Optical-fiber microcavities reach angstrom-scale precision

Using heat and light to subtly vary the local radius and refractive index of a glass fiber is a simple and surprisingly reproducible way to create and tune a microresonator.

A long-standing goal of optical science is the development of devices such as miniature buffers, microlasers, optical switches, and filters that can be assembled into an all-optical computer—no electrons needed. The fundamental structure required for all of those circuit elements is the microresonator, whose highly reflecting walls can confine a light signal in a tiny volume for up to hundreds of microseconds.

If fashioned into dielectric toroids, disks, or spheres a few tens of microns in diameter, the resonators store the light in the form of whispering-gallery modes, so called because the optical waves circulate by total internal reflection around a perimeter just as acoustic waves do around the dome of Saint Paul's Cathedral in London. The resonance condition occurs whenever the roundtrip distance is an integral number of wavelengths.

The appeal of such whispering-gallery-mode resonators, particularly those made of silica, lies in their extremely high *Q* factor. A measure of light's survival time, the *Q* factor is the ratio of the resonance wavelength to its spectral width and can, in such resonators, reach up to 10^{10}. That's at least two orders of magnitude higher than is achievable from state-of-the-art lithographically etched resonators such as those in photonic crystals. The higher the *Q*, the longer the circulation of light, the sharper the resonance spectrum, and the more sensitive to refractive-index variations is each resonator in a circuit.

Getting individual resonators in series to efficiently pass signals between them, however, requires matching the resonance wavelengths in each one. But precise tuning has proven difficult because the resonances are determined by a microcavity's shape and size. In one promising approach developed over the past decade, researchers locally heat a glass fiber using a flame or carbon dioxide laser and pull it so that a bulge forms between tapered sections. In 2009 Arno Rauschenbeutel and colleagues from Johannes Gutenberg University used the technique to create a bottle-shaped microresonator whose local curvature axially confined whispering-gallery modes and could be accurately reproduced to within 2 μm.[1]

Misha Sumetsky and his colleagues at OFS Laboratories (a former branch of Bell Labs) now report an improvement in that precision to an astonishing 2 Å by modifying the approach in a way that could hardly be simpler: They opted not to stretch the fiber but merely to heat it locally to anneal away residual stresses.[2] Mode-confining bulges still appeared but with surprising fidelity.

As proof of principle, the group fabricated a series of five nearly identical resonators that are each separated, like beads on a string, by 100 μm of ordinary silica fiber.[3] Panel a of the figure illustrates the concept.

It's a SNAP

Sumetsky hit on the method last year after trying to determine the natural variations in the diameter of an optical fiber. University of Bath physicists Tim Birks, Jonathan Knight, and Tim Dimmick had proposed an experiment a decade earlier to resolve the issue:[4] Light sent through a thin, tapered microfiber perpendicular to the main fiber will evanescently couple to that fiber and excite modes at wavelengths that satisfy the resonance condition. The ra-

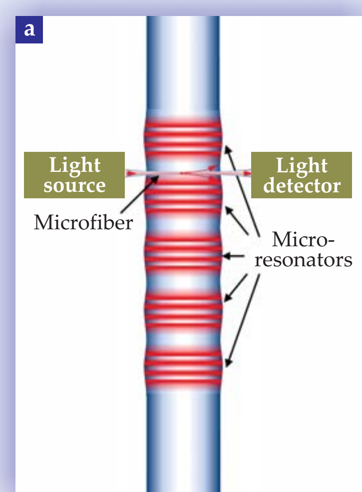

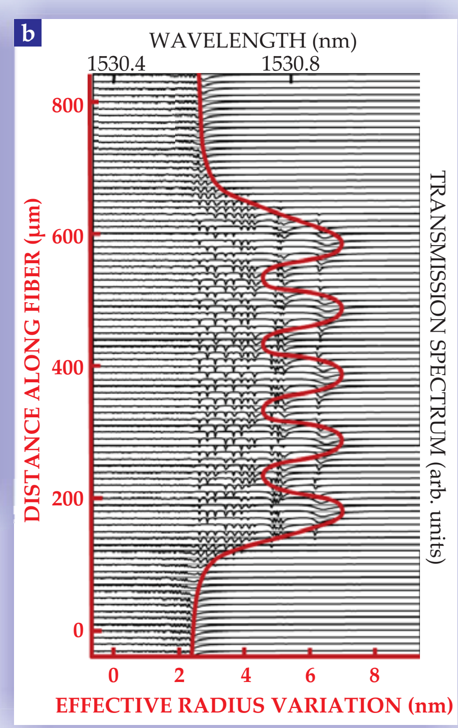

(a) A chain of five microresonators along a strip of glass fiber can be created by subtly modulating the fiber's radius and refractive index—together expressed as an "effective radius"—with a series of brief heatings from a laser. Light introduced inside the fiber via a narrower, perpendicular microfiber excites whispering gallery modes that are confined radially by total internal reflection and axially by a harmonic potential fixed by the curvature of each resonator. The modes' intensity maxima are shown in red. **(b)** The microfiber can also be used to measure the fiber's undulating profile. Transmission spectra taken every 10 μm along the fiber's axis reveal the resonant wavelengths (black) in each potential well. The bold envelope (red) surrounding the allowed modes maps the variation in effective radius for each resonator. The coupling between resonators is evident in the tunneling of modes through potential barriers. (Adapted from ref. 3.)

dius can thus be inferred from the microfiber's transmission spectrum, whose dips in intensity reveal the resonance wavelengths. Shifts in the radius are directly proportional to shifts in the wavelengths.

As Sumetsky and colleagues translated the microfiber along the main fiber's axis a few microns at a time, they noticed only exceedingly subtle shifts in the resonance wavelengths, an indication that even standard, off-the-shelf silica fiber could be uniform to within a few angstroms over millimeter lengths. Armed with that insight, they set out to fabricate extremely shallow bulges—a change in diameter of a few nanometers—to deliberately confine the resonant modes.

As hot glass is drawn under a known tension into fiber, it preserves that tension as it cools, an effect discovered at OFS Labs several years ago. The researchers found that they could incrementally release the frozen stress by locally heating the fiber for just five seconds to a few hundred degrees—well below silica's melting point. As the fiber got hotter with increases in laser power, the concomitant additional relaxation increased the fiber's "effective radius"—a parameter capturing the combined effect of changes in actual radius and refractive index. That the fiber relaxed in so reproducible a way for a given laser power was, Sumetsky admits, "just our luck and completely unexpected." The simple, quick, inexpensive, and reproducible nature of the technology prompted the researchers to coin the name "surface nanoscale axial photonics," or SNAP.

The OFS researchers exposed ordinary 36-μm-diameter fiber to five separate laser bursts, each generating a 2.5-nm modulation in effective radius. Altogether, the exposures produced a chain of quantum-well microresonators spanning close to a half-millimeter of fiber, as shown in panel b of the figure. Although the quantum wells are three-dimensional, the slow axial propagation of light through the chain follows a one-dimensional Schrödinger equation, whose potential is given by the shallow radius variation. The potential barriers between wells turned out to be low enough that the resonators coupled together, evidenced in the splitting of one band of modes and their tunneling from one resonator to the next.

To measure how well the radius variations—and thus the resonance wavelengths—of each resonator matched those of its neighbors, the researchers took transmission spectra every 10 μm along the microfiber axis. Birks's earlier prescription for equating a shift in radius with a shift in a resonance wavelength works ideally outside a microcavity. But inside, resonant modes oscillate back and forth axially as helices between two turning points. Accordingly, the resonance dips don't shift in wavelength with radius but remain nearly constant as the microfiber is scanned across each potential well. The widths of the resonances, however, narrow abruptly at the barrier regions, where coupling to the microfiber is weakest. So to gauge the tuning across different resonators, the researchers had to directly compare wavelengths of their narrowest resonances. The variation in those wavelengths showed that the radii were identical to within 2 Å.

Chalcogenides

Heat isn't the only process that can reproducibly vary the effective radius of silica fiber so subtly. Sumetsky and company found that control over the change in effective radius was even finer—to within 1 Å—when they exposed photosensitive germanium-doped silica to UV light.

And silica isn't the only glass fiber that can be formed into microresonators using laser light. Benjamin Eggleton and colleagues at the University of Sydney have demonstrated a similarly subtle localization of electromagnetic modes in a cavity they created from a half-millimeter strip of arsenic trisulfide fiber.[5]

As part of the chalcogenide class of amorphous semiconductors, As_2S_3 forms weak interatomic bonds—at least compared with those in oxides like silica—which makes it transparent into the mid-IR region of the spectrum. And like doped silica, the chemical bonding changes when the material is exposed to light with a wavelength near the band edge. That photosensitivity has led to applications in solar cells, IR sensors, and phase-change memories in compact disks.[6] The material is also highly polarizable, its refractive index a hundred times more nonlinear than that of doped silica.

Like Sumetsky's team, Eggleton and company exploited that photosensitivity using green laser light and monitored the resonance wavelengths of circulating whispering-gallery modes. As in silica, the change in the effective radius of As_2S_3 was dramatic and quick. But it was reversed in sign: The University of Sydney researchers measured a change of −0.24 nm/min, large enough that a single 12-minute light exposure created a mode-trapping microcavity.

(The negative sign meant they had to irradiate outside the developing cavity.) Moreover, because the material remained cool during the exposure, Eggleton was able to dynamically monitor changes in the transmission spectra.

How the effective radius of a chalcogenide fiber may change with heat and ion doping remains an interesting and open question.

Mark Wilson

References

1. M. Pöllinger et al., *Phys. Rev. Lett.* **103**, 053901 (2009).
2. M. Sumetsky et al., *Opt. Lett.* **36**, 4824 (2011).
3. M. Sumetsky et al., http://arxiv.org/abs/1112.5175v1.
4. T. A. Birks, J. C. Knight, T. E. Dimmick, *IEEE Photonics Technol. Lett.* **12**, 182 (2000).
5. F. Luan et al., *Opt. Lett.* **36**, 4761 (2011).
6. B. J. Eggleton et al., *Nat. Photonics* **5**, 141 (2011).

The Large Hadron Collider yields tantalizing hints of the Higgs boson

There's not yet enough data for a convincing sighting of standard particle theory's long-sought capstone. But there's less and less room for it to hide.

On 13 December, a standing-room-only audience at CERN and webcast viewers worldwide heard representatives of the two principal detector teams at the laboratory's Large Hadron Collider (LHC) report the year-end status of their searches for the Higgs boson. Interest in a possible first sighting of this last undiscovered particle required by particle theory's standard model (SM) is unusually widespread.

In the SM, the quantum field of the Higgs boson serves to break the underlying symmetry of the model's unification of the electromagnetic and weak interactions. It bestows the requisite large masses on the $W^{\pm}$ and Z^0 bosons that mediate the short-range weak interactions, while leaving their sibling, the photon, massless and free ranging. Coupling to the Higgs field also accounts for the nonvanishing masses of the fundamental fermions: the quarks and leptons.

The SM Higgs boson H has neither charge nor spin. (One speaks of an SM Higgs to distinguish it from the many elaborations beyond the SM that theorists have in their pockets, just in case.) The SM does not specify the Higgs mass M_H, but it does predict the particle's production cross section and decay rates as functions of M_H. Precision measurements of electroweak parameters and null results of searches at other colliders had pretty much restricted M_H to the range 114–156 GeV. (For comparison, the proton weighs about 1 GeV.) So that has been the mass region to which the teams running the gargantuan ATLAS and CMS detectors at the LHC have given particular attention in their searches, which began early in 2010.

By April 2011, the LHC was running well enough with colliding beams of 3.5-TeV protons for CERN director-general Rolf Heuer to predict a statistically robust determination of the existence, or nonexistence, of an SM Higgs by the end of 2012. But since April, the collider's performance has far surpassed expectations—providing each detector a total of 5×10^{14} proton–proton collisions by the end of October. That unanticipated abundance, though still not enough for a conclusive Higgs discovery, has yielded the tantalizing hints that were the highlights of the 13 December colloquium.

ATLAS and CMS spokespersons Fabiola Gianotti and Guido Tonelli, reporting their data on the expected Higgs decay modes,[1,2] showed "intriguing" mass peaks consistent with M_H near 125 GeV. But the peaks were also consistent with statistical fluctuations of 2.3 standard deviations for ATLAS and 1.9 standard deviations for CMS. Those quoted statistical significances do not, however, take into account the encouraging fact that both teams saw very similar results. "But even taken together, they can't be regarded as a discovery," cautioned Gianotti. "That," added Tonelli, "will require the LHC's 2012 run."

Looking for Higgs decays

The SM expectation is that 5×10^{14} pp collisions would produce almost 10^4 Higgs bosons in each detector. But with a lifetime of only 10^{-22} seconds, an H could be seen and weighed only through its decay products, identified against an overwhelming background of hadronic collision debris.

In the M_H region under principal scrutiny, the only decay modes that would have both adequate detection rates and mass resolution are $H \to \gamma\gamma$ and $H \to Z^0Z^0 \to 2(\ell^+\ell^-)$. The detectors are good at distinguishing high-energy gammas (γ) and charged leptons ($\ell^{\pm}$) from the much more abundant hadrons. Still, for the two "golden" modes combined, the excess over the expected non-Higgs background would be only a few dozen events in each detector over the entire LHC run so far.

Indeed both teams found that their observed M_H distributions for the $\gamma\gamma$ and ZZ modes showed peaks above background near 125 GeV. Each peak, with roughly the expected amplitude, is, by itself, statistically marginal. But they all fall within a few GeV of 125 GeV—a spread consistent with the measurement uncertainties.

The case for an SM Higgs sighting was further bolstered by a third decay mode: $H \to W^+W^- \to \ell^+\nu\,\ell^-\nu$. Its mass resolution is degraded by the two invisible neutrinos (ν), but both teams find that the WW data contribute significantly to the broad excess of events they found in the mass region 110–140 GeV. The combined ATLAS data for the three decay modes peak at 126 GeV, and the CMS peak for the combined modes is centered at 124 GeV.

As a function of M_H, the ATLAS team's summary figure on page 17 plots the 95%-confidence upper limit on the Higgs-production cross section σ deduced from the data and normalized to the SM predictions $\sigma_{SM}(M_H)$. Those upper limits serve to exclude an SM Higgs heavier than 131 GeV. Similarly, the CMS data exclude M_H above 128 GeV.

The figure's green and yellow swath shows the SM expectation (two standard deviations wide) for what upper limits the experimenters could extract from the data sample *if* there were no Higgs in that mass region. One gets an idea of the statistical significance of the peak at 126 GeV by noting how far it rises above the swath. More LHC running will move the swath downward, raising a true Higgs peak farther above what a two-standard-deviation statistical fluke could mimic.

"Though we can't yet claim a discov-