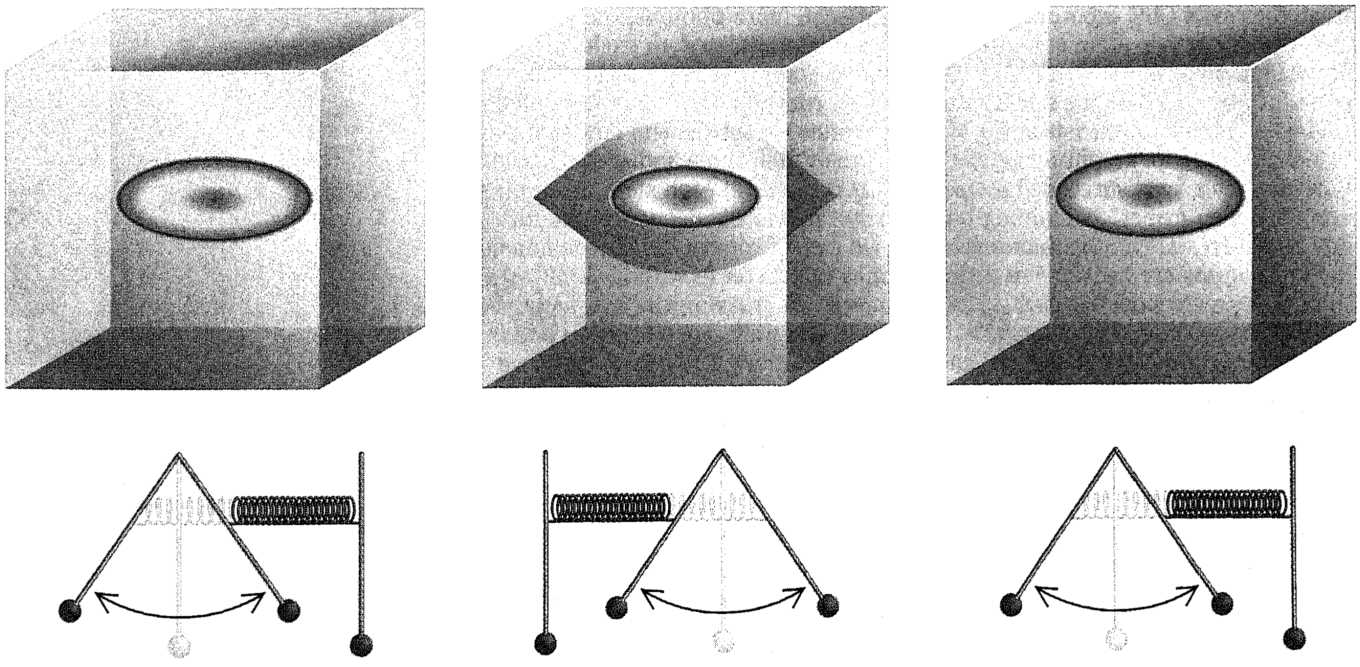

ATOM IN A CAVITY with highly reflective walls can be modeled by two weakly coupled pendulums. The system oscillates between two states. In one, the atom is excited, but there is no photon in the cavity (*left* and *right*). In the other, the atom is de-excited, and the cavity contains a photon (*center*). The atom and the cavity continually exchange energy.

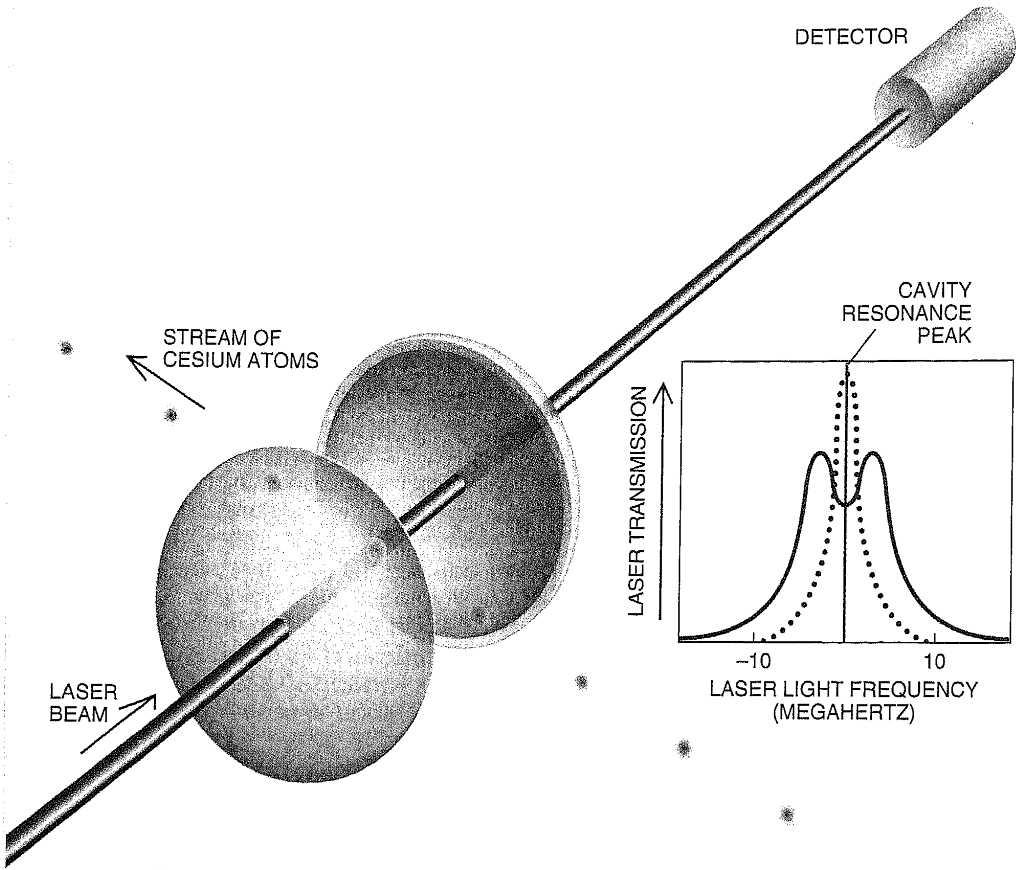

LASER BEAM TRANSMISSION through a cavity made of two closely spaced spherical mirrors is altered by the passage of individual atoms. When the cavity is empty, transmission peaks at a frequency set by the cavity dimensions (*dotted curve*). When an atom resonant with the cavity enters, however, the atom and cavity form a coupled-oscillator system. Transmission peaks at two separate frequencies corresponding to the "eigenmodes" of the atom-cavity system. The distance between the peaks marks the frequency at which the atom and cavity exchange energy.

tively, consists of a tuned cavity and an atomic medium that can undergo transitions whose wavelength matches the length of the cavity. When energy is supplied to the medium, the radiation field inside the cavity builds up to a point where all the excited atoms undergo stimulated emission and give out their photons in phase. A maser usually contains a very large number of atoms, collectively coupled to the radiation field in a large, resonating structure. In contrast, the cavity QED experiments operate on only a single atom at a time in a very small box. Nevertheless, the principles of operation are the same.

Indeed, in 1984 physicists at the Max Planck Institute for Quantum Optics in Garching, Germany, succeeded in operating a "micromaser" containing only one atom. To start up the micromaser, Rydberg atoms are sent one at a time through a superconducting cavity. These atoms are prepared in a state whose favored transition matches the resonant frequency of the cavity (between 20 and 70 gigahertz). In the Garching micromaser the atoms all had nearly the same velocity, so they spent the same time inside the cavity.

This apparatus is simply another realization of the atom-cavity coupled oscillator; if an atom were to remain inside the cavity indefinitely, it would exchange a photon with the cavity at some characteristic rate. Instead, depending on the atom's speed, there is some fixed chance that an atom will exit unchanged and a complementary chance that it will leave a photon behind.

If the cavity remains empty after the first atom, the next one faces an identical chance of exiting the cavity in the same state in which it entered. Eventually, however, an atom deposits a photon; then the next atom in line encounters sharply altered odds that it will emit energy. The rate at which atom and field exchange energy depends on the number of photons already present—the more photons, the faster the atom is stimulated to exchange additional energy with the field. Soon the cavity contains two photons, modifying the odds for subsequent emission even further, then three and so on at a rate that depends at each step on the number of previously deposited photons.

In fact, of course, the photon number does not increase without limit as atoms keep crossing the resonator. Because the walls are not perfect reflectors, the more photons there are, the greater becomes the chance that one of them will be absorbed. Eventually this loss catches up to the gain caused by atomic injection.

About 100,000 atoms per second can pass through a typical micromaser (each remaining perhaps 10 microseconds); meanwhile the photon lifetime within the cavity is typically about 10 milliseconds. Consequently, such a device running in steady state contains about 1,000 microwave photons. Each of them carries an energy of about 0.0001 electron volt; thus, the total radiation stored in the cavity does not exceed one tenth of one electron volt. This amount is much smaller than the electronic excitation energy stored in a single Rydberg atom, which is on the order of four electron volts.

Although it would be difficult to measure such a tiny field directly, the atoms passing through the resonator provide a very simple, elegant way to monitor the maser. The transition rate from one Rydberg state to the other depends on the photon number in the cavity, and experimenters need only measure the fraction of atoms leaving the maser in each state. The populations of the two levels can be determined by ionizing the atoms in two small detectors, each consisting of plates with an electric field across them. The first detector operates at a low field to ionize atoms in the higher-energy state; the second operates at a slightly higher field to

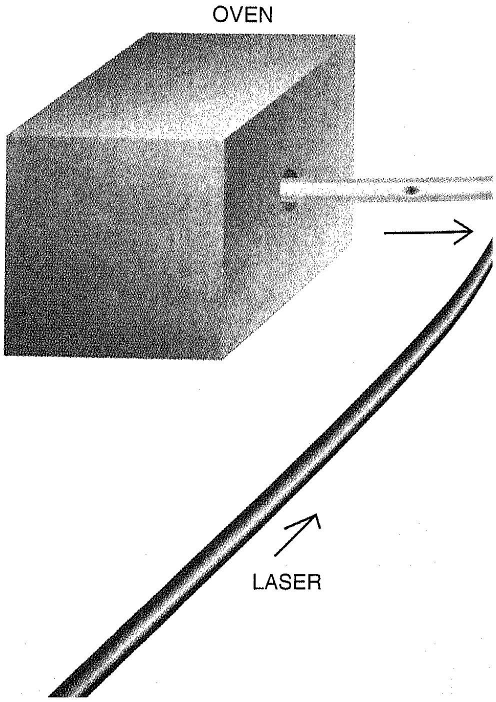

ionize atoms in the lower-lying state (those that have left a photon behind in the cavity).

With its tiny radiation output and its drastic operational requirements, the micromaser is certainly not a machine that could be taken off a shelf and switched on by pushing a knob. It is nevertheless an ideal system to illustrate and test some of the principles of quantum physics. The buildup of photons in the cavity, for example, is a probabilistic quantum phenomenon—each atom in effect rolls a die to determine whether it will emit a photon—and measurements of micromaser operation match theoretical predictions.

An intriguing variation of the micromaser is the two-photon maser source. Such a device was operated for the first time five years ago by our group at ENS. Atoms pass through a cavity tuned to half the frequency of a transition between two Rydberg levels. Under the influence of the cavity radiation, each atom is stimulated to emit a pair of identical photons, each bringing half the energy required for the atomic transition. The maser field builds up as a result of the emission of successive photon pairs.

The presence of an intermediate energy level near the midpoint between the initial and the final levels of the transition helps the two-photon process along. Loosely speaking, an atom goes from its initial level to its final one via a "virtual" transition during which it jumps down to the middle level while emitting the first photon; it then jumps down again while emitting the second photon. The intermediate step is virtual because the energy of the emitted photons, whose frequency is set by the cavity, does not match the energy differences between the intermediate level and either of its neighbors. How can such a paradoxical situation exist? The Heisenberg uncertainty principle permits the atom briefly to borrow enough energy to emit a photon whose energy exceeds the difference between the top level and the middle one, provided that this loan is paid back during the emission of the second photon.

Like all such quantum transactions, the term of the energy loan is very short. Its maximum duration is inversely proportional to the amount of borrowed energy. For a mismatch of a few billionths of an electron volt, the loan typically lasts a few nanoseconds. Because larger loans are increasingly unlikely, the probability of the two-photon process is inversely proportional to this mismatch.

The micromaser cavity makes two-photon operation possible in two ways. It inhibits single-photon transitions that are not resonant with the cavity, and it strongly enhances the emission of photon pairs. Without the cavity, Rydberg atoms in the upper level would radiate a single photon and jump down to the intermediate level. This process would deplete the upper level before two-photon emission could build up.

Although the basic principle of a two-photon micromaser is the same as that of its simple one-photon cousin, the way in which it starts up and operates differs significantly. A strong fluctuation, corresponding to the unlikely emission of several photon pairs in close succession, is required to trigger the system; as a result, the field builds up only after a period of "lethargy." Once this fluctuation has occurred, the field in the cavity is relatively strong and stimulates emission by subsequent atoms, causing the device to reach full power (about 10^{-18} watt) rapidly. A two-photon laser system recently developed by a group at Oregon State University operates along a different scheme but displays essentially the same metastable behavior.

The success of micromasers and other similar devices has prompted cavity QED researchers to conceive new experiments, some of which would have been dismissed as pure science fiction only a few years ago. Perhaps the most remarkable of these as yet hypothetical experiments are those that deal with the forces experienced by an atom in a cavity containing only a vacuum or a small field made of a few photons.

The first thought experiment starts with a single atom and an empty cavity tuned to a transition between two of the atom's states. This coupled-oscillator system has two nonstationary states: one corresponds to an excited atom in an empty cavity, the other to a de-excited atom with one photon. The system also has two stationary states, obtained by addition or subtraction of the nonstationary ones—addition of the nonstationary states corresponds to the in-phase oscillation mode of the two-pendulum model, and subtraction of the states corresponds to the

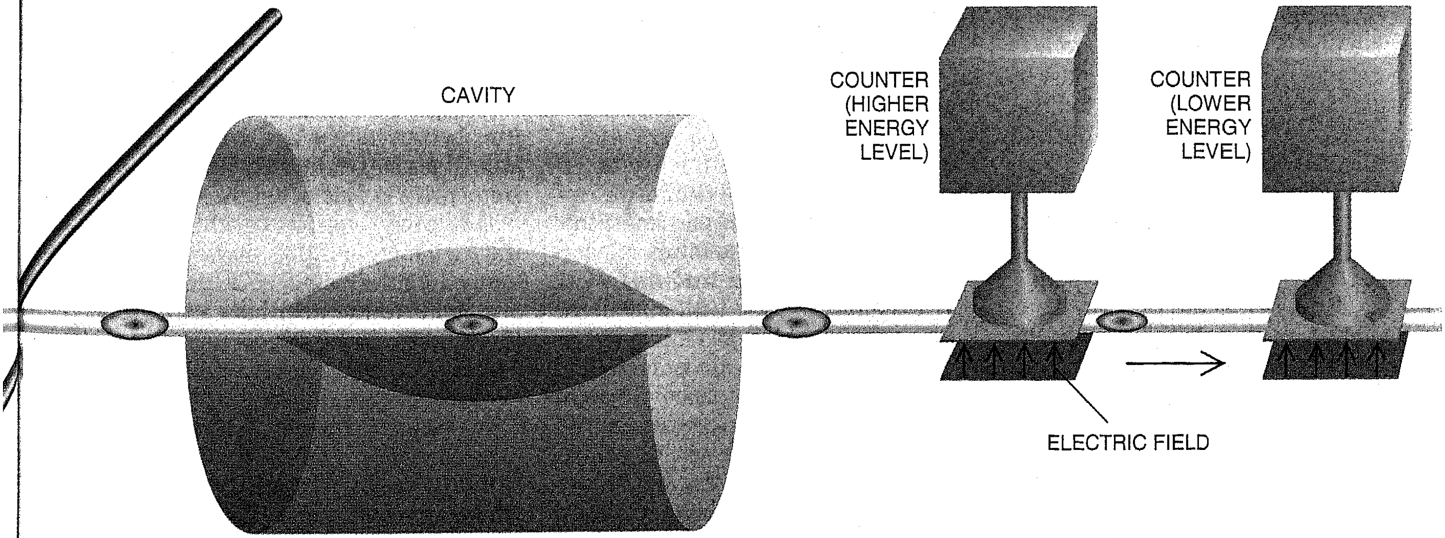

MICROMASER uses an atomic beam and a superconducting cavity to produce coherent microwave radiation. A laser beam (*left*) strikes atoms coming out of an oven and excites them into high-energy Rydberg states. The atoms pass one at a time through a cavity tuned to the frequency of a transition to a lower-energy state; the field builds up as successive atoms interact with the cavity and deposit photons in it. The micromaser field can be inferred from the readings of counters that monitor the number of atoms leaving the cavity in either the higher- or lower-energy state.

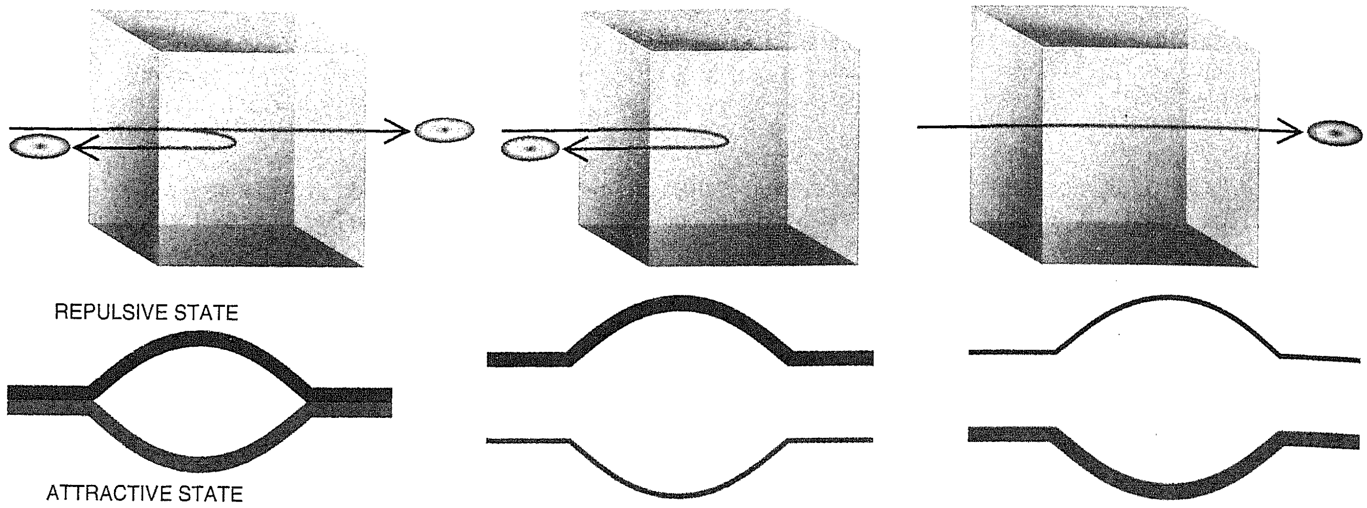

EMPTY CAVITY can repel or attract slow-moving, excited atoms. The strength of the coupling between an atom and a tuned cavity typically vanishes at the walls and reaches a maximum in the center. (Curves at the bottom show the energy of the atom-cavity system as a function of the atom's position within the cavity.) The change in energy results in a force on atoms moving through the cavity. If the cavity wavelength matches the atomic transition exactly, this force can be either attractive or repulsive (*left*). If the atomic transition has a slightly higher frequency than the resonant frequency of the cavity, the force will be repulsive (*center*); if the transition has a lower frequency, the force will be attractive (*right*).

out-of-phase mode. These stationary states differ in energy by a factor equal to Planck's constant, *h*, times the exchange frequency between the atom and the cavity.

This exchange frequency is proportional to the amplitude of the cavity's resonant vacuum field. Typically this field vanishes at the walls and near the ports by which the atom enters and leaves the cavity. It reaches a maximum at the cavity center. As a result, the atom-cavity coupling (and thus the energy difference between the system's two stationary states) is zero when the atom enters and leaves the cavity and goes to a maximum when the atom reaches the middle of the cavity.

The fundamental laws of mechanics say, however, that for a change in the relative position of two objects to lead to a change in energy, a force must be exerted between these objects. In other words, the atom experiences a push or a pull, albeit an infinitesimal one, as it moves through the empty cavity. If the system is prepared in the higher-energy state, its energy reaches a maximum at the center—the atom is repelled. If the system is in the lower-energy state, the interaction attracts the atom to the cavity center. These forces have been predicted independently by our group and by a group at Garching and the University of New Mexico.

For Rydberg atoms in a microwave cavity with a typical exchange frequency of 100 kilohertz, the potential energy difference is about one ten-billionth of an electron volt. This corresponds to a temperature of a few microkelvins and to the kinetic energy of an atom moving with a velocity of a few centimeters per second. If the speed of the incoming atom is less than this critical value, the potential barrier caused by the atom-cavity interaction will reflect the atom back, or, conversely, the potential well will be deep enough to trap it near the cavity center. Atoms in such slow motion can now be produced by laser cooling [see "Laser Trapping of Neutral Particles," by Steven Chu; SCIENTIFIC AMERICAN, February 1992]; these tiny forces may yet be observed.

If a very slow moving, excited atom is sent into a resonant, empty cavity, these forces result in a kind of atomic beam splitter. The nonstationary initial state of the system consists of the sum of the repelling and attractive states—a superposition of the two stationary atom-cavity wave functions. Half corresponds to an atom reflected back at the cavity entrance, and the other half corresponds to an atom passing through; either outcome occurs with equal probability.

To prepare a pure attractive or repelling state, one should detune the cavity slightly from the atomic transition. When the transition is a bit more energetic than the photon that the cavity can sustain, the state with an excited atom and no photon has a little more energy than the one with a de-excited atom and one photon. When the atom enters the cavity, the exchange coupling works to separate the two states, so that the state with an excited atom and no photon branches unambiguously into the higher-energy steady state, in which the atom is repelled. The same trick just as easily makes an attractive state if the cavity photon energy is slightly higher than the atomic transition.

This evolution of the atom-cavity system relies on the so-called adiabatic theorem, which says that if a quantum system's rate of change is slow enough, the system will continuously follow the state it is initially prepared in, provided the energy of that state does not coincide at any time with that of another state. This adiabaticity criterion is certainly met for the very slow atoms considered here.

These atom-cavity forces persist as long as the atom remains in its Rydberg state and the photon is not absorbed by the cavity walls. This state of affairs can typically last up to a fraction of a second, long enough for the atom to travel through the centimeter-size cavity.

The forces between atom and cavity are strange and ghostly indeed. The cavity is initially empty, and so in some way the force comes from the vacuum field, which suggests that it is obtained for nothing. Of course, that is not strictly true, because if the cavity is empty, the atom has to be initially excited, and some price is paid after all.

The force can also be attributed to the exchange of a photon between the atom and the cavity. Such a view is analogous to the way that electric forces between two charged particles are ascribed to the exchange of photons or the forces between two atoms in a molecule to the exchange of electrons.

Another interpretation of the atom-cavity vacuum attraction and repulsion, based on a microscopic analysis, shows that these phenomena are in fact not essentially different from the electrostatic forces whose demonstration was a society game in the 18th-century French court. If one charges a needle and brings small pieces of paper into its vicinity, the pieces stick to the metal. The strong electric field at the tip polarizes the pieces, pulling their electrons onto one side and leaving a net positive charge on the other, essentially making small electric dipoles. The attraction between the needle and the charges on the near side of the paper exceeds the repulsion between the needle and those on the far side, creating a net attractive force.

The atom and the cavity contain the same ingredients, albeit at a quantum level. The vacuum field bounded by the cavity walls polarizes the Rydberg atom, and the spatial variations of the field produce a net force. The atomic dipole and the vacuum field are oscillating quantities, however, and their respective oscillations must maintain a constant relative phase if a net force is to continue for any length of time. As it turns out, the photon exchange process does in fact lock the atomic dipole and the vacuum fluctuations.

The tiny force experienced by the atom is enhanced by adding photons to the cavity. The atom-cavity exchange frequency increases with the field intensity, so that each photon adds a discrete quantum of height to the potential barrier in the repelling state and a discrete quantum of depth to the potential well in the attractive state. As a result, it should be possible to infer the number of photons inside the cavity by measuring the time an atom with a known velocity takes to cross it or, equivalently, by detecting the atom's position downstream of the cavity at a given time.

One could inject perhaps a dozen or so photons into a cavity and then launch through it, one by one, Rydberg atoms whose velocity is fixed at about a meter per second. The kinetic energy of these atoms would be greater than the atom-cavity potential energy, and they would pass through the cavity after experiencing a slight positive or negative delay, depending on the sign of the atom-cavity detuning. To detect the atom's position after it has passed through the cavity, researchers could fire an array of field ionization detectors simultaneously some time after the launch of each atom. A spatial resolution of a few microns should be good enough to count the number of photons in the cavity.

Before measurement, of course, the photon number is not merely a classically unknown quantity. It also usually contains an inherent quantum uncertainty. The cavity generally contains a field whose description is a quantum wave function assigning a complex amplitude to each possible number of photons. The probability that the cavity stores a given number of photons is the squared modulus of the corresponding complex amplitude.

The laws of quantum mechanics say that the firing of the detector that registers an atom's position after it has crossed the cavity collapses the ambiguous photon-number wave function to a single value. Any subsequent atom used to measure this number will register the same value. If the experiment is repeated from scratch many times, with the same initial field in the cavity, the statistical distribution of photons will be revealed by the ensemble of individual measurements. In any given run, however, the photon number will remain constant, once pinned down.

This method for measuring the number of photons in the cavity realizes the remarkable feat of observation known as quantum nondemolition. Not only does the technique determine perfectly the number of photons in the cavity, but it also leaves that number unchanged for further readings.

Although this characteristic seems to be merely what one would ask of any measurement, it is impossible to attain by conventional means. The ordinary way to measure this field is to couple the cavity to some kind of photodetector, transforming the photons into electrons and counting them. The absorption of photons is also a quantum event, ruled by chance; thus, the detector adds its own noise to the measured intensity. Furthermore, each measurement requires absorbing photons; thus, the field irreversibly loses energy. Repeating such a procedure therefore results in a different, lower reading each time. In the nondemolition experiment, in contrast, the slightly nonresonant atoms interact with the cavity field without permanently exchanging energy.

Quantum optics groups around the world have discussed various versions of quantum nondemolition experiments for several years, and recently they have begun reducing theory to practice. Direct measurement of an atom's delay is conceptually simple but not very sensitive. More promising variants are based on interference effects involving atoms passing through the cavity—like photons, atoms can behave like waves. They can even interfere with themselves. The so-called de Broglie wavelength of an atom is inversely proportional to velocity; a rubidium atom traveling 100 meters per second, for example, has a wavelength of 0.45 angstrom.

If an atom is slowed while traversing the cavity, its phase will be shifted by an angle proportional to the delay. A delay that holds an atom back by a mere 0.22 angstrom, or one half of a de Broglie wavelength, will replace a crest of the matter wave by a trough. This shift can readily be detected by atomic interferometry.

If one prepares the atom itself in a superposition of two states, one of which is delayed by the cavity while the other is unaffected, then the atomic wave packet itself will be split into two parts. As these two parts interfere with each other, the resulting signal yields a measurement of the phase shift of the matter wave and hence of the photon number in the cavity. Precisely this experiment is now under way at our laboratory in Paris, using Rydberg atoms that are coupled to a superconducting cavity in an apparatus known as a Ramsey interferometer.

Such an apparatus has many potential uses. Because the passing atoms can monitor the number of photons in a cavity without perturbing it, one can witness the natural death of photons in real time. If a photon disappears in the cavity walls, that disappearance would register immediately in the atomic interference pattern. Such experiments should provide more tests of quantum theory and may open the way to a new generation of sensors in the optical and microwave domains.

FURTHER READING

RADIATIVE PROPERTIES OF RYDBERG STATES IN RESONANT CAVITIES. S. Haroche and J.-M. Raimond in *Advances in Atomic and Molecular Physics,* Vol. 20, pages 350–409; 1985.

THE SINGLE ATOM MASER AND THE QUANTUM ELECTRODYNAMICS IN A CAVITY. H. Walther in *Physica Scripta,* Vol. T23, pages 165–169; 1988.

CAVITY QUANTUM ELECTRODYNAMICS. S. Haroche and D. Kleppner in *Physics Today,* Vol. 42, No. 1, pages 24–30; January 1989.

CAVITY QUANTUM ELECTRODYNAMICS. E. A. Hinds in *Advances in Atomic, Molecular, and Optical Physics,* Vol. 28, pages 237–289; 1991.

CAVITY QUANTUM OPTICS AND THE QUANTUM MEASUREMENT PROCESS. P. Meystre in *Progress in Optics,* Vol. 30. Edited by E. Wolf. Elsevier Science Publishers, 1992.

CAVITY QUANTUM ELECTRODYNAMICS. S. Haroche in *Fundamental Systems in Quantum Optics.* Proceedings of Les Houches Summer School, Session LIII. Edited by J. Dalibard, J.-M. Raimond and J. Zinn-Justin. North-Holland, 1992.