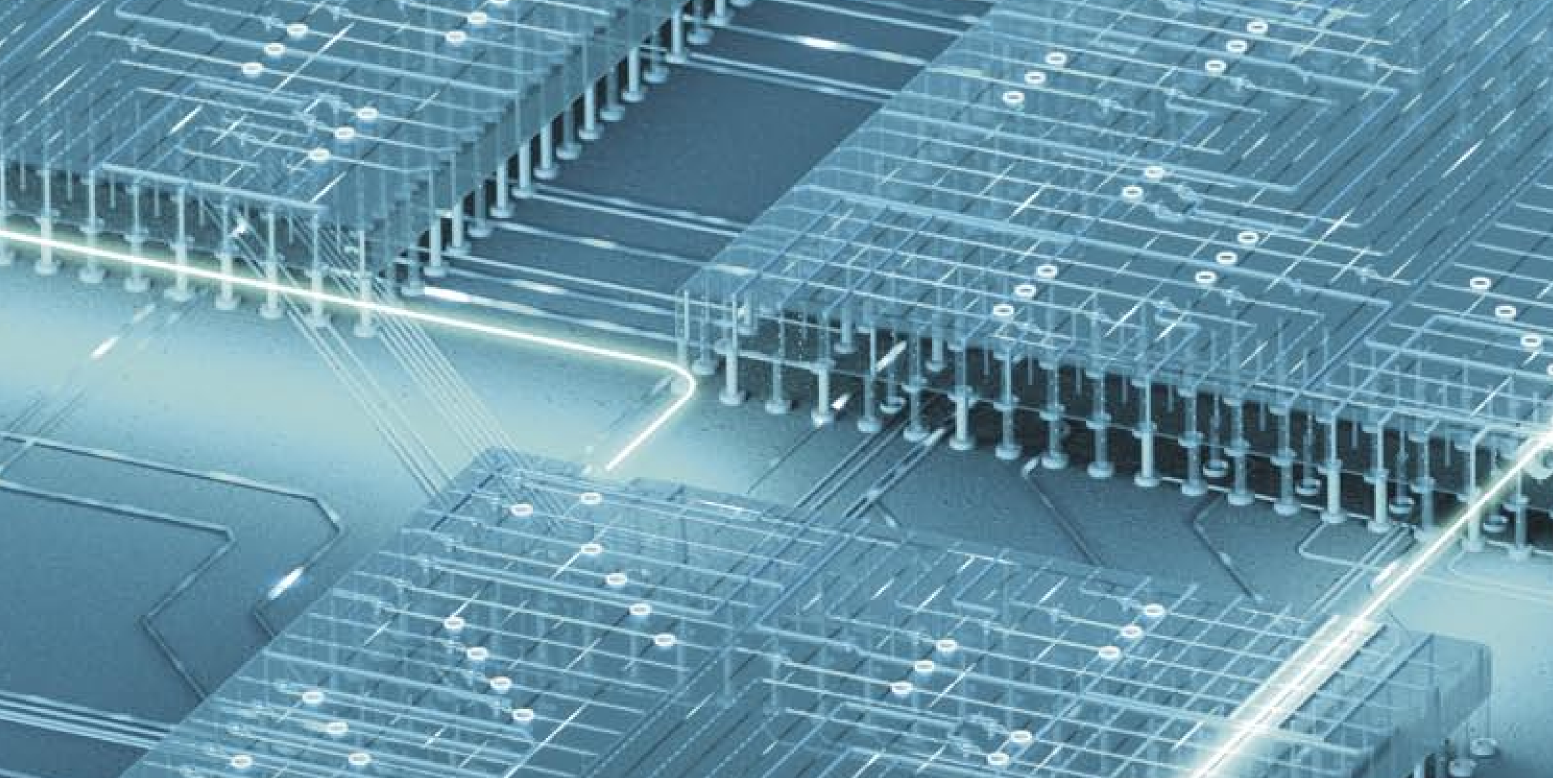

MAKING SIL

Scientists have at last persuaded silicon to emit laser beams. In a few

LOW-COST SILICON CHIPS ENABLE ENGINEERS to manipulate streams of electrons so that they can perform the myriad functions that make our computers, cellular phones and other consumer electronics so useful. If integrated silicon circuits could similarly create and control beams of light, they could make possible a range of inexpensive new technologies suited to many other applications. But for decades, silicon's very nature has thwarted scientists' dogged efforts to transform the material into a source of the necessary concentrated light.

Now several research groups, starting with mine, have coaxed silicon to produce laser light. The advance could have enormous implications for electronic devices that incorporate lasers and optical amplifiers, which currently rely on lasing materials that are far more costly and less common than silicon.

Replacing traditional copper interconnects and cables with optical conduits could raise data-trans-

U.C.L.A. (*laser glow*); BRYAN CHRISTIE DESIGN

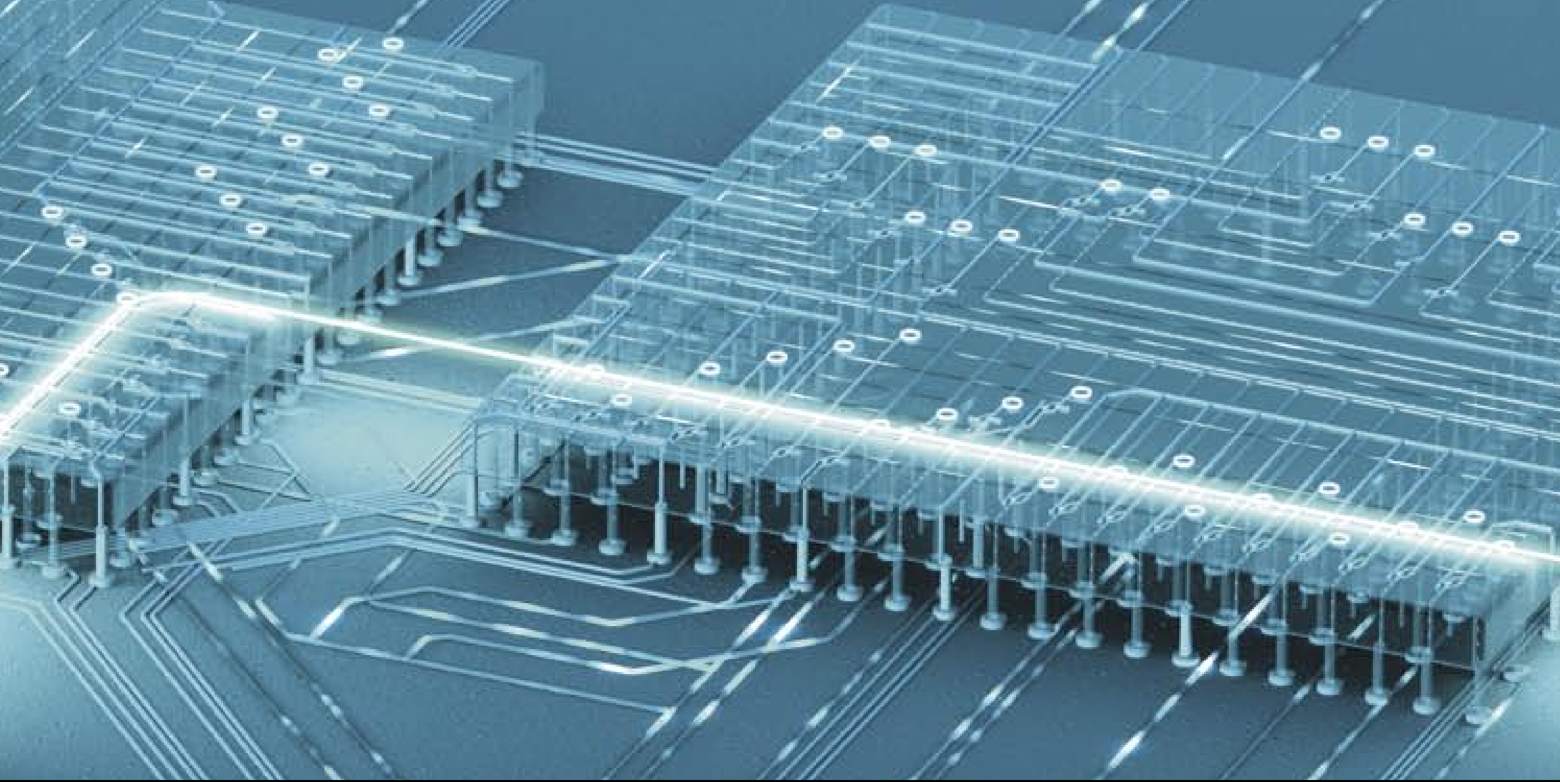

IC N LASE

years, computers and other devices will handle light as well as electrons

BY BAHRAM JALALI

fer speed limits by orders of magnitude beyond the capabilities of current technology. For example, cable modems, the workhorses of household Internet connections, are currently limited to a transfer rate of about one megabyte a second. Optical devices built on silicon chips could effortlessly transfer oversize digital files, such as high-definition videos, at rates as high as 10 gigabits a second—a 10,000-fold improvement. Compact sensors that include integrated circuits with silicon lasers could combine the capabilities of a diagnostic lab-on-a-chip with wireless communications to detect pollutants, chemical warfare agents or explosives as parts of extensive environmental monitoring and security networks. In a promising military applica-

LASER BEAM ***(red and white glow in title)*** **was the first ever produced by a silicon device; its infrared light is invisible to the eye and is false-colored here. Silicon lasers that are integrated into microchips** ***(background)*** **could make low-cost computing with light practical.**

THE CHALLENGE OF GETTING SILICON TO LASE

Silicon offers enormous promise for low-cost computing with light, but its very nature makes it an unlikely lasing medium.

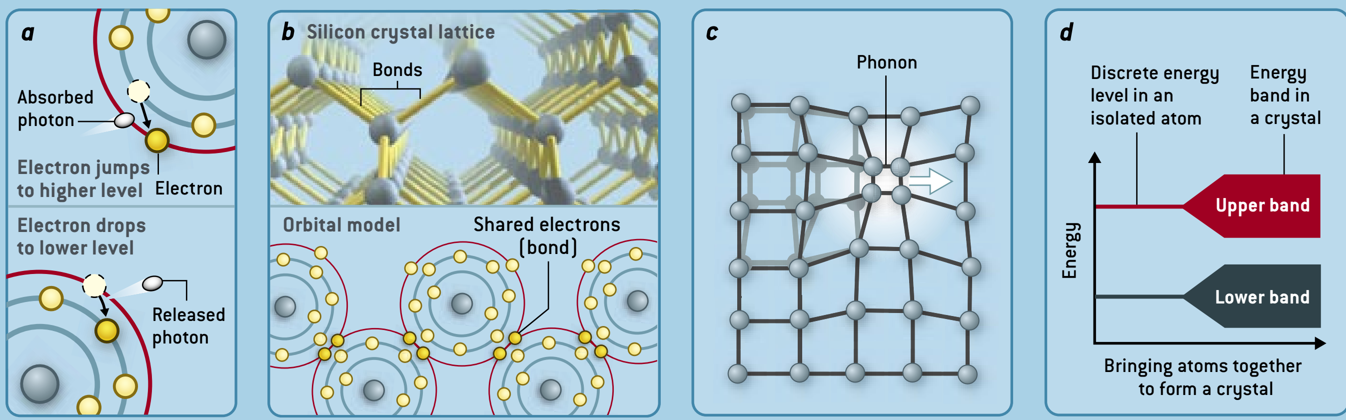

The lasing process is based on the quantum behavior of electrons in the outer orbitals of the atoms in a suitable material. An outer-shell electron in a single atom is energized (or "pumped") after absorbing a photon—an elementary quantum unit of light—which raises it to a higher orbit and energy level (*a*). An energized electron releases a photon when it drops down to a lower level.

In a solid, atoms form bonds by sharing these outer electrons (*b*). To achieve light amplification, the prerequisite to lasing, external energy sources pump the shared electrons to higher energy levels. As the energized electrons release photons, those photons, in turn, stimulate further photon emissions—amplifying light. Photons can also be amplified when they collide with excited phonons—which are quantized atomic vibrations of the crystal lattice (*c*).

When single atoms bond together into crystals, the character of the shared electrons' energy levels changes to broader bands (*d*) because of the effect on the electromagnetic environment of many nearby atoms. Thus, a pumped electron in a crystal jumps from one band to another.

When the energies and momenta of electrons in a common lasing medium such as gallium arsenide are plotted on a graph, the energy bands line up vertically because they share the same momenta (*e*). (A band describes the possible quantum states that electrons can take; each state has a quantity that can be identified with a classical momentum and that must be conserved during collisions.) The bands in silicon, in contrast,

tion, silicon lasers might be able to mislead the infrared sensors of heat-seeking antiaircraft missiles to provide a cost-effective countermeasure against them.

Laser Primer

WHY HAS IT TAKEN so long to teach silicon this new trick? Unlike materials commonly used as host media for lasing, such as the gallium arsenide used in DVD players, silicon is not naturally organized to promote the two-step process that results in a coherent light beam. It cannot emit light efficiently when energized (the first requirement), and whatever light it does produce is incapable of amplifying light into a laser beam by "stimulating" it to create more photons. ("Laser" stands for *l*ight *a*mplification by *s*timulated *e*mission of *r*adiation.)

In a laser, an external energy source, usually light or an electric current, "pumps" the electrons in the atoms of the host medium to a higher energy level, which physicists call an upper (or excited) state. When these atoms return to their normal (ground) state, the extra energy is released as photons of light (elementary quantum units of electromagnetic radiation that exist simultaneously as a wave and a particle). Albert Einstein dubbed this process "spontaneous emission": a phenomenon that yields photons that travel outward in all directions at random, giving rise to low-intensity, diffused light—much like that cast by a fluorescent bulb. When one of these emitted photons passes through a group of previously pumped electrons in the host material, it triggers, or stimulates, the electrons to discharge their surplus energy all at once, a concept first proposed in a paper published by Einstein in 1917. The resulting photons travel together in the same direction in synchrony, forming a highly directional light beam. As the beam travels through other excited atoms in the medium, its photons, in turn, stimulate the emission of yet more photons in a cascade. The effect

Overview/*Silicon Lasers*

- Scientists have long sought a silicon chip that can handle light as deftly as it does electrons, but silicon does not produce light easily, especially concentrated laser light. Such an advance could lead to ultrafast digital data transfers, new sensor networks and many other innovations.
- After many years of work, researchers are finally making silicon lase by using several different materials-based techniques. The birth of a novel hybrid technology—silicon electrophotonics—is at hand.

JEN CHRISTIANSEN (*this page and opposite page*); KEN EWARD *BioGrafx* (*b*)

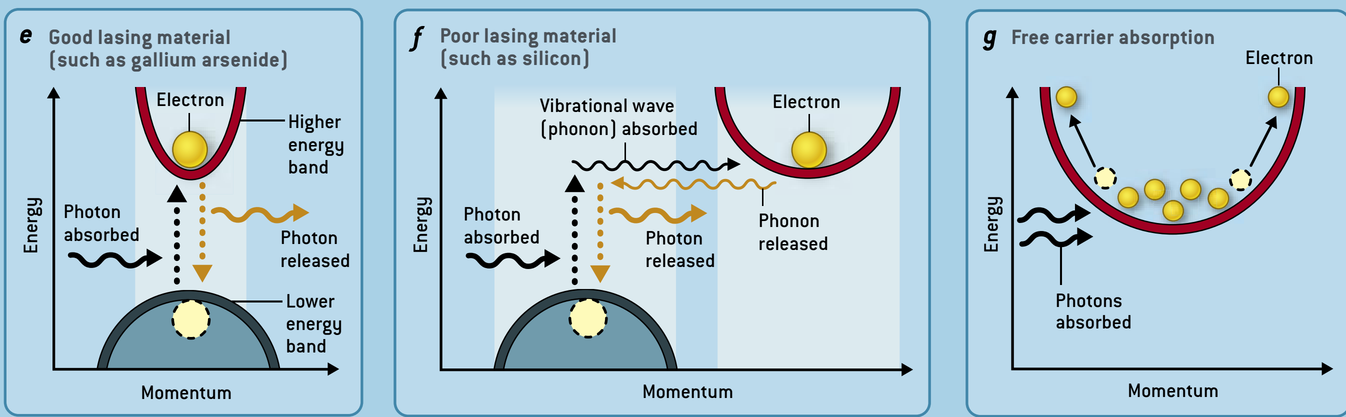

have different momenta, which means that energy from an absorbed photon alone is not enough to cause an electron to jump to a higher band (*f*). Instead the electron must wait until a phonon with the right extra momentum shows up to broker the transfer of energy. Unfortunately, these electrons often lose their surplus energy as heat before a suitable phonon arrives, which leads silicon to emit light inefficiently.

Silicon's low emission efficiency allows an intraband phenomenon called free carrier absorption to obstruct light amplification and lasing. When a passing photon interacts with an energized electron (a free carrier) in an upper band, one of two competing processes can occur. Either the photon can stimulate the emission of another photon, causing the electron to fall to a lower band, or the electron can simply absorb the photon, which merely moves it higher in the same band (*g*), an event that does not produce another photon and thus does not promote light amplification and lasing.

Gallium arsenide's upper bands hold comparatively few electrons. When depicted on a graph, the upper band is narrow, with steeply sloped sides. Because gallium arsenide has a high emission rate (it amplifies light efficiently because its bands line up), its photon emissions easily outpace its absorptions, and thus the material amplifies light. The wider, less-steep-sided upper bands in silicon require more electrons to fill them up. Silicon, with its low emission rate (caused by its indirect lineup) and high free carrier absorption rate, is unable to amplify light.

is analogous to the way the mass of an avalanche grows as it courses down a snow-covered mountainside.

Einstein's prediction of stimulated emission did not garner much interest until the 1950s, when physicists began to realize its potential applications in optical or photonic devices. In 1958 Charles Townes and Arthur Schawlow proposed partially surrounding a light-amplifying material with mirrors to reflect some of the photons it generated back inside. They showed that the stimulation process then would feed on itself (as in a chain reaction). This approach, once fully developed, could create a powerful light flow with a well-defined wavelength—a laser beam. Just two years later Theodore Maiman demonstrated the first operational laser, made by optically pumping a ruby crystal with a powerful lamp.

Silicon has proved considerably less pliant than ruby crystals or other subsequently developed lasing media. In semiconductors—materials whose electrical performance lies midway between an excellent conductor such as copper and an insulator such as rubber or certain ceramics—electrons exist in energy bands, which are ranges of energy levels, or states, that electrons can occupy.

The energy band describes the range of levels that electrons are "forbidden" or "allowed" to inhabit according to quantum theory; the gap between permitted bands (the band gap) is a range of energy levels that the electron is forbidden to occupy. An electron in an atom's outer orbital can take on energy (letting it jump to a higher band) by absorbing a photon or can release energy (dropping back down) by emitting one. Physicists classify these interactions as a kind of scattering event.

Imagine the energy bands as a series of buckets in which electrons sit [*see box above*]. Normally, almost all the electrons remain in the lower energy band, or bucket, leaving the upper band nearly empty. But if a photon with energy equal to or greater than the band gap collides with an electron, it can raise the electron to the upper band. The electron hops from the lower bucket to a higher one. Light absorption, the name for this effect, is the basis for the way solar cells convert light into electricity.

For the material to produce photons, it must receive enough energy to pump many electrons from the lower band to the upper band, causing a so-called population inversion (compared with the usual distributions present in the bands). It is unnecessary to pump the entire electron population; only the part near the top of the lower band need be affected. Engineers often excite electrons directly by forcing electric current through a semiconductor diode. Illuminating the substance with an external light source, as Maiman did, can also pump electrons.

AS PHOTONS TRAVEL THROUGH EXCITED ATOMS, THEY STIMULATE THE EMISSION OF YET MORE PHOTONS IN A CASCADE.

The electrons in the upper band eventually release energy, thereby emitting photons. When the resulting photons speed through a semiconductor that has many electrons in the upper level (an inverted electron population), they stimulate emissions of still other photons. In the best case, semiconductor emissions match the energy absorbed.

Although electrons and photons trade energy in these emission and absorption (scattering) processes, the system's total energy is conserved—that is, energy credits equal energy debits, as required by the law of conservation of energy. But absorption and emission will occur only if momentum is also conserved, according to the law of conservation of momentum. Momentum, which for a photon traveling (as a wave) in a crystal is determined directly from its wavelength, can be thought of as a tendency for a photon to continue to travel in the same direction. Being packets of pure energy, photons do not have much momentum to contribute in scattering collisions, so the transfers work best when the upper and lower bands (the starting and ending points of the interband transactions) have the same momentum. This equality of momentum occurs in commonly used lasing materials such as gallium arsenide and indium phosphide, whose energy bands lie directly on top of each other when plotted on a graph comparing energy and momentum. Such a direct lineup allows energy to be traded directly between an electron and a photon [*see box on preceding two pages*]. Whether a substance has this so-called direct lineup is intrinsic to the arrangement of atoms in the material's crystal lattice.

Silicon, however, naturally has an indirect lineup as a result of a less than optimal atomic crystal structure—that is, the material suffers from a large difference in the momentum between its upper and lower bands. (A band describes the possible quantum states that electrons can take; each state has a quantity that can be identified with a classical momentum and that must be conserved during collisions.) Thus, electrons cannot easily exchange energy with a photon and still conserve momentum. Instead they must wait until a vibrational wave of the silicon's crystal lattice, called a phonon, with just the right momentum shows up to provide the necessary extra momentum to broker the transfer of energy. Unfortunately, the electrons in silicon often lose their pump energy as heat while waiting for a suitable phonon to arrive. As a result, silicon exhibits a low emission efficiency; only about one excited electron in a million will successfully release a photon. Common lasing media such as gallium arsenide, in comparison, feature emission efficiencies some 10,000 times larger.

An indirect band gap limits the efficiency of a silicon laser, but it does not prohibit lasing by itself. Two other factors, also intrinsic to silicon, are involved. The first is free carrier absorption, a process that happens within a given energy band. Imagine a group of electrons (the free carriers) that have been pumped to an upper band. When a passing photon interacts with an energized electron, two events can occur—one favorable, the other unfavorable. The photon can cause the electron to drop down to a lower band by stimulating emission of another photon, which feeds the process of light amplification. Or the photon can be absorbed by the electron, which then merely moves higher in the upper band, a process that fails to generate another photon and thus does not result in amplification of light. The rates at which these two competing effects occur depend on the number of pumped electrons that reside in the upper energy band.

The upper bands (or buckets) in good lasing materials such as gallium arsenide are narrow and have steep sides, so they tend to hold relatively few electrons. In contrast, silicon features wider, less-steep-sided upper bands that require more electrons to fill up. When pumped, silicon has a large tendency to support free carrier absorption. Because gallium arsenide has a high emission rate (it amplifies light efficiently because its bands line up), its total photon emissions easily outpace its absorptions. Silicon, with its low emission rate (caused by its indirect lineup) and high free carrier absorption rate, is unable to amplify light.

An esoteric process known as Auger recombination also impedes silicon lasing. In this phenomenon, rather than emitting light, an electron in the upper band loses its energy to other electrons that subsequently give up the excess energy as heat. The amount of wasted light energy depends on the number of electrons present in the upper band. Silicon undergoes more Auger recombination than does gallium arsenide because it needs more electrons to be pumped into the upper band to overcome its low light-emission efficiency.

Teaching Silicon to Lase

IN THE PAST FIVE YEARS, researchers have begun to find ways around silicon's built-in hurdles. One method to enhance light emission takes advantage of an intriguing phenomenon called quantum confinement, which occurs when electron movement is constrained in one or more directions. In a three-dimensional restriction, called a quantum cage, an electron becomes agitated when the cage's size shrinks. This effect occurs as a result of the Heisenberg uncertainty principle, which states that localizing an electron makes its velocity, and hence its momentum (which equals mass times velocity), more random. This condition effectively relaxes the momentum con-

THE AUTHOR

BAHRAM JALALI is professor of electrical engineering at the Henry Samueli School of Engineering and Applied Science at the University of California, Los Angeles. He is also trustee of the California Science Center. Jalali spends much of his leisure time kayaking and sailing with his three children off the southern coast of California.

servation requirement governing electron-photon energy transfer, which boosts the rate of light emission by the semiconductor.

To make a quantum cage for silicon, researchers can create a thin film of silica (silicon dioxide) glass with tiny pieces of crystalline silicon embedded in it. These nanocrystals, which can be pumped by illuminating them with an external light source, are only a few atoms wide, so they can achieve quantum confinement. In 2000 Lorenzo Pavesi's group at the University of Trento in Italy first reported evidence of optically amplified silicon nanocrystals. The physics community initially greeted the result with skepticism, but Philip Fauchet of the University of Rochester and others subsequently confirmed the effect. Although this approach has yet to produce a laser, it inspired other innovations that have yielded encouraging results.

One advance that exploits quantum confinement makes use of rare earth elements, such as erbium, that scientists know to be good light emitters. Device makers routinely add erbium to the glass in optical fibers to create light-pumped amplifiers and lasers for telecommunications networks. Francesco Priolo of the University of Catania in Italy and Salvatore Coffa of Geneva-based STMicroelectronics have led the research into this method as a means for improving silicon's optical performance. Coffa's group has demonstrated light-emitting diodes (LEDs) that operate at room temperature with efficiencies as high as those of gallium arsenide devices.

The STMicroelectronics LED is a metal-glass-semiconductor sandwich in which a voltage that is maintained between the metal and the semiconductor accelerates electrons through the glass. As they move through, these electrons pump the electrons of erbium atoms in the glass, causing them to emit light. In this case, the quantum confinement in nanocrystals plays a relatively modest role, that of enhancing the conductivity of the glass so that the voltage required to establish the electron flow is reduced. Although LED technology is extremely useful, it produces diffused light (via spontaneous emission) rather than laser light generated by stimulated emission. The STMicroelectronics researchers do, however, anticipate demonstrating true lasing in erbium-doped silicon sometime soon.

Most recently, James Xu's group at Brown University observed lasing at low temperatures (–230 Celsius—too low for common use) in a piece of nanostructured silicon [*see illustration on page 65*]. They produced this effect by first forming an array of closely placed physical holes (110 nanometers apart) on the surface of a thin film of silicon and then pumping them optically. Xu and his team attribute the laser emissions they moni-

RUNNING A SILICON LASER

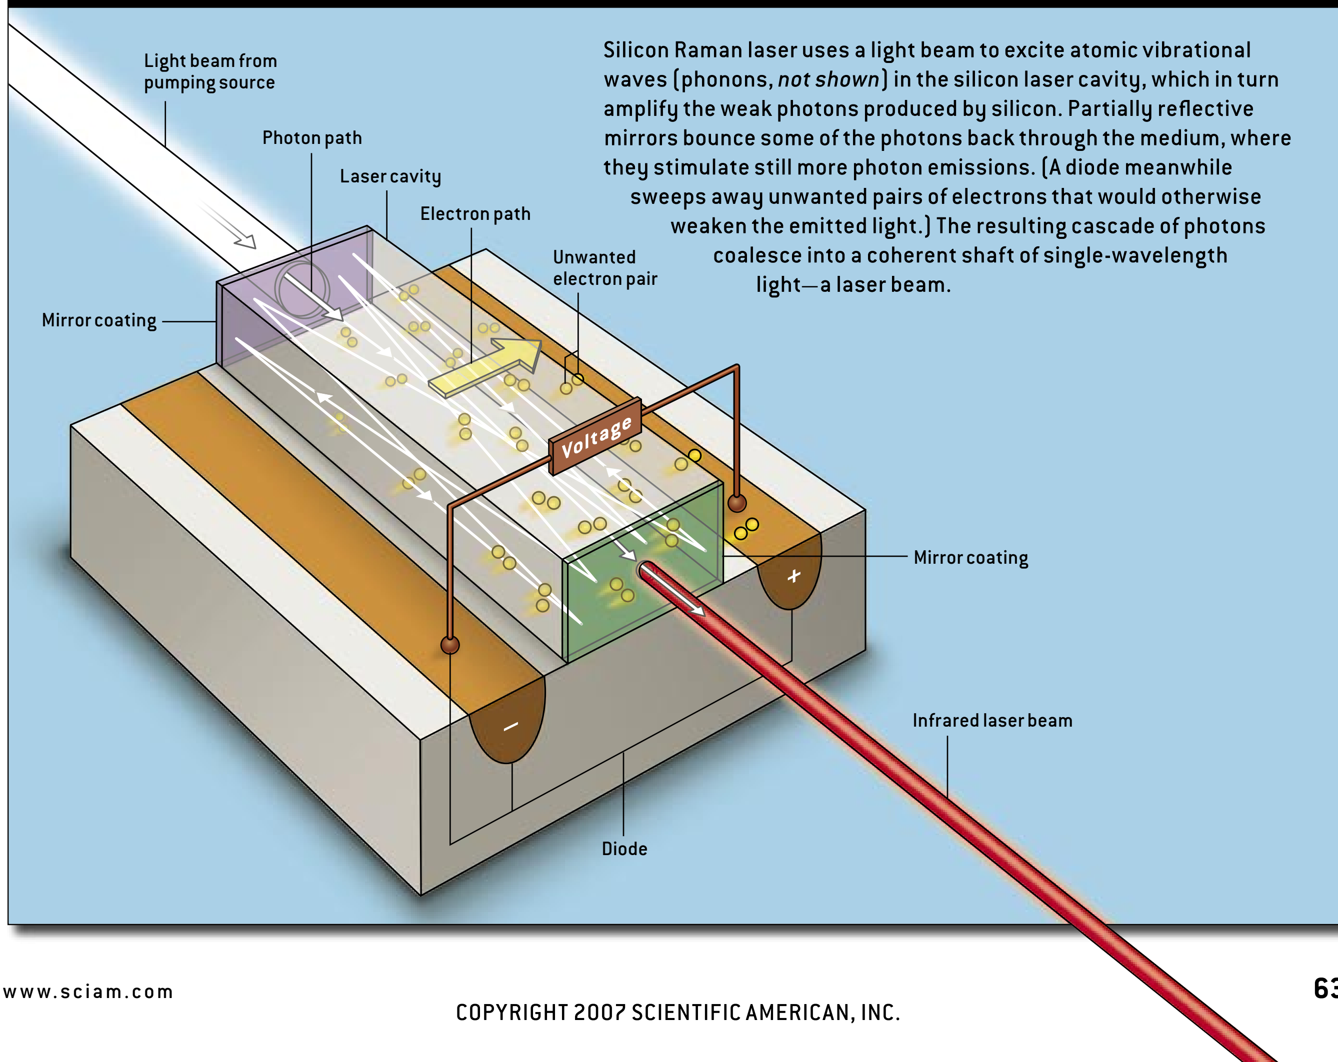

Silicon Raman laser uses a light beam to excite atomic vibrational waves (phonons, *not shown*) in the silicon laser cavity, which in turn amplify the weak photons produced by silicon. Partially reflective mirrors bounce some of the photons back through the medium, where they stimulate still more photon emissions. (A diode meanwhile sweeps away unwanted pairs of electrons that would otherwise weaken the emitted light.) The resulting cascade of photons coalesce into a coherent shaft of single-wavelength light—a laser beam.

JEN CHRISTIANSEN

SURPRISINGLY, OUR SILICON DEVICE CONVERTED PUMP ENERGY TO LIGHT NEARLY AS EFFICIENTLY AS CONVENTIONAL LASERS.

tored to electrons localized on lattice defects that occur naturally on crystal surfaces of the silicon nanostructures. They further ascribed the enhanced emissions to the quantum-based uncertainty in momentum produced by the tight localization of electrons. These structures create exciting possibilities for nanoscale silicon lasers, which exploit not just optical lasing in silicon but also the element's ability to function as sophisticated mirrors and filters that can manipulate the generated light. Such devices could be useful in future communications networks [see "Photonic Crystals: Semiconductors of Light," by Eli Yablonovitch; Scientific American, December 2001].

Silicon Learns to Lase

PUMPING ELECTRONS into the upper energy band of a semiconductor crystal is not the only way to amplify light. Researchers are also following other routes toward silicon lasers. For example, if one adds energy to the phonons in a crystalline semiconductor, a weak beam of light traveling through the lattice can pick up this energy and become amplified. Feeding some of the intensified light back into the crystal then produces a laser.

In 2002 and 2003, with support from the Defense Advanced Research Projects Agency, my group at the University of California, Los Angeles, showed that a silicon chip could generate and amplify light using this technique. In 2004 we demonstrated the first silicon laser. As with Maiman's apparatus, we pumped our device optically, which is usually a rather inefficient process. Surprisingly, however, our silicon device converted pump energy to light nearly as effectively as today's conventional lasers. Shortly afterward, we embedded the laser device in a diode and succeeded in switching it on and off electrically.

Scientists call the interaction of light with phonons the Raman effect. They employed it extensively in the late 1960s and the 1970s to probe the physical properties of many materials, including silicon. More recently, engineers have harnessed the effect to make optical fibers function as amplifiers and lasers. Because several kilometers of fiber are needed for this purpose, earlier researchers had failed to see it as a practical route to a silicon laser chip. Our team, however, realized that everyone overlooked that the Raman effect in silicon can be 10,000 times larger than in optical fibers, which are made of glass. This much bigger response follows from the well-ordered atomic structure in a silicon crystal (finally, an inherent property of silicon that assists its ability to lase). The random atomic arrangement in the amorphous glass of optical fibers keeps the Raman effect small.

A Raman laser requires optical pumping. To avoid generating electrons in the silicon's upper energy band that would preclude light emission (the free carrier absorption problem), our group excited the silicon using infrared light with a wavelength of 1,500 nanometers. This technique kept the photon energy less than the band-gap level—thus, it remained insufficient to elevate an electron into the upper band. Occasionally, however, two photons will pool their energies and manage to boost an electron into the upper band. Although these kinds of pumped electrons are relatively few, they sap the system of energy.

Raman-based lasers are not the only ones subject to this kind of energy loss. In 2006 Alexander Gaeta and Michal Lipson of Cornell University demonstrated a potentially useful device that amplifies light by mixing it with a more powerful light beam. This amplifier, and its yet to be demonstrated laser counter-

Silicon Laser Image Amplifier

In an optical fiber (or waveguide) with a cross section that is much larger than the wavelength of some incoming light, any pattern in the light goes in and out of focus as it travels down the light pipe as a result of constructive and destructive interference among the lightwaves reflecting off the waveguide's walls. This focusing effect combines with optical amplification to simultaneously focus and amplify an image as the light passes down the waveguide (*warmer colors toward the right*). Researchers at the University of California, Los Angeles, and Northrop Grumman are jointly developing a device in which the Raman effect (the interaction of photons and phonons) amplifies an optical image as it propagates through a thick silicon waveguide. This image amplifier should improve the sensitivity of laser-based remote sensing and imaging systems that scientists use for environmental monitoring.

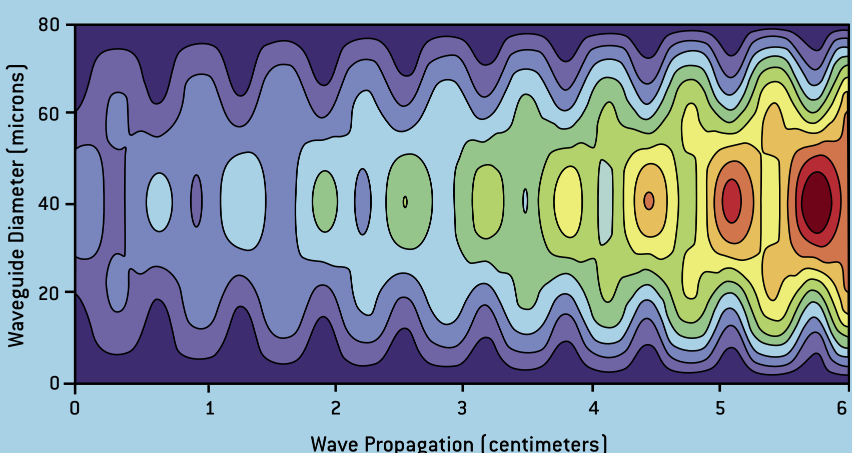

U.C.L.A./NORTHROP GRUMMAN

part, will experience the same losses as a Raman-based system.

To avoid such losses, our first laser operated in a pulsed mode that did not allow electrons to accumulate and drain the system of energy. For continuous laser operation, one can apply an electric field (created by an adjacent diode) to sweep lingering electrons away. Researchers at the Chinese University of Hong Kong suggested trying this approach, and Haisheng Rong and his coworkers at Intel demonstrated it in 2005. Recent investigations indicate that this method is only partially effective because the rate at which electrons can be removed is limited by the maximum velocity the particles can attain in silicon (one-thousandth the speed of light). It also requires significant electrical power to achieve. Fortunately, we know some tricks that can improve the efficiency of the silicon laser: bombarding the silicon with protons or adding small amounts of platinum tends to force the electrons to return rapidly to the lower energy band rather than soaking up scarce photons via free carrier absorption.

These procedures reduce the number of electrons in the upper band, which minimizes their reabsorption of light. Removing the electrons solves only part of the problem, though. The devices still lose pump energy when these electrons are generated unintentionally. By borrowing a trick that underpins the operation of solar cells, my team showed in 2006 that silicon Raman lasers can generate electrical power by harvesting the lost pump energy. Electrons that are generated by the unintentional two-photon absorption flow through the silicon to generate electricity. We learned that we could arrange the electron flow in such a way that the device's power consumption, given by the product of its electric current and its voltage, is negative, which means it actually generates power. The collected electrical power can then drive electronic circuits that reside on the same chip.

This difficulty vanishes altogether if one starts with an optical pumping wavelength that is longer than about 2,300 nanometers, as my research group later demonstrated. The resulting photon energy is so low that even a pair of photons does not possess enough energy to raise an electron into the upper band, which would be unhelpful in a Raman laser. We have found that silicon becomes an excellent lasing medium, arguably one of the best, when pumped with infrared wavelengths of 2,300 to about 7,000 nanometers (at which point other forms of deleterious effects start to appear). This spectrum lies beyond the reach of existing semiconductor lasers, so silicon laser technology permits the development of new applications. Among all laser materials, silicon offers one of the best combinations of thermal conductivity (to pass on unwanted heat) and resistance to damage from high levels of optical power, making it ideal for generating superintense laser beams.

Scientists have also developed a promising hybrid approach to producing a silicon-based laser that relies on adding a piece of gallium arsenide or indium phosphide to the top of a silicon substrate. The silicon research community has traditionally resisted hybrid techniques because the addition of other materials changes the electrical properties of silicon, so those materials are viewed as contaminants. Recent encouraging results, however, obtained by groups working at the University of Michigan at Ann Arbor and separately by a team of investigators at Intel and the University of California, Santa Barbara, have led to renewed interest in this approach. If research can overcome the problems of material incompatibility, this method may provide another near-term commercial path to a silicon-based laser.

The unrelenting pursuit of the silicon laser has finally begun to pay off. At last, the field seems to have reached the critical mass that will allow silicon to challenge traditional laser materials. This progress should make the convergence of electronics and photonics all but inevitable. Although it is too early to know the precise trajectory that this new electrophotonic technology will take, the new applications that will be made possible by silicon lasers are likely to have a dramatic impact on our everyday lives. SA

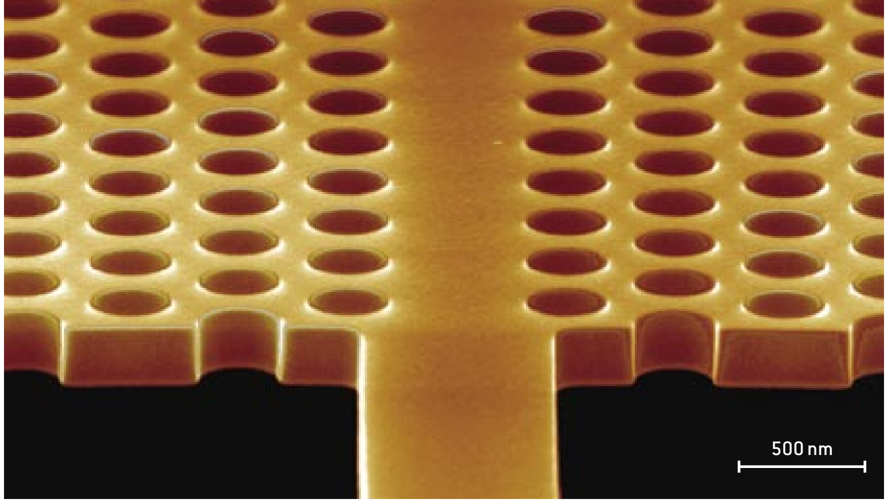

LASING AT LOW TEMPERATURE was demonstrated by James Xu's research group at Brown University in a thin film of silicon similar to that shown above. The surface of the team's device features nanoscale holes placed only 110 nanometers apart. Lasing occurs because electrons are quantum-confined in electron cages at the silicon surfaces.

IBM RESEARCH

MORE TO EXPLORE

Subtle Is the Lord: The Science and the Life of Albert Einstein. Abraham Pais. Oxford University Press, 1982.

The Incredible Shrinking Transistor. Yuan Taur in *IEEE Spectrum,* Vol. 36, No. 7, pages 25–29; July 1999.

Demonstration of a Silicon Raman Laser. Ozdal Boyraz and Bahram Jalali in *Optics Express,* Vol. 12, No. 21, pages 5269–5273; October 2004.

Computing at the Speed of Light. W. Wayt Gibbs in *Scientific American,* Vol. 291, No. 5, pages 80–87; November 2004.

Optical Gain and Stimulated Emission in Periodic Nanopatterned Crystalline Silicon. S. G. Cloutier, P. A. Kossyrev and J. Xu in *Nature Materials,* Vol. 4, No. 12, pages 887–891; December 2005.

Optical Interconnects: The Silicon Approach. Edited by L. Pavesi and G. Guillot. Springer Series in Optical Sciences. Springer, 2006.

Silicon Photonics: The State of the Art. Graham Reed. John Wiley & Sons, 2007.