

Here, the line defect (a line of missing holes in the PC slab, which is usually used as a waveguide), is partially modified so that light guided in the waveguide is captured. Leak paths of the cavity mode are eliminated by optimizing the arrangement of the holes in the PC.

The importance of a Q factor of one million can manifest itself in many different ways. It means that light propagation is effectively slowed by a factor of a million, or stored optical energy is about a million times larger. In the report by the Japanese team, surprisingly slow light propagation and long storage times are demonstrated with a storage lifetime of longer than a nanosecond.

Unfortunately, the slow-down factor is intrinsically linked to the operating bandwidth, which is a critical factor in signal processing and optical communications — the longer the delay, the narrower the bandwidth (around 200 MHz in this case). To get around this problem, two solutions have been proposed — photonic-bandgap

engineering⁴ and the dynamic tuning of cavity properties⁵.

The bandgap-engineering attempts to reach an optimum balance between the slow-down factor and operating bandwidth. However, the relationship between the two still exists and constrains what can be achieved. On the other hand, tuning Q potentially breaks this constraint and allows wide-bandwidth frequency signals to be processed. Here, Q is first set to be low so that wideband signals can be captured by the cavity. Shortly after capture, Q is suddenly increased so that light is temporarily 'stored' within the cavity on a timescale of the storage lifetime. When Q is suddenly dropped, the light is then released from the cavity and the initial signals are recovered. The process can be considered to be equivalent to the writing and readout operation of an optical memory element.

Tanabe and colleagues have shown, for the first time, that it is possible to rapidly tune the cavity Q factor. They demonstrated a sudden drop of Q with a switching

time on the order of 10–100 ps by using photo-pumping to control carrier-induced changes in the refractive index.

Clearly such light-based memory still has some way to go — it does not yet fully support writing processes — but nevertheless this is an important step towards dynamical optical memory. By overcoming the constraint between the achievable slow-down factor and the signal bandwidth, high- Q nanocavities could be applied not only to optical memory but also to modulators, switches, wavelength converters and other devices that harness the large amount of stored optical energy and corresponding highly efficient nonlinearity⁶.

References

1. Yablonovitch, E. *Phys. Rev. Lett.* **58**, 2059–2062 (1987).
2. Tanabe, T., Notomi, M., Kuramochi, E., Shinya, A. & Taniyama, H. *Nature Photon.* **1**, 49–52 (2007).
3. Baba, T., Mori, D., Inoshita, K. & Kuroki, Y. *IEEE J. Sel. Top. Quantum Electron.* **10**, 484–491 (2004).
4. Mori, D. & Baba, T. *Appl. Phys. Lett.* **85**, 1101–1103 (2004).
5. Yanik, M. F. & Fan, S. *Phys. Rev. Lett.* **92**, 083901 (2004).
6. Soljacić, M. & Joannopoulos, J. D. *Nature Mater.* **3**, 211–219 (2004).

LASER WAKEFIELDS


Surf, you're on camera

In the past 20 years laser technology has blossomed to the point where table-top plasma accelerators are now a common sight. In these machines, electrons surf on electric fields 1,000 times stronger than those offered by conventional accelerators. The fields are created inside plasmas travelling close to the speed of light, in the wake of an intense laser pulse. Until now, the features of these laser wakefields were invisible. But a clever camera has revealed details not seen before.

Picturing the plasma behind the laser beam is important because its structure tells us about the quality — the divergence and energy spread — of the electron beam produced by the accelerator. Existing methods used to image wakefields are slow and unwieldy; they only photograph the plasma at individual points in time, therefore many snapshots have to be combined together to get a complete

dynamic image. In contrast, the approach adopted by researchers from Texas and Michigan simplifies things by illuminating the entire plasma at once (*Nature Phys.* **2**, 749–753; 2006). The result is precise information about the temporal and spatial behaviour of the wake.

Their technique, frequency-domain holography, is a twist on conventional holography geared towards imaging structures travelling at luminal velocities. An intense 800-nm, 30-fs pump laser is focused into a jet of helium gas, creating a plasma and laser wakefield. Meanwhile two chirped 400-nm, 1-ps pulses — a 'reference' beam in front and a 'probe' beam trailing 3 ps behind — ride along with the pump beam. Because the probe waves overlap with the pump pulse, they become imprinted with phase shifts that reveal information about the plasma disturbance. A two-dimensional holographic image is



obtained by detecting interference between the probe and reference beams with a spectrometer and a CCD. The wake structure is recovered by Fourier-transforming this data.

The resulting images can distinguish features smaller than the plasma frequency. They show sharp ionization fronts, followed by a wake containing a He²⁺ core

and an outer He⁺ 'corona'. A 30-TW laser pulse generates a wake with curved 'horse-shoe' wavefronts, arising from nonlinear laser-plasma interactions. But this new camera does more than just offer us a window into the world of relativistic laser-plasmas; it takes us a step closer to using plasma accelerators to their full potential.

Amber Jenkins