

Slowing a beam of light to a halt may pave the way for new optical communications technology, tabletop black holes and quantum computers

Frozen LIGHT

EVERYONE KNOWS OF THE SPEED OF LIGHT AS one of the unshakable properties of the universe. It's not surprising, then, that experiments to radically alter light's speed require some serious equipment and hard work. Running such an experiment requires first a careful tune-up and optimization of the setup and then a long period of painstaking data-gathering to get a consistent set of measurements. At the Rowland Institute for Science in Cambridge, Mass., our original ultra-slow-light experiments took place in stints lasting 27 hours nonstop. Instead of breaking for meals, we learned to balance a slice of pizza in one hand, leaving the other clean to flip mirrors in and out on the optics table during 38 seconds of total darkness at a crucial stage of each run.

Our goal was to drastically slow down light, which travels through empty space at the universe's ultimate speed limit of nearly 300,000 kilometers a second. We saw the first sign of light

pulses slowing down in March 1998. As happens so often in experimental physics—because it can take so many hours to get all the components working together for the first time—this occurred in the wee hours of the morning, at 4 A.M. By July we were down to airplane speed. At that time I had to go to the Niels Bohr Institute in Copenhagen to teach a class. I remember sitting in the plane marveling that I was traveling “faster than light”—that I could beat one of our slow pulses to Denmark by a full hour.

Needless to say, I was restless during the week in Copenhagen and eager to get back to Cambridge to continue the light-slowness experiments. In the next month we reached 60 kilometers an hour and decided that it was time to publish. The real payoff for the hard work, prior to those results, was sitting in the lab in the middle of the night and observing the slow-light pulses, knowing that we were the first in history to

BY LENE VESTERGAARD HAU





FREEZING OF LIGHT begins with a process in which a carefully tuned laser beam renders an opaque material transparent to a second laser beam.

COPYRIGHT 2003 SCIENTIFIC AMERICAN, INC.

Compressing a kilometer-long laser pulse to one millionth of a meter sets off quantum shock waves in sodium atoms near absolute zero.

see light go so slowly that you could outpace it on a bicycle.

In the summer of 2000 we brought pulses of light to a complete halt within tiny gas clouds cooled to near absolute zero. We could briefly keep the pulses on ice, so to speak, and then send them back on their way.

As well as being of great intrinsic interest, slowing and freezing light have a number of applications. At sufficiently low temperatures the ultracold clouds of atoms used in our slow-light experiments form Bose-Einstein condensates, remarkable systems in which all the atoms gather in a single quantum state and act in synchrony. Further uses could involve sending a light pulse through a condensate as slowly as a sound wave, which we expect would cause a wave of atoms to “surf” on the light pulse, setting off oscillations of the entire condensate.

The slow and frozen light work also opens up new possibilities for optical communications and data storage and for quantum-information processing—that is, for quantum computers, which would utilize quantum phenomena to outperform conventional computers. The freezing-light system essentially converts between motionless forms of quantum information and photons flying around at the usual speed of light.

Getting Atoms into a State

MANY ORDINARY MATERIALS slow down light. Water, for instance, slows light to about 75 percent of its velocity in a vacuum. But that type of speed reduction, associated with a material's refractive index, is limited. Diamond, which has one of the highest refractive indices of a transparent material, slows light by a factor of only 2.4. Reducing light's speed by factors of tens of millions requires new effects that depend on quantum mechanics. My group produces the conditions for these effects in a cigar-shaped cloud of sodium atoms—typically 0.2

millimeter long and 0.05 millimeter in diameter—trapped in a magnetic field and cooled to within a millionth of a degree of absolute zero.

Sodium belongs to the family of alkali atoms, which have a single outermost, or valence, electron. The valence electron produces almost all the action: Different excited states of a sodium atom correspond to that electron's being promoted to larger orbits around the nucleus, with higher energies than its usual lowest energy state, or ground state. These states determine how the atom interacts with light—which frequencies it will absorb strongly and so on. In addition, both the valence electron and the atom's nucleus are magnets, in effect acting like tiny compass needles. The electron's magnetism is associated with its intrinsic angular momentum, or spin, a little like the association of the earth's rotational axis with magnetic north but with exact alignment. The precise energies of an atom's excited states depend on how the spins of the nucleus and the valence electron are aligned.

Although an atom can assume many states, we use just three to slow light. When we finish preparing and cooling the atom cloud, every atom is internally in state 1, its ground state: the valence electron is in its lowest orbit, and its spin is exactly opposite, or anti-aligned, with the nuclear spin. Also, the total magnetism of each atom is anti-aligned with the magnetic field that we use to hold the cloud in place. State 2 is very similar, only with the electron and nuclear spins aligned, which raises the atom's energy a little. State 3 has about 300,000 times more energy than state 2 (with state 1 as the reference level) and is produced by boosting the valence electron up to a larger orbit. Atoms relaxing from state 3 down to state 1 or 2 generate the characteristic yellow glow of sodium streetlights.

The pulse of laser light that we wish to slow (the “probe” pulse) is tuned to the energy difference between states 1 and 3. If we sent a pulse of that light into the cloud without doing any other preparation, the atoms would completely absorb the pulse and jump from state 1 to state 3. After a brief time, the excited atoms would relax by reemitting light, but at random and in all directions. The cloud would glow bright yellow, but all information about the original light pulse would be obliterated.

To prevent this absorption, we use electromagnetically induced transparency, a phenomenon first observed in 1990 by Stephen E. Harris's group at Stanford University. In electromagnetically induced transparency, a laser beam with a carefully chosen frequency shines on the cloud and changes it from being as opaque as a wall to being as clear as glass for light of another specific frequency.

The transparency-inducing laser beam, or coupling beam,

Overview/Stopping Light

- Nothing travels faster than light in a vacuum, but even light is slowed down in many media. Scientists have manipulated clouds of atoms with lasers so that pulses of light travel through the clouds at one 20-millionth of their normal speed—slower than highway traffic.
- A similar technique completely halts the pulses, turning them into a quantum imprint on the atoms. Later, another laser beam converts the frozen pulse back into a moving light pulse with all the properties of the original.
- The process of slowing and stopping light has many research and technological applications.

is tuned to the energy difference between states 2 and 3. The atoms, in state 1, cannot absorb this beam. As the light of the probe laser pulse, tuned to state 3, arrives, the two beams shift the atoms to a quantum superposition of states 1 and 2, meaning that each atom is in both states at once. State 1 alone would absorb the probe light, and state 2 would absorb the coupling beam, each by moving atoms to state 3, which would then emit light at random. Together, however, the two processes cancel out, like evenly matched competitors in a tug of war—an effect called quantum interference. The superposition state is called a dark state because the atoms in essence cannot see the laser beams (they remain “in the dark”). The atoms appear transparent to the probe beam because they cannot absorb it in the dark state. Which superposition is dark—what ratio of states 1 and 2 is needed—varies according to the ratio of light in the coupling and probe beams at each location. But once the system starts in a dark state (in this case, 100 percent coupling beam and 100 percent state 1), it adjusts to remain dark even when the probe beam lights up.

A similar cancellation process makes the refractive index exactly one—like empty space—for probe light tuned precisely to state 3. At very slightly different frequencies, however, the cancellation is less exact and the refractive index changes. A short pulse of light “sniffs out” this variation in the index because a pulse actually contains a small range of frequencies. Each of these frequency components sees a different refractive index and therefore travels at a different velocity. This velocity, that of a continuous beam of one pure frequency, is the phase velocity. The pulse of light is located where all these components are precisely in sync (or, more technically, in phase). In an ordinary medium such as air or water, all the components move at practically the same velocity, and the place where they are in sync—the location of the pulse—also travels at that speed. When the components move with the range of velocities that occurs in the transparent atoms, the place where they are in sync gets shifted progressively farther back; in other words, the pulse is slowed. The velocity of the pulse is called the group velocity, because the pulse consists of a group of beams of different frequencies.

This process differs in a number of important respects from the usual slowing of light by a medium with a refractive index greater than one: the group velocity is slowed, not the phase velocity; a steep variation of the refractive index with frequency, not a large value of the index itself, causes the slowing; and the coupling laser beam has to be on the entire time.


Ultracold Atoms for Freezing Light

THE MORE RAPIDLY the refractive index changes with frequency, the slower the pulse travels. How rapidly the index can change is limited by the Doppler effect: the incessant thermal motion of the atoms in the gas smears out each atomic state across a small range of energies. The Doppler effect is like the change in tone of a siren moving toward or away from you. Imagine the cacophony you would hear if many police cars were racing toward and away from you at various speeds.


My research group uses extremely cold atoms (which move

slowly) to minimize this Doppler spreading. Consequently, the energy states are sharply defined, and the variation of refractive index can be made very steep. Slow light in hot gases has since been obtained by Marlan O. Scully’s group at Texas A&M University and Dmitry Budker’s group at the University of California at Berkeley. The use of hot atoms removes the need to produce ultracold atoms, but it puts severe constraints on, for example, the geometry of the setup; the probe and coupling beams must propagate in exactly the same direction.

We chill our sodium atoms with a combination of laser beams, magnetic fields and radio waves. The atoms first emerge from a hot source as an intense beam, traveling about 2,600 kilometers an hour. A laser beam hits the atoms head-on and in a millisecond slows them to 160 kilometers an hour—a deceleration of 70,000 times gravity produced by a laser beam that wouldn’t burn your finger. Further laser cooling in an optical molasses—six beams bathing the atoms from all sides—chills the atoms to 50 millionths of a degree above absolute zero. In a few seconds we accumulate 10 billion atoms in the molasses. Next we turn off the laser beams, plunging the lab into total darkness, and turn on electromagnets, whose combined field holds the atom cloud like a trap. For 38 seconds we



WAYLAYING LIGHT: Before the light pulse [yellow] reaches the cloud of cold atoms [purple] that will freeze it, all the atoms’ spins [small arrows] are aligned and a coupling laser beam [red] renders the cloud transparent to the pulse [1, 2]. The cloud greatly slows and compresses the pulse [3], and the atoms’ states change; an imprint of the pulse is created in the cloud, and it accompanies the slow light. When the pulse is fully inside the cloud [4], the coupling beam is turned off [5], halting the imprint and the light; at zero velocity the light vanishes but the imprint stays. Later the coupling beam is turned on [6], regenerating the light pulse and setting the imprint and the light back in motion.



OPTICAL PROPERTIES induced in a cloud of atoms by a carefully tuned laser beam are the key to the light-slowing process. A coupling laser beam passing through the cloud makes it transparent to light of a precise frequency (*top*) and causes an associated sharp variation of its refractive index (*bottom*). The transparency allows properly tuned light to pass through the cloud without being absorbed, and the steeper the change in the refractive index, the slower the light travels.

cool the atoms through evaporation, kicking out the hotter atoms and leaving the cooler ones behind. Specially tuned radio waves help to speed the hot atoms on their way. This whole process—from hot atom beam to cold, trapped atoms—takes place inside a vacuum chamber pumped out to 10^{-14} (10 quadrillionths) of atmospheric pressure.

When we cool the cloud to about 500 billionths of a degree, it forms a Bose-Einstein condensate, a very odd state of matter in which the several million atoms left after the evaporative cooling behave in a completely synchronized fashion [see “The Coolest Gas in the Universe,” by Graham P. Collins; *SCIENTIFIC AMERICAN*, December 2000]. These ultracold atom clouds, freely suspended in the middle of a vacuum chamber by a magnetic field, are the coldest places in the universe. And yet the rest of the experimental setup, within one centimeter of the cloud, is at room temperature. Vacuum-sealed windows on the chamber let us see the atoms directly by eye during laser cooling: a cold atom cloud in the optical molasses looks like a little bright sun, five millimeters in diameter. Such easy optical access allows us to massage the atoms with laser beams and make them do exactly what we want.

When our cigar of cold atoms is in place, we illuminate it from the side with the coupling laser. Then we launch a probe pulse along the axis of the cigar. To measure the speed of the light pulse, we do the most direct measurement imaginable: we sit behind the atom cloud with a light detector and wait for the pulse to come out, to see how long it takes. Immediately after the pulse

has gone through, we measure the length of the cloud with yet another laser beam, shone from below to project the cloud’s shadow onto a camera. That length divided by the delay of the pulse gives us the velocity. The delays are in the range of microseconds to milliseconds; this might sound short, but it is equivalent to light taking a detour through kilometers of optical fiber wound in a coil—and our clouds are only 0.1 to 0.2 millimeter long.

When we slow a light pulse down by a factor of 20 million, more happens than just a change of speed. At the start our pulse of light is a kilometer long, racing through the air at nearly 300,000 kilometers a second. (Of course, our laboratory’s length is much less than a kilometer, but if we could place our laser that far away, its pulses would be that long in the air.) The pulse’s leading edge crosses the glass window into the vacuum chamber and enters our levitating speck of sodium atoms. Inside this tenuous cloud the light travels at 60 kilometers an hour. A cyclist on a racing bike could overtake such sluggish light.

Through the Gas, Darkly

WITH THE FRONT of the light pulse traveling so slowly and its tail still going full tilt through the air, the pulse piles into the gas like a concertina. Its length is compressed by a factor of 20 million to a mere twentieth of a millimeter. Even though the atom cloud is small, the pulse compresses so much that it fits completely inside—important for stopping light. One might expect the light’s intensity to increase greatly because the same amount of energy is crammed into a smaller space. This amplification does not happen, however; instead the electromagnetic wave remains at the same intensity. To put it another way, in free space the pulse contains 50,000 photons, but the slow pulse contains $1/400$ of a photon (the factor of 20 million again). What has happened to all the other photons and their energy? Some of that energy goes into the sodium atoms, but most of it is transferred to the coupling laser beam. We have monitored the intensity of the coupling laser to observe this energy transfer directly.

These transfers of energy also change the states of the sodium atoms where the pulse is passing by. At the front of the pulse the atoms are changed from their original state 1 to a superposition of states 1 and 2, the dark state. This state has the largest proportion of state 2 at the central, most intense part of the pulse. As the rear of the slow pulse leaves a region of atoms, the atoms revert to state 1. The spatial pattern of atomic dark states in the cloud mimics the shape of the compressed slow-light pulse and accompanies it through the gas as an imprint. When this imprint and the light pulse reach the end of the gas cloud, the light pulse sucks energy back out of the atoms and the coupling beam

THE AUTHOR

LENE VESTERGAARD HAU is a MacArthur Fellow and Gordon McKay Professor of Applied Physics and professor of physics at Harvard University. She received her Ph.D. in theoretical solid state physics from the University of Århus in Denmark. The author wishes to thank the team of Zachary Dutton, Christopher Slowe, Chien Liu, Cyrus H. Behroozi, Brian Busch and Michael Budde, as well as Stephen E. Harris of Stanford University, for an extremely fruitful collaboration.

to dash away through the air at its customary 300,000 kilometers a second, restored to its original kilometer of length.

The velocity of the slow light pulse depends on several parameters. Some of these are fixed once we choose our atom species and which excited states to use, but two of the variables are under our control: the density of the atom cloud and the intensity of the coupling laser beam. Increasing the cloud's density decreases the light's speed, but we can push that only so far, in part because very dense clouds leak atoms out of the magnetic trap too rapidly. The pulse speed is also reduced if the coupling laser beam is weaker. Of course, if the coupling laser is too feeble, the cloud will not be transparent and will absorb the pulse. Nevertheless, we can still achieve the ultimate in slowing without losing the pulse to absorption if we turn off the coupling laser beam while the compressed slowed pulse is contained in the middle of the gas.

In response, the light pulse comes to a grinding halt and

turns off. But the information that was in the light is not lost. Those data were already imprinted on the atoms' states, and when the pulse halts, that imprint is simply frozen in place, somewhat like a sound recorded on a magnetic tape. The stopping process does not compress the pattern of states, because the imprinted atoms in unison force the light pulse to turn off, unlike the earlier stage in which the pulse gradually entered the gas.



The frozen pattern imprinted on the atoms contains all the information about the original light pulse. We effectively have a hologram of the pulse written in the atoms of the gas. This hologram is read by turning the coupling laser back on. Like magic, the pulse reappears and sets off in slow motion again, along with the imprint in the atoms' states, as if nothing had interrupted it.

We can store the light for several milliseconds, long enough for a pulse to travel hundreds of kilometers in air. The pulse does degrade the longer it is stored, because even though the gas

A BENCHTOP GUIDE TO STOPPING LIGHT

EXPERIMENTAL SETUP

Three laser beams and an ultracold cloud of sodium atoms (*size exaggerated*) in a high vacuum lie at the heart of the slow-light experiment. The coupling beam interacts with the cloud, making it transparent but molasseslike to a pulse of the probe beam. A photomultiplier tube measures the pulse's time of arrival to better-than-microsecond precision. The imaging beam then measures the length of the cloud by projecting its shadow onto a camera. Not shown are the system that delivers and cools a new ultracold cloud for each pulse, electromagnets whose combined field holds the atoms in place, and additional details of the optics.



WHAT STOPPED LIGHT LOOKS LIKE

The precise times of detection of light pulses reveal the slowing and stopping of light. With no atom cloud present, the input pulse is detected at time "zero" (*top*). Slowing of the pulse by a cloud is revealed by the pulse's delay (*dotted curve*). To stop a pulse, the coupling beam (*bottom*) is turned off while the slowed pulse is inside the cloud. The time that the pulse is stopped—about 40 microseconds—adds to its delay. The slowed pulse loses intensity because the cloud is not perfectly transparent, but the pulse is stopped and revived with 100 percent efficiency.

CREATING QUANTUM SHOCK WAVES

IT TAKES A HURRICANE

FROZEN LIGHT can be used to examine the nature of superfluids, which flow without friction when close to absolute zero. To probe a Bose-Einstein condensate—a sodium superfluid—a light pulse was stopped at a roadblock within the condensate (*at the large black hemisphere in micrographs*), compressing a kilometer-long pulse to 0.001 millimeter. Within the compressed light pulse region, the sodium atoms were converted from their ground state to a slightly energized state, which leaves a narrow void in the condensate. This action causes density “dimples” that propagate out toward the boundaries, at the speed of sound in the condensate, which almost immediately forms quantum shock waves (moving to the right, at time 0 milliseconds). The fronts of the shock waves curve up (at 0.5 millisecond), creating quantized vortices with zero atom density at their centers (two white voids at 2.5 milliseconds)—the superfluid analogue of the eye of a hurricane. At 5 milliseconds there is a long sausage-shaped region (*black*) with high atom density, which relaxes over time and represents the presence of a large collective excitation of the condensate.



atoms are very cold they still move a bit, causing the pattern of dark states to diffuse slowly. In addition, collisions between atoms can disrupt the dark superposition states. After some milliseconds, the revived output pulse will begin to be weaker than the original. Yet these limits show that cold atoms allow for long storage times of immensely compressed optical information. And storage times are maximized with condensates in which atom collisions tend not to destroy the dark states.

We can also play some tricks. If the coupling beam is turned back on at a higher intensity, the output pulse will be brighter but shorter than the one we sent in. Turning the coupling beam on and off quickly several times regenerates the stored pulse in several pieces. Such manipulations demonstrate the degree of control that we have over stored pulses and may be extremely useful in future applications.

Since our initial observation in 2000 of stopped light, it has been obtained in a hot gas by Ronald L. Walsworth and Mikhail D. Lukin of the Harvard-Smithsonian Center for Astrophysics in Cambridge, Mass., and in a cooled, doped solid by Philip R. Hemmer, then at the Air Force Research Laboratory in Hanscom, Mass.

Quantum Shock Waves

SLOW AND STOPPED LIGHT open up many interesting experiments. In our latest trial we turned everything around: rather than using cold atoms and condensates to slow light, we used slow light pulses to probe the odd properties of Bose-Einstein condensates.

A condensate is a superfluid—it flows through a tube without friction, just as electric current flows through a superconductor without resistance. Helium turns superfluid when it is cooled to very low temperatures. Once a superfluid gets going, it flows forever with no need for further energy input. Superfluids are described by a characteristic length scale called the healing length, the minimum distance over which a superfluid can adjust (heal itself) to an external perturbation. In our condensates, the healing length is about 0.001 millimeter. So if we can compress our light pulses to that scale, we should be able


to cause dramatic excitations of the superfluid, helping us to learn about the superfluid state.

We have succeeded in compressing a light pulse this drastically by creating a “roadblock” that forces the pulse to stop at a particular point in space, where it becomes totally compressed and localized. This is possible because when the probe light pulse is propagating at right angles to the coupling laser, we can spatially manipulate the intensity of the coupling laser along the propagation direction of the incoming light pulse. The speed of the pulse is controlled by the coupling laser intensity: the lower the intensity, the lower the speed.


To create the roadblock, we illuminate only the front half of the condensate with the coupling laser. When the light pulse enters the atom cloud, it slows to bicycle speed as before. Then, as the pulse runs into the roadblock region in the middle of the condensate, the pulse really slows down and compresses. Within the light pulse region at the roadblock, the atoms are almost entirely in state 2 (the dark state) because the coupling intensity is very low. Outside, the atoms are all in state 1. So we have a way of creating very localized defects in the condensate. We have taken direct images of this process in our laboratory.

As described earlier, for an atom in state 2, the spin of the valence electron is aligned opposite to the electron spin of an atom in state 1. This has severe consequences, because the electromagnet is used to trap atoms and hold them in place; atoms in state 2 will get kicked out of the magnet, as if a magnetic north pole was facing another north pole. So the localized defect of atoms in state 2 will zip out of the magnet in less than 0.5 millisecond. That process is by itself interesting and gives rise to what is called a pulsed atom laser.

Let's concentrate, however, on what is left in the trap. We end up with a condensate of state 1 atoms, with a hole punched in the middle. As you might imagine, condensates do not like holes punched in their middles, and the response we observed is fascinating. Two density “dimples” are created. They start propagating out toward the condensate boundaries, at the speed of sound in the condensate, about 1 to 10 millimeters a second. During this process, the back edge of the dimples steep-



0.5 millisecond



2.5 milliseconds



5 milliseconds

ens because there is a dramatic variation of the atom density and thus the sound speed across the dimple structures. The back part of each dimple, where the density is high, will catch up to the central part, where the density is the lowest. Such a steepening of the back edge is what in a normal fluid, like water, would lead to shock wave formation. But because a condensate is a superfluid, the steepening forms the superfluid analogue of that phenomenon—a quantum shock wave, a quantized rotation pattern, a quantized vortex shaped like a smoke ring.

Our lab has photographed this process [see illustration above]. The central parts of the quantized vortices appear as white dots. In these places there is no atom density and the light just goes right through the condensate, as in the eye of a hurricane.

This process shows how the superfluid nature of the condensate will break down during the formation of these dramatic excitations. The vortices even seem to have a life of their own: they move around and sometimes bounce off one another like billiard balls. At other times they collide and seem to explode into a bunch of sound waves. Because a slow light pulse always creates vortices in pairs of opposite circulation, they are very much like particle and antiparticle pairs that can annihilate. We are excited about probing this rich system further.

Black Holes and Computers

CONDENSATES CAN BE PRODUCED in a vortex state wherein the gas rotates, reminiscent of water going down the drain. Ulf Leonhardt of the University of St. Andrews in Scotland has suggested that a pulse of slow light traveling through a vortex would find itself dragged along with the gas—very similar to a phenomenon believed to occur near black holes. With slow light, we can perhaps study this and some other black hole phenomena in the laboratory.

Slow light also enables a new kind of nonlinear optics that occurs, in particular, when one laser beam alters the properties of another beam. Nonlinear optics is a huge field of research, both of fundamental interest and with applications from imaging to telecommunications. Extremely intense beams are usually needed to achieve nonlinear optical effects, but with slow light

the corresponding phenomena can be produced with a very small number of photons. Such effects could be useful for creating ultrasensitive optical switches.

Another application for slow and stopped light could be quantum computers, in which the usual 1's and 0's are replaced with quantum superpositions of 1's and 0's called qubits. Such computers, if they can be built, could solve certain problems that would take an ordinary computer a very long time. Two broad categories of qubits exist: those that stay in one place and interact with one another readily (such as quantum states of atoms) and those that travel rapidly from place to place but are difficult to make interact in the ways needed in a quantum computer (photons). The slow-light system, by transforming flying photons into stationary dark state patterns and back, provides a robust way to convert between these types of qubits, a process that could be essential for building large-scale quantum computers. We can imagine imprinting two pulses in the same atom cloud, allowing the atoms to interact and then reading out the result by generating new output light pulses.

Even if frozen light doesn't prove to be the most convenient and versatile component for building a quantum computer, it has opened up more than enough research applications to keep us—and other groups—busy for many more all-night sessions in the years to come. SA

MORE TO EXPLORE

Electromagnetically Induced Transparency. Stephen E. Harris in *Physics Today*, Vol. 50, No. 7, pages 36–42; July 1997.

Light Speed Reduction to 17 Metres per Second in an Ultracold Atomic Gas. Lene Vestergaard Hau, S. E. Harris, Zachary Dutton and Cyrus H. Behroozi in *Nature*, Vol. 397, pages 594–598; February 18, 1999.

Observation of Coherent Optical Information Storage in an Atomic Medium Using Halted Light Pulses. Chien Liu, Zachary Dutton, Cyrus H. Behroozi and Lene Vestergaard Hau in *Nature*, Vol. 409, pages 490–493; January 25, 2001.

Observation of Quantum Shock Waves Created with Ultra-Compressed Slow Light Pulses in a Bose-Einstein Condensate. Zachary Dutton, Michael Budde, Christopher Slowe and Lene Vestergaard Hau in *Science*, Vol. 293, pages 663–668; July 27, 2001.