THE LIGHT THAT SHIN

by CHARLES H. TOWNES

A Nobel laureate recounts the invention of the laser and the birth of quantum electronics.

ON JULY 21, 1969, astronauts Neil Armstrong and Edwin "Buzz" Aldrin set up an array of small reflectors on the moon, facing them toward Earth. At the same time, two teams of astrophysicists, one at the University of California's Lick Observatory and the other at the University of Texas's McDonald Observatory, were preparing small instruments on two big telescopes. Ten days later, the Lick team pointed its telescope at the precise location of the reflectors on the moon and sent a small pulse of power into the hardware they had added to it. A few days after that, the McDonald team went through the same steps. In the heart of each telescope, a narrow beam of extraordinarily pure red light emerged from a synthetic ruby crystal, pierced the sky, and entered the near vacuum of space. The two rays were still only a thousand yards wide after traveling 240,000 miles to illuminate the moon-based reflectors. Slightly more than a second after each light beam hit its target, the crews in California and Texas detected its faint reflection. The brief time interval between launch and detection of these light pulses permitted calculation of the distance to the moon to within an inch—a measurement of unprecedented precision.

The ruby crystal for each light source was the heart of a laser (an acronym for light amplification by stimulated emission of radiation), which is a device first demonstrated in 1960, just nine years earlier. A laser beam reflected from the moon to measure its distance is only one dramatic illustration of the spectacular quality of laser light. There are many other more mundane, practical uses such as in surveying land and grading roads, as well as myriad everyday uses—

Will & Deni McIntyre, P4219F Photo Researchers

Lasers are commonly used in eye surgery to correct defective vision.

ranging from compact disc players to grocery-store checkout devices.

The smallest lasers are so tiny one cannot see them without a microscope. Thousands can be built on semiconductor chips. The bigge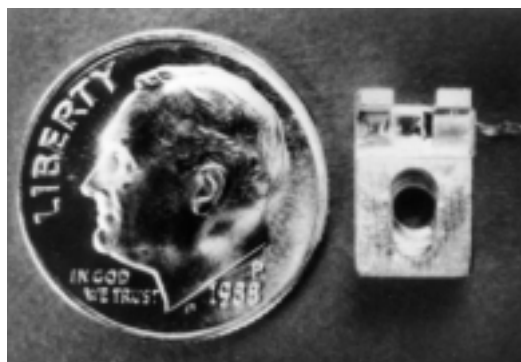st lasers consume as much electricity as a small town. About 45 miles from my office in Berkeley is the Lawrence Livermore National Laboratory, which has some of the world's most powerful lasers. One set of them, collectively called NOVA, produces ten individual laser beams that converge to a spot the size of a pinhead and generate temperatures of many millions of degrees. Such intensely concentrated energies are essential for experiments that can show physicists how to create conditions for nuclear fusion. The Livermore team hopes thereby to find a way to generate electricity efficiently, with little pollution or radioactive wastes.

For several years after the laser's invention, colleagues liked to tease me about it, saying, "That's a great idea, but it's a solution looking for a problem." The truth is, none of us who worked on the first lasers ever imagined how many uses there might eventually be. But that was not our motivation. In fact, many of today's practical technologies have resulted from basic scientific research done years to decades before. The people involved, motivated mainly by curiosity, often have little idea as to where their research will eventually lead.

An early small laser made of semiconducting material (right) compared with a U.S. dime.

LIKE THE TRANSISTOR, the laser and its progenitor the maser (an acronym for microwave amplification by stimulated emission of radiation) resulted from the application of quantum mechanics to electronics after World War II. Together with other advances, they helped spawn a new discipline in applied physics known since the late 1950s as "quantum electronics."

Since early in the twentieth century, physicists such as Niels Bohr, Louis de Broglie, and Albert Einstein learned how molecules and atoms absorb and emit light—or any other electromagnetic radiation—on the basis of the newly discovered laws of quantum mechanics. When atoms or molecules absorb light, one might say that parts of them wiggle back and forth or twirl with new, added energy. Quantum mechanics requires that they store energy in very specific ways, with precise, discrete levels of energy. An atom or a molecule can exist in either its ground (lowest) energy state or any of a set of higher (quantized) levels, but not at energies between those levels. Therefore they only absorb light of certain wavelengths, and no others, because the wavelength of light determines the energy of its individual photons (see box on right). As atoms or molecules drop from higher back to lower energy levels, they emit photons of the same energies or wavelengths as those they can absorb. This process is usually spontaneous, and this kind of light is normally emitted when these atoms or molecules glow, as in a fluorescent light bulb or neon lamp, radiating in nearly all directions.

Einstein was the first to recognize clearly, from basic thermodynamic principles, that if photons can be absorbed by atoms and boost them to higher energy states, then light can also prod an atom to give up its energy and drop down to a lower level. One photon hits the atom, and two come out. When this happens, the emitted photon takes off in precisely the same direction as the light that stimulated the energy loss, with the two waves exactly in step (or in the same "phase"). This "stimulated emission" results in coherent amplification, or amplification of a wave at exactly the same frequency and phase.
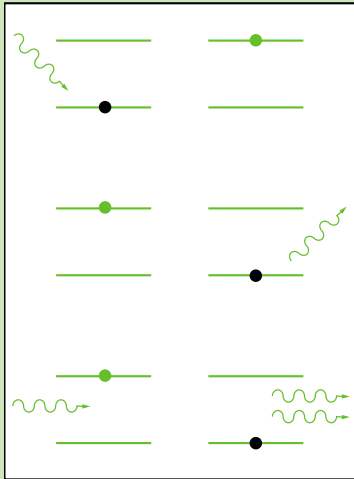
Both absorption and stimulated emission can occur simultaneously. As a light wave comes along, it can thus excite some atoms that are in lower energy states into higher states and, at the same time, induce some of those in higher states to fall back down to lower states. If there are more atoms in the upper states than in the lower states, more light is emitted than absorbed. In short, the light gets stronger. It comes out brighter than it went in.

The reason why light is usually absorbed in materials is that substances almost always have more atoms and molecules sitting in lower energy states than in higher ones: more photons are absorbed than emitted. Thus we do not expect to shine a light through a piece of glass and see it come out the other side brighter than it went in. Yet this is precisely what happens with lasers.

The trick in making a laser is to produce a material in which the energies of the atoms or molecules present have been put in a very abnormal condition, with more of them in excited states than in the

How Lasers Work

BECAUSE OF QUANTUM mechanics, atoms and molecules can exist only in discrete states with very specific values of their total energy. They change from one state to another by absorbing or emitting photons whose energies correspond to the difference between two such energy levels. This process, which generally occurs by an electron jumping between two adjacent quantum states, is illustrated in the accompanying drawings.

Stimulated emission of photons, the basis of laser operation, differs from the usual absorption and spontaneous emission. When an atom or molecule in the "ground" state absorbs a photon, it is raised to a higher energy state (top). This excited state may then radiate spontaneously, emitting a photon of the same energy and reverting back to the ground state (middle). But an excited atom or molecule can also be stimulated to emit a photon when struck by an approaching photon (bottom). In this case, there is now a second photon in addition to the stimulating photon; it has precisely the same wavelength and travels exactly in phase with the first.

Lasers involve many atoms or molecules acting in concert. The set of drawings (right) illustrates laser action in an optical-quality crystal, producing a cascade of photons emitted in one direction.
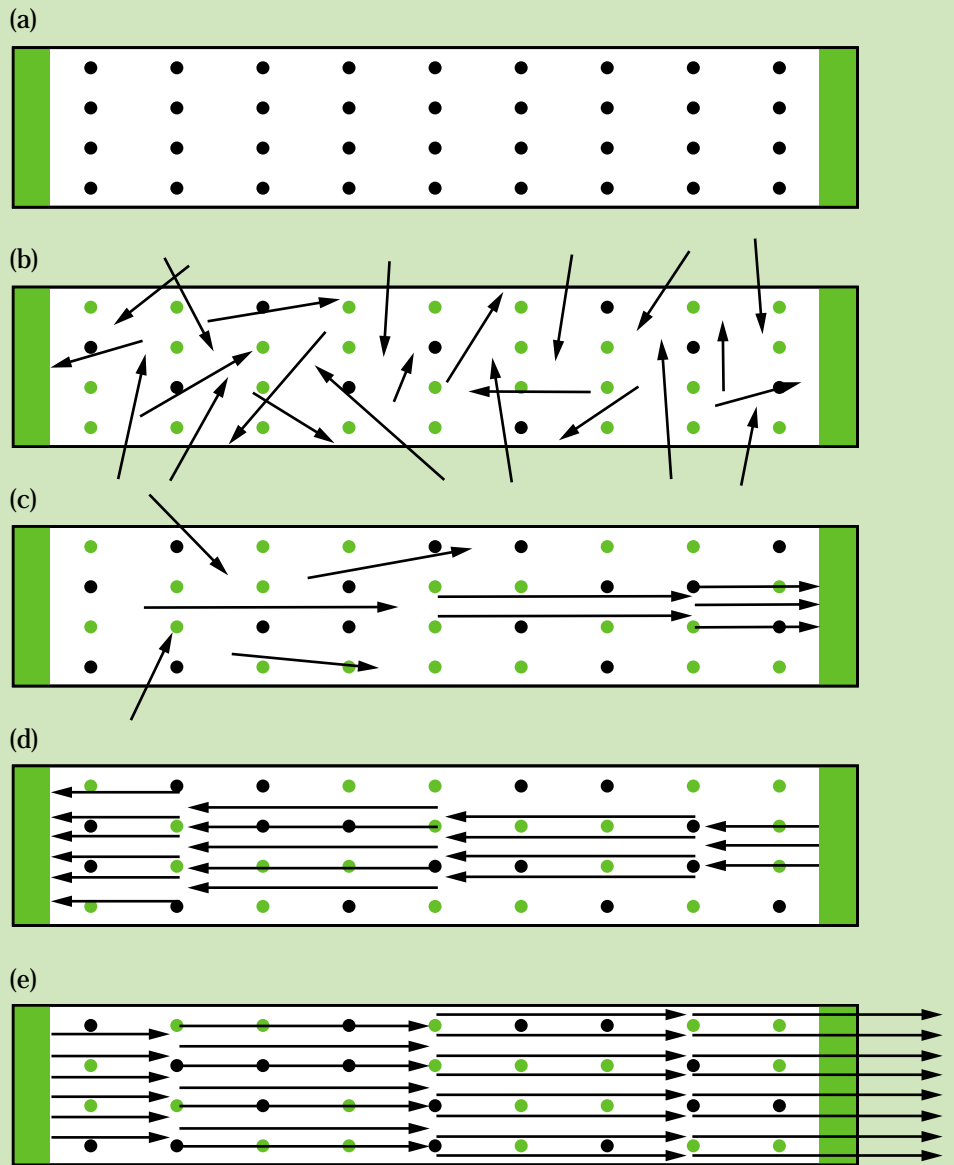
(a) Before the cascade begins, the atoms in the crystal are in their ground state. (b) Light pumped in and absorbed by these atoms raises most of them to the excited state. (c) Although some of the spontaneously emitted photons pass out of the crystal, the cascade begins when an excited atom emits a photon parallel to the axis of the crystal. This photon stimulates another atom to contribute a second photon, and (d) the process continues as the cascading photons are reflected back and forth between the parallel ends of the crystal. (e) Because the right-hand end is only partially reflecting, the beam eventually passes out this end when its intensity becomes great enough. This beam can be very powerful.

ground, or lower, states. A wave of electromagnetic energy of the proper frequency traveling through such a peculiar substance will pick up rather than lose energy. The increase in the number of photons associated with this wave represents amplification—light amplification by stimulated emission of radiation.

If this amplification is not very large on a single pass of the wave through the material, there are ways to beef it up. For example, two parallel mirrors—between which the light is reflected back and forth, with excited molecules (or atoms) in the middle—can build up the wave. Getting the highly directional laser beam out of the device is just a matter of one mirror being made partially transparent, so that when the internally reflecting beam gets strong enough, a substantial amount of its power shoots right on through one end of the device.

The way this manipulation of physical laws came about, with the many false starts and blind alleys on the way to its realization, is the subject of my book, *How the Laser Happened*. Briefly summarized in this article, it also describes my odyssey as a scientist and its unpredictable and perhaps natural path to the maser and laser. This story is interwoven with the way the field of quantum electronics grew, rapidly and strikingly, owing to a variety of important contributors, their cooperation and competitiveness.

DURING THE LATE 1940s and early 1950s, I was examining molecules with microwaves at Columbia University—doing microwave spectroscopy. I tried to find a way to generate waves shorter than those produced by the klystrons and magnetrons developed for radar in World War II. One day I suddenly had an idea—use molecules and the stimulated emission from them. With graduate student Jim Gordon and postdoc Herb Zeiger, I decided to experiment first with ammonia (NH_3) molecules, which emit radiation at a wavelength of about 1 centimeter. After a couple of years of effort, the idea worked. We christened this device the "maser." It proved so interesting that for a while I put off my original goal of trying to generate even shorter wavelengths.

In 1954, shortly after Gordon and I built our second maser and showed that the frequency of its microwave radiation was indeed remarkably pure, I visited Denmark and saw Niels Bohr. As we were walking along the street together, he asked me what I was doing. I described the maser and its amazing performance.

"But that is not possible!" he exclaimed. I assured him it was.

Similarly, at a cocktail party in Princeton, New Jersey, the Hungarian mathematician John von Neumann asked what I was working on. I told him about the maser and the purity of its frequency.

"That can't be right!" he declared. But it was, I replied, telling him it had already been demonstrated.

Such protests were not offhand opinions about obscure aspects of physics; they came from the marrow of these men's bones. Their objections were founded on principle—the Heisenberg uncertainty principle. A central tenet of quantum mechanics, this principle is among the core achievements that occurred during

a phenomenal burst of creativity in the first few decades of the twentieth century. As its name implies, it describes the impossibility of achieving absolute knowledge of all the aspects of a system's condition. There is a price to be paid in attempting to measure or define one aspect of a specific particle to very great exactness. One must surrender knowledge of, or control over, some other feature.

A corollary of this principle, on which the maser's doubters stumbled, is that one cannot measure an object's frequency (or energy) to great accuracy in an arbitrarily short time interval. Measurements made over a finite period of time automatically impose uncertainty on the observed frequency.

To many physicists steeped in the uncertainty principle, the maser's performance, at first blush, made no sense at all. Molecules race through its cavity, spending so little time there—about one ten-thousandth of a second—that it seemed impossible for the frequency of the output microwave radiation to be so narrowly defined. Yet that is exactly what happens in the maser.

There is good reason that the uncertainty principle does not apply so simply here. The maser (and by analogy, the laser) does not tell you anything about the energy or frequency of any specific molecule. When a molecule (or atom) is stimulated to radiate, it must produce exactly the same frequency as the stimulating radiation. In addition, this radiation represents the average of a large number of molecules (or atoms) working together. Each individual molecule remains anonymous, not

accurately measured or tracked. The precision arises from factors that mollify the apparent demands of the uncertainty principle.

I am not sure that I ever did convince Bohr. On that sidewalk in Denmark, he told me emphatically that if molecules zip through the maser so quickly, their emission lines must be broad. After I persisted, however, he relented.

"Oh, well, yes," he said; "Maybe you are right." But my impression was that he was just trying to be polite to a younger physicist.

After our first chat at that Princeton party, von Neumann wandered off and had another drink. In about fifteen minutes, he was back.

"Yes, you're right," he snapped. Clearly, he had seen the point.

NOVA, the world's largest and most powerful laser, at the Lawrence Livermore National Laboratory in California. Its 10 laser beams can deliver 15 trillion watts of light in a pulse lasting 3 billionths of a second. With a modified single beam, it has produced 1,250 trillion watts for half a trillionth of a second. (Courtesy Lawrence Livermore National Laboratory)

He seemed very interested, and he asked me about the possibility of doing something similar at shorter wavelengths using semiconductors. Only later did I learn from his posthumous papers, in a September 1953 letter he had written to Edward Teller, that he had already proposed producing a cascade of stimulated infrared radiation in semiconductors by exciting electrons with intense neutron bombardment. His idea was almost a laser, but he had not thought of employing a reflecting cavity nor of using the coherent properties of stimulated radiation.

I N THE LATE SUMMER of 1957, I felt it was high time I moved on to the original goal that fostered the maser idea: oscillators that work at wavelengths appreciably shorter than 1 millimeter, beyond what standard electronics could achieve. For some time I had been thinking off and on about this goal, hoping that a great idea would pop into my head. But since nothing had spontaneously occurred to me, I decided I had to take time for concentrated thought.

A major problem to overcome was that the rate of energy radiation from a molecule increases as the fourth power of the frequency. Thus to keep molecules or atoms excited in a regime to amplify at a wavelength of, say, 0.1 millimeter instead of 1 centimeter requires an increase in pumping power by many orders of magnitude. Another problem was that for gas molecules or atoms, Doppler effects increasingly broaden the emission spectrum as the wavelength gets smaller. That means there is less amplification per molecule available to

The development of the laser followed no script except to hew to the nature of scientists groping to understand, to explore, and to create.

drive any specific resonant frequency inside a cavity.

As I played with the variety of possible molecular and atomic transitions, and the methods of exciting them, what is well-known today suddenly became clear to me: it is just as easy, and probably easier, to go right on down to really short wavelengths—to the near-infrared or even optical regions—as to go down one smaller step at a time. This was a revelation, like stepping through a door into a room I did not suspect existed.

The Doppler effect does indeed increasingly smear out the frequency of response of an atom (or molecule) as one goes to shorter wavelengths, but there is a compensating factor that comes into play. The number of atoms required to amplify a wave by a given amount does not increase, because atoms give up their quanta more readily at higher frequencies. And while the power needed to keep a certain number of atoms in excited states increases with the frequency,

the total power required—even to amplify visible light—is not necessarily prohibitive.

Not only were there no clear penalties in such a leapfrog to very short wavelengths or high frequencies, this approach offered big advantages. In the near-infrared and visible regions, we already had plenty of experience and equipment. By contrast, wavelengths near 0.1 mm and techniques to handle them were relatively unknown. It was time to take a big step.

Still, there remained another major concern: the resonant cavity. To contain enough atoms or molecules, the cavity would have to be much longer than the wavelength of the radiation—probably thousands of times longer. This meant, I feared, that no cavity could be very selective for one and only one frequency. The great size of the cavity, compared to a single wavelength, meant that many closely spaced but slightly different wavelengths would resonate.

By great good fortune, I got help and another good idea before I proceeded any further. I had recently begun a consulting job at Bell Labs, where my brother-in-law and former postdoc Arthur Schawlow was working. I naturally told him that I had been thinking about such optical masers, and he was very interested because he had also been thinking in that very direction.

We talked over the cavity problem, and Art came up with the solution. I had been considering a rather well-enclosed cavity, with mirrors on the ends and holes in the sides only large enough to provide a way to pump energy into the gas and to kill some of the directions in

which the waves might bounce. He suggested instead that we use just two plates, two simple mirrors, and leave off the sides altogether. Such arrangements of parallel mirrors were already used at the time in optics; they are called Fabry–Perot interferometers.

Art recognized that without the sides, many oscillating modes that depend on internal reflections would eliminate themselves. Any wave hitting either end mirror at an angle would eventually walk itself out of the cavity—and so not build up energy. The only modes that could survive and oscillate, then, would be waves reflected exactly straight back and forth between the two mirrors.

More detailed studies showed that the size of the mirrors and the distance between them could even be chosen so that only one frequency would be likely to oscillate. To be sure, any wavelength that fit an exact number of times between the mirrors could resonate in such a cavity, just as a piano string produces not just one pure frequency but also many higher harmonics. In a well-designed system, however, only one frequency will fall squarely at the transition energy of the medium within it. Mathematically and physically, it was "neat."

Art and I agreed to write a paper jointly on optical masers. It seemed clear that we could actually build one, but it would take time. We spent nine months working on and off in our spare time to write the paper. We needed to clear up the engineering and some specific details, such as what material to use, and to clarify the theory. By August 1958 the manuscript was complete, and the Bell

Labs patent lawyers told us that they had done their job protecting its ideas. *The Physical Review* published it in the December 15, 1958, issue.

In September 1959 the first scientific meeting on quantum electronics and resonance phenomena occurred at the Shawanga Lodge in New York's Catskill Mountains. In retrospect, it represented the formal birth of the maser and its related physics as a distinct subdiscipline. At that meeting Schawlow delivered a general talk on optical masers.

Listening with interest was Theodore Maiman from the Hughes Research Laboratories in Culver City, California. He had been working with ruby masers and says he was already thinking about using ruby as the medium for a laser. Ted listened closely to the rundown on the possibility of a solid-state ruby laser, about which Art was not too hopeful. "It may well be that more suitable solid materials can be found," he noted, "and we are looking for them."

James Gordon, Nikolai Basov, Herbert Zeiger, Alexander Prokhorov, and author Charles Townes (left to right) attending the first international conference on quantum electronics in 1959.

Maiman felt that Schawlow was far too pessimstic and left the meeting intent on building a solid-state ruby laser. In subsequent months Ted made striking measurements on ruby, showing that its lowest energy level could be partially emptied by excitation with intense light. He then pushed on toward still brighter excitation sources. On May 16, 1960, he fired up a flash lamp wrapped around a ruby crystal about 1.5 cm long and produced the first operating laser.

The evidence that it worked was somewhat indirect. The Hughes group did not set it up to shine a spot of light on the wall. No such flash of a beam had been seen, which left room for doubt about just what they had obtained. But the instruments had shown a substantial change in the fluorescence spectrum of the ruby; it became much narrower, a clear sign that emission had been stimulated, and the radiation peaked at just the wavelengths that were most favored for such action. A short while later, both the Hughes researchers and Schawlow at Bell Labs independently demonstrated powerful flashes of directed light that made spots on the wall—clear intuitive proof that a laser is indeed working.

I**T IS NOTEWORTHY** that almost all lasers were first built in industrial labs. Part of the reason for industry's success is that once its importance becomes apparent, industrial laboratories can devote more resources and concentrated effort to a problem than can a university. When the goal is clear, industry can be effective. But the first lasers were invented and built by young scientists recently hired after their university research in the field of microwave and radio spectroscopy—students of Willis Lamb, Polykarp Kusch, and myself, who worked together in the Columbia Radiation Laboratory, and of Nicholas Bloembergen at Harvard. The whole field of quantum electronics grew out of this approach to physics.

The development of the maser and laser followed no script except to hew to the nature of scientists groping to understand, explore and create. As a striking example of how important technology applied to human interests can grow out of basic university research, the laser's development fits a pattern that could not have been predicted in advance.

What research planner, wanting a more intense light source, would have started by studying molecules with microwaves? What industrialist, looking for new cutting and welding devices, or what doctor, wanting a new surgical tool as the laser has become, would have urged the study of microwave spectroscopy? The whole field of quantum electronics is truly a textbook example of broadly applicable technology growing unexpectedly out of basic research.