



The Story Behind “Stopped Light”

Ronald Walsworth, Susanne Yelin, and Mikhail Lukin

Opaque media can be made transparent by use of quantum mechanical interference, enabling fascinating phenomena with potential applications in both classical and quantum information processing. In particular, recent laboratory demonstrations have shown that light pulses can be slowed down and effectively brought to a standstill in coherently prepared atomic ensembles, thereby reversibly transferring information between light and matter with potentially high efficiency.

The speed of light is intriguing: it is fast beyond human imagination, and yet its finite value sets the ultimate speed limit for the communication of information. Recently, laboratory demonstrations have shown that light pulses can be greatly slowed down, and even effectively halted and re-started, using coherently prepared atomic media. Here, we discuss the story behind this so-called “stopped light.” In particular, we describe how the recent demonstrations used nonlinear optical spectroscopy techniques, developed in the optics community over the past several decades, to map the information carried by a light pulse into a cloud of atoms by controlling the group velocity of the pulse. It is now possible to envision some potential applications of “slow” and “stopped” light, particularly in the area of quantum information processing. This emerging field is based on the idea of exploiting the basic laws of quantum mechanics to enable dramatic improvements in computing speed and communication security. Surprisingly, perhaps, “slowing” and “stopping” light pulses may serve to enhance the processing of information.

Electromagnetically induced transparency

Quantum interference can be used to make an opaque medium transparent. Usually, the propagation of a pulse or beam of light through a resonant medium is strongly affected by absorption. However, when two light beams are used, it is possible to create a window of transparency which allows both beams to propagate with (almost) no absorption. A typical model of such “electromagnetically induced transparency” (EIT) is shown in Fig. 1(a).

An atomic medium has two lower states, $|g_1\rangle$ and $|g_2\rangle$, which are often Zeeman or hyperfine levels. Each of these lower levels is coupled to a common electronically excited state, $|e\rangle$, by one of the two light beams (often called “signal” and “control” beams). The light beams drive the atoms into an intensity-dependent antisymmetric superposition of the lower states, such that the two absorption paths $|g_1\rangle \rightarrow |e\rangle$ and $|g_2\rangle \rightarrow |e\rangle$ interfere destructively, and the atoms can never reach the excited state.¹ For this reason, the two light beams can propagate through the medium without absorption. We will refer to $|g_1\rangle$

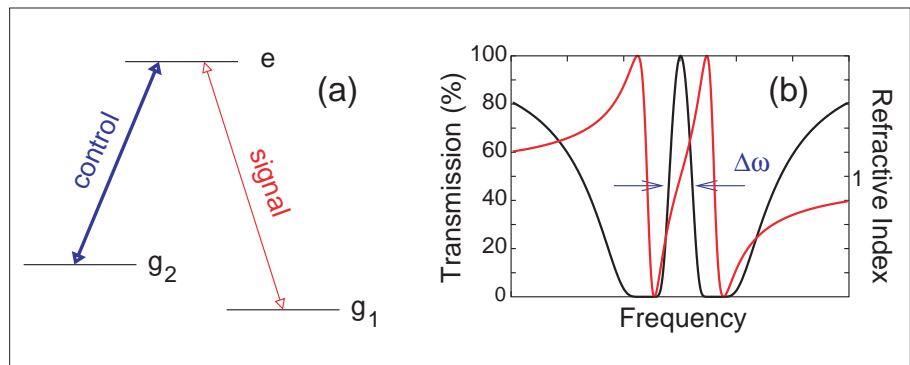


Figure 1. (a) Prototype atomic system for EIT. (b) Spectrum of transmission and refractive index corresponding to EIT. Rapid variation of the refractive index (red curve) causes a reduction of group velocity. “ $\Delta\omega$ ” signifies the width of the transparency frequency window.




Figure 2. A light pulse enters an EIT medium and exhibits a spatial compression due to a reduction in group velocity, while photons are converted into flipped spins. The “slow light” pulse (red curves) and spin waves (green circles) then propagate together. From top to bottom: pulse before entering medium, shortly after entering, shortly before exiting, after exiting.

and $|g_2\rangle$ as “spin states,” since they usually represent atomic states with different angular momentum, or spin.

The interference effect resulting in atom-light decoupling is extremely fragile. Even a very small mistuning of the lasers can destroy it, causing large absorption. The width of the transparency window [“ $\Delta\omega$ ” in Fig. 1(b)], i.e., the light frequency range for which transparency is created, is usually considerably narrower than the width of the original absorption line. (This is particularly true for an optically dense or Doppler-broadened medium.)

The robustness of the interference can, however, be improved by increasing the intensity of the control field. In particular, the width of the transparency window for the signal light is proportional to the control beam intensity. Also, the center of the transparency window is at the exact two-photon resonance of the two light fields, with the transition between the lower levels, i.e., $\hbar\omega_{\text{signal}} - \hbar\omega_{\text{control}} = E_{g_2} - E_{g_1}$.

“Slow light” in EIT media

Speaking of a single “speed of light” is incorrect. There are actually different speeds of propagation for the different aspects of a light pulse. For example, the *phase velocity* is the speed at which each maximum of the rapidly oscillating electric field travels for each frequency component that makes up a light pulse. In the atomic vapor EIT experiment we will discuss in this article, the phase velocity is very close to the well-known “speed of light,” $c = 3 \times 10^8$ m/sec. On the other hand, the *group velocity* is the speed at which the peak of a macroscopic pulse of light moves, and is usually relevant to the speed of information transfer. The group velocity v_g can be quite different from c . Whereas the phase velocity v_{ph} is slowed only by the index of refraction of a medium through which light moves, $v_{\text{ph}} = c/n$, the group velocity is inversely proportional to the dispersion of the medium, i.e., the variation of $n(\omega)$ with the light frequency ω . The very narrow transparency window created by EIT also creates a large variation of $n(\omega)$ with frequency [see Fig. 1(b), red curve]. It follows that v_g can be very small in an EIT

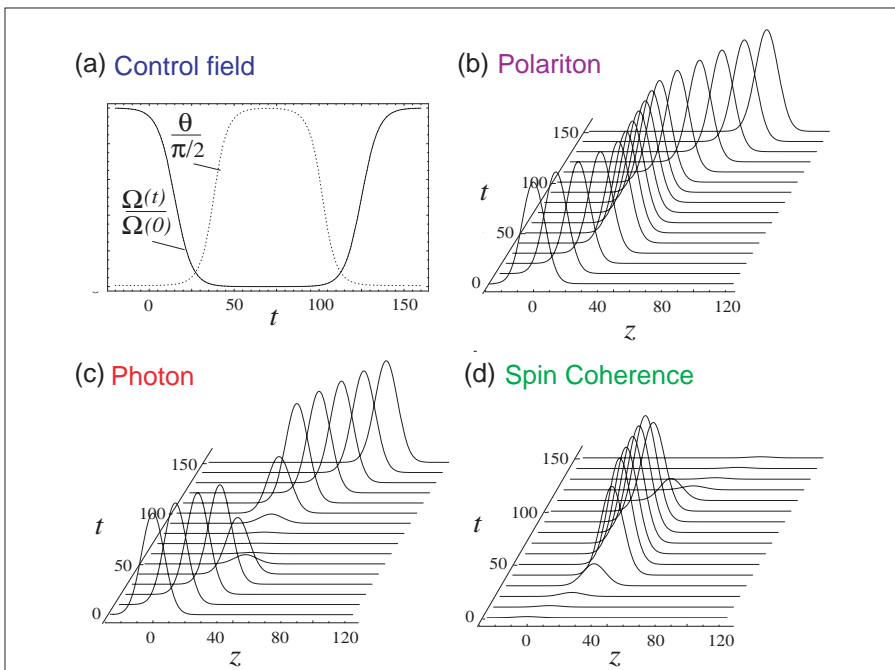


Figure 3. A dark-state polariton can be stopped and re-accelerated by ramping the control field intensity as shown in (a). The broken line shows the mixing angle between photonic and spin states θ . The coherent amplitudes of the polariton ψ , the electric field E of the photon, and the spin coherence S , are plotted in (b-d).

medium. The width of the transparency window, and thus v_g , is a function of the atomic density and the control beam intensity, and is therefore under experimental control. In particular, v_g decreases nearly linearly with both quantities.

A signal pulse of light entering an atomic EIT medium undergoes the following dynamics (see Fig. 2). The atoms are pumped into one of their ground states and coupled by the control beam. Upon entering the atomic storage cell, the front edge of the signal pulse abruptly slows down because of the medium’s low group velocity. Since the front edge is slowed first and the rear edge last, the signal pulse ends up considerably compressed by a factor of c/v_g once it is fully in the EIT medium. However, the peak amplitude of the signal pulse is unchanged by this spatial compression. Consequently, the signal pulse contains much less energy when in the EIT medium. Since the atomic gas gains only negligible amounts of energy (the excited state is never populated!), the excess energy must be carried away somehow: this job is performed by the control beam. A fraction of the photons that make up the signal pulse are converted into a “spin wave” via a two-photon

transition that maps the signal pulse into a coherent superposition of the atomic states, $|g_1\rangle$ and $|g_2\rangle$. In so doing, the excess energy of the signal photons is expended in the creation of new control beam photons. The resulting atomic spin wave travels together with the signal pulse at the reduced v_g in a coupled excitation known as a “dark-state polariton.” Upon reaching the far end of the cell, the signal pulse is fully reconstituted via the reverse process: i.e., pulse expansion and acceleration of v_g back to c . The only difference between this released signal pulse and a pulse which has propagated only through vacuum is that the EIT medium has induced a delay in the released pulse of $\tau = (1/v_g - 1/c)L$, where L is the cell length. This pulse delay process is often referred to as “slow light.”

The subject of slow light attracted great attention with the remarkable experiment conducted by Lene Hau and co-workers in 1999, in which an ultracold gas of Na atoms was used to slow light pulses to a group velocity of 17 m per second. Furthermore, this work made use of a large optical density to compress the pulse entirely into the medium. Soon thereafter, the groups of Marlan Scully and Dmitry Budker demonstrated very slow group ve-

locities using warm atomic vapor. All these experiments built upon important earlier work on EIT,² which had observed pulse delays considerably exceeding the pulse duration.¹

Even in an ideal EIT system, however, there is an important limitation to the slow light technique. Only signal frequencies that lie within the narrow transparency window can propagate unabsorbed through the medium. Since the window width, as well as the group velocity, are proportional to the control intensity, the maximum pulse delay is inversely proportional to the pulse length to be slowed. (The proportionality constant depends on the optical density of the EIT medium.) It is therefore not possible to reduce the group velocity to zero with the technique described above. One more twist to the EIT bag of tricks is needed.

“Stopping light”

To bring a light pulse to a complete stop (i.e., $v_g = 0$), “dynamic EIT” is employed.^{2,3} Once a light pulse is compressed into an EIT medium and propagating as a dark-state polariton, its properties can be modified by simply changing the intensity of the control beam: if the control beam intensity is decreased, then the group velocity is further slowed. This also implies that the contribution of photons to the polariton state is reduced. In particular, if the control beam is turned off after the signal pulse has been compressed into the EIT medium, two things happen: the polariton becomes purely atomic, and its group velocity is reduced to zero. At this point, information originally carried by the photons (pulse shape, amplitude, length, polarization, etc.) is fully mapped onto the long-lived ground or spin states of the atoms, $|g_1\rangle$ and $|g_2\rangle$. As long as this “light-storage” or “trapping” process is sufficiently smooth, the frequency spectrum of the polariton will narrow continuously as the pulse slows, so that it always remains within the transparency window. The stored light pulse can be easily retrieved by simply turning the control beam back on and re-accelerating the stopped dark-state polariton. For this reason, there is, in principle, no loss associated with the trapping procedure.

The light-storage process is illustrated in Fig. 3, which shows the evolution of the signal light pulse, spin coherence, and polariton, when the control beam is turned

off and on. The amplitude of the signal pulse decreases as it is decelerated, while the spin coherence grows. The amplitude of the polariton is constant, but its speed is changed, even slowing to zero. The procedure is reversed when the control beam is turned back on.

The smoothness (i.e., adiabatic) conditions for storing light are as simple as they are crucial: the entire light pulse should be within the EIT medium at the beginning of the trapping procedure; and the light pulse's frequency spectrum before it enters the EIT medium should be contained within the transparency window before the control beam is turned off.

An essential characteristic of the dynamic EIT technique is that almost no photon energy or linear momentum is stored in the EIT medium. Instead, both are transferred into (or borrowed from) the control beam in such a way that all other information that describes an entire light pulse is coherently converted into a low-energy spin wave. This near-perfect information-mapping efficiency is a key feature that distinguishes the dynamic EIT approach for storing light from earlier optical techniques such as photon echoes, and potentially enables applications in quantum information science.

A second distinguishing feature of the technique discussed here is that it involves many atoms in an optically dense system. This result can be compared to those achieved in the fascinating experiments carried out by the groups of Jeff Kimble and Gerhard Rempe, which involved single atoms coupled to a high-Q cavity to enhance the effective photon absorption cross-section. In our case, the atom-light coupling is enhanced by the number of atoms ($N \sim 10^{10}$ in atomic vapor experiments to date). Since it is not known which atom scatters a photon, the signal light pulse couples to collective spin states; that is, a whole ensemble of atoms “guards” a single spin excitation. For this reason, if one atom is lost from the collective atomic state in which the light-pulse information has been stored, the resulting state differs from the initial state only by a very small factor, which scales as $1/N$.

But this description represents an ideal scenario. In practice, several loss or “deco-

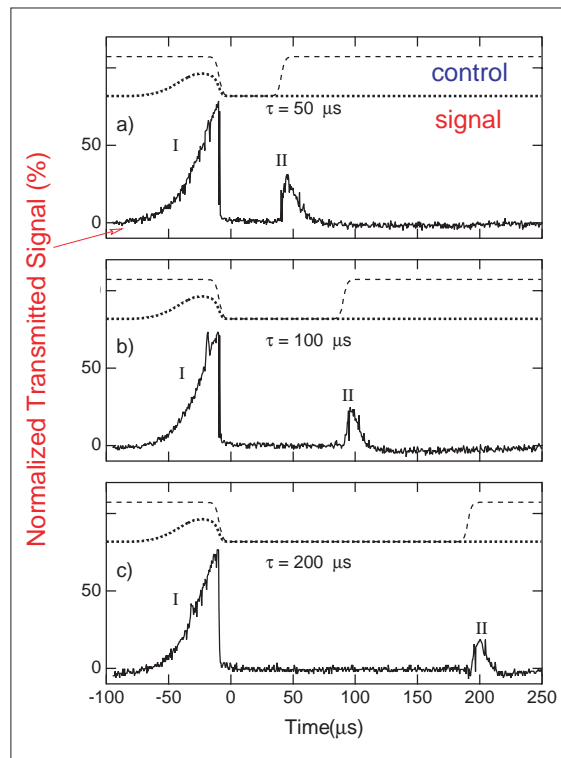


Figure 4. Observed light-pulse storage in a ^{87}Rb vapor cell. Examples are shown for storage times of (a) 50 μs , (b) 100 μs , and (c) 200 μs . Shown above the data in each graph are calculated representations of the timing of the applied control field (dashed line) and input signal pulse (dotted line).

herence” mechanisms must be taken into account. For this reason, the lifetime of the coherence created between the two atomic ground states, ($|g_1\rangle$ and $|g_2\rangle$), is always finite. This implies that after a characteristic length of time, τ_{coh} , the atoms will end up either in state $|g_1\rangle$ or in state $|g_2\rangle$, rather than in the “dark-state” superposition of the two; the transparency will then be lost and the dark-state polariton will dissipate.

Experimental demonstrations

Recent experiments have demonstrated that weak pulses of laser light can be “stopped” in coherently prepared atomic media, and then released after a storage interval of up to ~ 1 ms. Hau and co-workers trapped light pulses using ultracold atomic sodium slightly above the point of Bose-Einstein condensation.³ A stored-light experiment carried out at the same time by our group used warm rubidium vapor.⁴ Very recently, similar effects in solid-state media have been reported.⁵

In our warm rubidium experiments, the control beam and the signal pulse were represented, respectively, by the two helicities of circularly polarized light (σ_+ and σ_-), derived from a single laser beam by carefully controlling the light polarization. These light fields couple pairs of Zeeman sublevels, ($|g_1\rangle$, $|g_2\rangle$), of electronic ground state ($5^2S_{1/2}$) ^{87}Rb atoms, with magnetic quantum numbers differing by two, via the excited ($5^2P_{1/2}$) state. In this configuration, the effective two-photon detuning of the control beam and signal pulse can be controlled by applying an external magnetic field, which causes Zeeman level shifts between $|g_1\rangle$ and $|g_2\rangle$. (For comparison, the stored-light experiments involving ultracold Na atoms employed hyperfine sublevels, with control and signal beams separated in frequency by about 1.7 GHz.) Since the control beam in our system was always much stronger than the signal pulse, most

of the relevant atoms were in the $5^2S_{1/2}$, $F=2$, $M_F=+2$ magnetic sublevel. In this case, the states $|g_1\rangle$, $|g_2\rangle$ of the simplified three-level EIT model correspond, respectively, to $|F=2, M_F=0\rangle$ and $|F=2, M_F=+2\rangle$. By use of a fast Pockels cell, we slightly rotated the polarization of the input light to create a weak pulse of σ_- signal light out of the strong σ_+ control beam. The light-storage experiments were carried out with a glass cell containing rubidium at temperatures of ~ 70 – 90°C , which corresponds to atomic vapor densities of $\sim 10^{11}$ – 10^{12} cm^{-3} . To ensure long lifetimes of the atomic Zeeman coherences, we magnetically shielded the Rb cell and filled it with about 5 torr of He buffer gas.

Typical input Σ -signal pulses had a duration of ~ 10 – 30 μs , corresponding to a spatial length of several kilometers in free space. Upon entrance into the Rb cell, the signal pulse was spatially compressed to a length of a few centimeters because of the reduction in group velocity. To trap, store, and release the signal pulse, we used an acousto-optic modulator to turn off the control beam, over about 3 μs , while much of the signal pulse was contained in the Rb cell. After a certain interval, we turned the control field on again, thereby releasing the stored portion of the signal pulse.

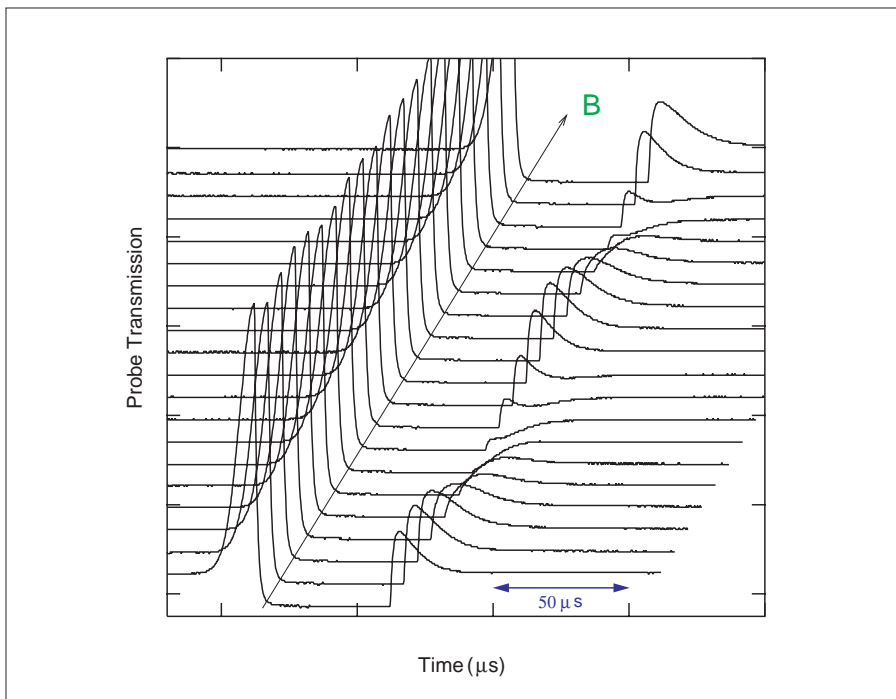


Figure 5. Results of optical interferometric measurements of released stored-light pulses demonstrate that this process is phase coherent. A magnetic field was pulsed during the storage interval with increasing strength from trace A to E, such that the accumulated phase difference between the output signal pulse and the control beam varied from approximately 0 to 4π .

An example of the observed storage process is shown in Fig. 4. Typically, two time-resolved σ_- signal pulses were registered by the photodetector. First, a fraction of the signal pulse left the cell before the control beam was turned off, which resulted in an observed signal that was not affected by the storage operation (peak I in each plot of Fig. 4). Because of the slow group velocity in the Rb vapor ($v_g \sim 1$ km/s), as compared to free-space propagation, this untrapped light was delayed by about $30 \mu\text{s}$. The second observed signal pulse was light that was stored in atomic excitations for a time interval τ_{store} . Note that the released signal pulse was detected only after the control beam was turned back on (peak II in each plot of Fig. 4). In our experiments, up to 50% of the input light pulse has been trapped for short storage times. We observed that the amplitude of the released signal pulse decreased as τ_{store} increased due to various loss mechanisms, such as atomic diffusion out of the laser beam. We could resolve released light pulses without signal averaging for storage intervals of up to $\tau_{\text{store}} \sim 0.5$ ms.

An important feature of the light-storage technique is that it is a coherent

process, that is, phase information is preserved in the mapping of information between light and atoms. As described in Ref. 6, we demonstrated this property by applying a pulsed magnetic field during the light-storage interval to vary controllably the phase of the collective Zeeman coherence in the Rb vapor. We then converted the spin excitations back into light and used optical interferometry to show that the phase shift had been mapped onto the light. To form an interferometer for the two fields, we adjusted our optics (a $\lambda/2$ plate) such that a small fraction ($< 10\%$) of the control beam was mixed into the signal light detection channel.

Figure 5 shows 20 stored light experiments for which we increased the Zeeman phase shift by approximately 0.2π for each successive run. Trace A in Fig. 5 shows the result for $\Phi \sim 0$, i.e., maximum constructive interference between the output signal light and the control beam. As we increased the pulsed magnetic field to change the phase by π , we observed destructive interference (e.g., trace B). As we increased the pulsed magnetic field further, we alternatively observed constructive and destructive interference as expect-

ed at $\Phi \sim 2\pi, 3\pi, 4\pi$, etc. (traces C-E in Fig. 5). We observed up to 10 periods of phase accumulation (i.e., $\Phi \sim 20\pi$) without loss of coherence.

Potential applications of “slow” and “stopped” light

EIT has already had a major impact on the field of optical science, although commercial applications have yet to emerge. One potential area of application is all-optical switching and signal processing in optical communication. The most serious roadblocks on this front are materials and speed issues. Good optical control requires long coherence times, and for this reason the majority of experiments have made use of atomic vapors that have relatively slow response. For practical communication systems, solid-state devices are desirable because of their low cost and the possibility of integration with existing technologies. Work in this area is being pursued by many groups.¹

In the exciting, emerging field of quantum information processing, “slow light” and “stopped light” may have key roles to play. We note that once a dark-state polariton is converted into a purely atomic excitation in a small sample, logic operations can be accomplished by promoting atoms into excited states with strong atom-atom interactions. In this case, the ability to interconvert quantum information between photons and atoms is essential for performing operations involving distant units, as well as for the scalability of such systems. Likewise, potential applications in the communication of quantum information over long distances might be implemented by combining an EIT-based quantum memory and linear optical elements.

In coming years, experiments will undoubtedly shed new light on the opportunities and practical limitations for classical signal processing and for quantum-state manipulation using EIT and its remarkable manifestations, including “stopped light.”

Acknowledgments

Our work has been supported by NSF, ONR, and NASA.

Ronald Walsworth and Susanne Yelin are with the Harvard-Smithsonian Center for Astrophysics, Cambridge, Massachusetts. Mikhail Lukin is with the Department of Physics, Harvard University, Cambridge, Massachusetts.