From Hecht, E., "Optics" 4th Edition Addison Wesley 2002

13 Modern Optics: Lasers and Other Topics

13.1 Lasers and Laserlight

During the early 1950s a remarkable device known as the *maser* came into being through the efforts of a number of scientists. Principal among these people were Charles Hard Townes of the United States and Alexandr Mikhailovich Prokhorov and Nikolai Gennadievich Basov of the USSR, all of whom shared the 1964 Nobel Prize in Physics for their work. The maser, which is an acronym for Microwave Amplification by Stimulated Emission of Radiation, is, as the name implies, an extremely low-noise, microwave amplifier.* It functioned in what was then a rather unconventional way, making direct use of the quantum-mechanical interaction of matter and radiant energy. Almost immediately after its inception, speculation arose as to whether or not the same technique could be extended into the optical region of the spectrum. In 1958 Townes and Arthur L. Schawlow prophetically set forth the general physical conditions that would have to be met in order to achieve Light Amplification by Stimulated Emission of Radiation. And then in July of 1960 Theodore H. Maiman announced the first successful operation of an optical maser or **laser**—certainly one of the great milestones in the history of Optics, and indeed in the history of science, had been achieved.

The laser is a quantum-mechanical device that manages to produce its "marvelous light" by taking advantage of the subtle ways in which atoms interact with electromagnetic radiation. To gain a solid, if only introductory, understanding of how the laser works and what makes its emissions so special, we'll first lay out some basic theory about ordinary thermal sources, such as lightbulbs and stars. That will require an introduction to blackbody radiation, but those insights are also basic to any treatment of the interaction of EM-radiation and matter. To that will be added a discussion of the Boltzmann distribution (p. 586) as applied to atomic energy levels. With this to stand on, we can appreciate the central notion of stimulated emission via the Einstein A and B coefficients (p. 586); the rest, more or less, follows.

*See James P. Gorden, "The Maser," *Sci. Am.* 199, **42** (December 1958).

13.1.1 Radiant Energy and Matter in Equilibrium

It shouldn't surprise anyone that if physics was to be turned on its head, it would be done while trying to figure out what light (i.e., radiant energy) was all about. Quantum theory had its earliest beginnings back in 1859 with the study of a seemingly obscure phenomenon known as **blackbody radiation**. That year, Charles Darwin published *The Origin of Species*, and Gustav Robert Kirchhoff proffered an intellectual challenge that would lead to a revolution in physics.

Kirchhoff was involved in analyzing the way bodies in thermal equilibrium behave in the process of exchanging radiant energy. This *thermal radiation* is electromagnetic energy emitted by all objects, the source of which is the random motion of their constituent atoms. He characterized the abilities of a body to emit and absorb electromagnetic energy by an *emission coefficient* ε_λ and an *absorption coefficient* α_λ. Epsilon is the energy per unit area per unit time emitted in a tiny wavelength range around λ (in units of $W/m^2/m$): thermal radiation comprises a wide range of frequencies, and an energy-measuring device by necessity admits a band of wavelengths. *Alpha is the fraction of the incident radiant energy absorbed per unit area per unit time in that wavelength range*; it's unitless. The emission and absorption coefficients depend on both the nature of the surface of the body (color, texture,

etc.) and the wavelength—a body that emits or absorbs well at one wavelength may emit or absorb poorly at another.

Consider an isolated chamber of some sort in thermal equilibrium at a fixed temperature T. Presumably, it would be filled with radiant energy at a myriad of different wavelengths—think of a glowing furnace. Kirchhoff assumed there was some formula, or *distribution function* $I_\lambda(\lambda)$, which depends on T and which provides values of the ***energy per unit area per unit time at each wavelength***; call it the ***spectral flux density*** within the cavity (or ***spectral exitance*** when it leaves it). He concluded that the *total* amount of energy at all wavelengths being absorbed by the walls versus the amount emitted by them must be the same, or else T would change, and it doesn't. Furthermore, Kirchhoff argued that if the walls were made of different materials (which behave differently with T), that same balance would have to apply for *each* wavelength range individually. The energy absorbed at λ, namely, $\alpha_\lambda I_\lambda$, must equal the energy radiated, ε_λ, *and this is true for all materials no matter how different.* **Kirchhoff's Radiation Law** is therefore

$$\frac{\varepsilon_\lambda}{\alpha_\lambda} = I_\lambda \qquad (13.1)$$

wherein the distribution I_λ, in units of $J/m^3 \cdot s$ or W/m^3, is a universal function the same for every type of cavity wall regardless of material, color, size, and shape and is only dependent on T and λ. That's quite extraordinary! Still, the British ceramist Thomas Wedgwood had commented long before (1792) that objects in a fired kiln all turned glowing red together along with the furnace walls, regardless of their size, shape, or material constitution.

Although Kirchhoff could not provide the energy distribution function in general, he did observe that a perfectly absorbing body, one for which $\alpha_\lambda = 1$, will appear black and, in that special case, $I_\lambda = \varepsilon_\lambda$. Moreover, the distribution function for a perfectly black object is the same as for an isolated chamber at that same temperature (visualize such a blackbody at equilibrium inside a hot oven). The radiant energy distribution at equilibrium within an isolated cavity is in every regard the same, "as if it came from a completely blackbody of the same temperature." Therefore ***the energy that would emerge from a small hole in the chamber should be identical to the radiation coming from a perfectly black object at the same temperature.***

The scientific community accepted the challenge of experimentally determining I_λ, but the technical difficulties were great and progress came slowly. The basic setup (Fig. 13.1*a*) is simple enough, although coming up with a reliable source was a daunting problem for a long time. Data must be extracted that is independent of the construction of the specific detector, and so the best thing to plot is the radiant energy per unit time, which enters the detector per unit area (of the entrance

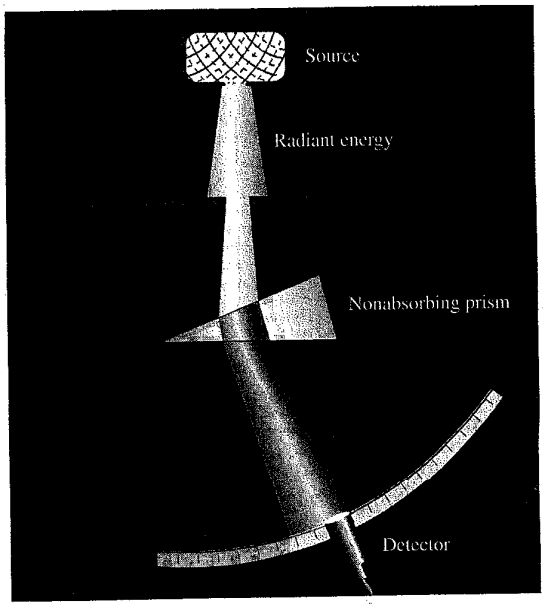

(a)

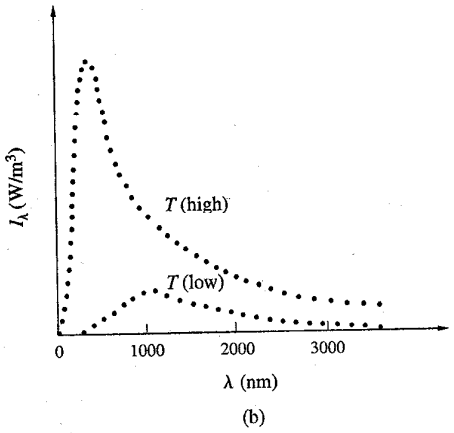

(b)

Figure 13.1 (*a*) A basic experimental setup for measuring blackbody radiation. (*b*) Values of I_λ at successive wavelengths as measured by a detector. Each curve corresponds to a specific source temperature.

window) per unit wavelength range (admitted by the detector). The kind of curves that were ultimately recorded are shown in Fig. 13.1*b*, and each is a plot of I_λ at a specific temperature.

Stefan-Boltzmann Law

In 1865 John Tyndall published some experimental results, including the determination that the total energy emitted by a heated platinum wire was 11.7 times greater when operating at 1200°C (1473 K) than it was at 525°C (798 K). Rather amazingly, Josef Stefan (1879) noticed that the ratio of $(1473\ \text{K})^4$ to $(798\ \text{K})^4$ was 11.6, nearly 11.7, and he surmised that the rate at which energy is radiated is proportional to T^4. In this observation Stefan was quite right (and quite lucky); Tyndall's results were actually far from those of a blackbody. Still, the conclusion was subsequently given a theoretical foundation by L. Boltzmann (1884). His was a traditional treatment of the radiation pressure exerted on a piston in a cylinder using the laws of thermodynamics and Kirchhoff's Law. The analysis progressed in much the same way one would treat a gas in a cylinder, but instead of atoms, the active agency was electromagnetic waves. The resulting Stefan–Boltzmann Law for blackbodies (which is correct, though nowadays we would derive it differently) is

$$P = \sigma A T^4 \tag{13.2}$$

where P is the total radiant power at all wavelengths, A is the area of the radiating surface, T is the absolute temperature in kelvins, and σ is a universal constant now given as

$$\sigma = 5.67033 \times 10^{-8}\ \text{W/m}^2\cdot\text{K}^4$$

The total area under any one of the blackbody-radiation curves of Fig. 13.1*b* for a specific T is the power per unit area, and from Eq. (13.2) that's just $P/A = \sigma T^4$.

Real objects are not perfect blackbodies; carbon black has an absorptivity of nearly one, but only at certain frequencies (obviously including the visible). Its absorptivity is much lower in the far infrared. Nonetheless, most objects resemble a blackbody (at least at certain temperatures and wavelengths)—you, for instance, are nearly a blackbody for infrared. Because of that, it's useful to write a similar expression for ordinary objects. This can be done by introducing a multiplicative factor called the total emissivity (ε), which relates the radiated power to that of a blackbody for which $\varepsilon = 1$, at the same temperature, thus

$$P = \varepsilon\sigma A T^4$$

Table 13.1 provides a few values of ε (at room temperature), where $0 < \varepsilon < 1$. Note that emissivity is unitless.

If an object with a *total absorptivity* of α is placed in an enclosure such as a cavity or a room having an emissivity ε_e and a temperature T_e, the body will radiate at a rate $\varepsilon\sigma A T^4$ and absorb energy inside the enclosure at a rate $\alpha(\varepsilon_e\sigma A T_e^4)$. Yet at any temperature at which the body and enclosure are in equilibrium (i.e., $T = T_e$), these rates must be equal; hence, $\alpha\varepsilon_e = \varepsilon$ and that has to be true for all temperatures. The net power radiated (when $T > T_e$) or absorbed (when $T < T_e$) by the body is then

$$P = \varepsilon\sigma A(T^4 - T_e^4)$$

All bodies not at zero kelvin radiate, and the fact that T is raised to the fourth power makes the radiation highly sensitive to temperature changes. When a body at 0°C (273 K) is brought up to 100°C (373 K), it radiates about 3.5 times the previous power. Increasing the temperature increases the net power radiated; that's why it gets more and more difficult to increase the temperature of an object. (Try heating a steel spoon to 1300°C.) Increasing the temperature of an object also shifts the emitted distribution of energy among the various wavelengths present. At the moment when the filament of a lightbulb "blows," the resistance, current, and temperature rise; it goes from its normal operating reddish-white color to a bright flash of blue-white.

TABLE 13.1 Some Representative Values of Total Emissivity*

Material	ε
Aluminum foil	0.02
Copper, polished	0.03
Copper, oxidized	0.5
Carbon	0.8
White paint, flat	0.87
Red brick	0.9
Concrete	0.94
Black paint, flat	0.94
Soot	0.95

*T = 300 K, room temperature.

Wien Displacement Law

Perhaps the last notable success in applying classical theory to the problem of blackbody radiation came in 1893 at the hands of the German physicist and Nobel laureate Wilhelm Otto Fritz Franz Wien (1864–1928), known to his friends as Willy. He derived what is today called the **Displacement Law**. Each blackbody curve reaches a maximum height at a value of wavelength (λ_{max}) that is particular to it and therefore to the absolute temperature T. At that wavelength, the blackbody radiates the most energy. Wien was able to show that

$$\lambda_{max} T = \text{constant} \tag{13.3}$$

where the constant was found experimentally to be 0.002 898 m·K. The peak wavelength is inversely proportional to the temperature. *Raise the temperature, and the bulk of the radiation shifts to shorter wavelengths and higher frequencies* (see the dashed curve in Fig. 13.2). As a glowing coal or a blazing star gets hotter, it goes from IR warm to red-hot to blue-white. A person or a piece of wood, both only roughly blackbodies, radiates for the most part in the infrared and would only begin to glow faintly in the visible at around 600°C or 700°C, long after either had decomposed. The bright cherry red of a chunk of "red" hot iron sets in at around 1300°C.

In 1899 researchers greatly advanced the state of experimentation by using, as a source of blackbody radiation, a small hole in a heated cavity (Fig. 13.3). Energy entering such an aperture reflects around inside until it's absorbed. (The pupil of the eye appears black for precisely the same reason.) A near-perfect absorber is a near-perfect emitter, and the region of a small hole in the face of an oven is a wonderful source of *blackbody radiation.*

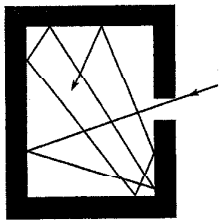

Figure 13.3 Radiant energy entering a tiny hole in a chamber will rattle around with little chance of ever emerging through the aperture, and so the hole looks black. In reverse, the aperture of a heated chamber appears as a blackbody source.

It was at this point in time that classical theory began to falter. All attempts to fit the entire radiation curve (Fig. 13.2) with some theoretical expression based on electromagnetism led only to the most limited successes. Wien produced a formula that agreed with the observed data fairly well in the short wavelength region but deviated from it substantially at large λ. Lord Rayleigh and later Sir James Jeans (1877–1946) developed a description in terms of the standing-wave modes of the field within the enclosure. But the resulting *Rayleigh–Jeans formula* matched the experimental curves only in the very long wavelength region. The failure of classical theory was totally inexplicable; a turning point in the history of physics had arrived.

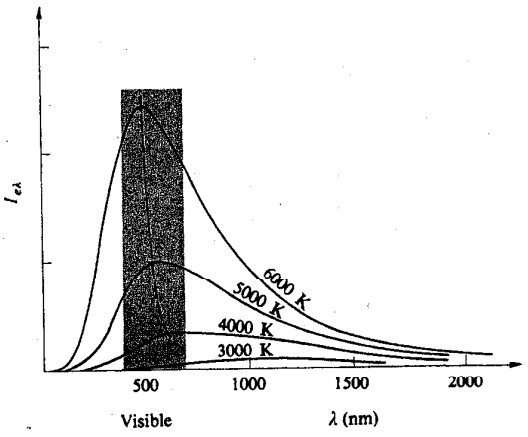

Figure 13.2 Blackbody radiation curves. The hyperbola passing through peak points corresponds to Wien's Law.

Planck Radiation Law

Max Karl Ernst Ludwig Planck at 42 was the somewhat reluctant father of quantum theory. Like so many other theoreticians at the turn of the century, he, too, was working on blackbody radiation. But Planck would succeed not only in producing Kirchhoff's distribution function, but also in turning physics upside-down in the process. We cannot follow the details of his derivation here; besides, the original version was wrong. (Bose and Einstein corrected it years later.) Still, it had such a powerful impact that it's worth looking at some of the features that are right.

Planck knew that if an arbitrary distribution of energetic molecules was injected into a constant-temperature chamber, it would ultimately rearrange itself into the Maxwell–Boltzmann distribution of speeds as it inevitably reached equilibrium. Presumably, if an arbitrary distribution of radiant energy

is injected into a constant-temperature cavity, it, too, will ultimately rearrange itself into the Kirchhoff distribution of energies as it inevitably reaches equilibrium.

In October 1900, Planck produced a distribution formula that was based on the latest experimental results. This mathematical contrivance, concocted "by happy guesswork," fit all the data available. It contained two fundamental constants, one of which (h) would come to be known as **Planck's Constant**. That much by itself was quite a success, even if it didn't explain anything. Although Planck had no idea of it at the time, he was about to take a step that would inadvertently revolutionize our perception of the physical Universe.

Naturally enough, Planck set out to construct a theoretical scheme that would logically lead to the equation he had already devised. He assumed that the radiation in a chamber interacted with simple microscopic oscillators of some unspecified type. These vibrated on the surfaces of the cavity walls, absorbing and reemitting radiant energy independent of the material. (In fact, the atoms of the walls do exactly that. Because of their tightly packed configuration in the solid walls, the atoms interact with a huge number of their neighbors. That completely blurs their usual characteristic sharp resonance vibrations, allowing them to oscillate over a broad range of frequencies and emit a continuous spectrum.) Try as he might, Planck was unsuccessful. At that time, he was a devotee of E. Mach, who had little regard for the reality of atoms, and yet the obstinate insolubility of the problem ultimately led Planck to "an act of desperation." He hesitantly turned to Boltzmann's "distasteful" statistical method, which had been designed to deal with the clouds of atoms that constitute a gas.

Boltzmann, the great proponent of the atom, and Planck were intellectual adversaries for a while. And now Planck was forced to use his rival's statistical analysis, which—ironically—he would misapply. If Boltzmann's scheme for counting atoms was to be applied to something continuous, such as energy, some adjustments would have to be made in the procedure. Thus, according to Planck, the total energy of the oscillators had to be thought of, at least temporarily, as apportioned into "energy elements" so that they could be counted. These energy elements were given a value proportional to the frequency ν of the resonators. Remember that he already had the formula he was after, and in it there appeared the term $h\nu$. Planck's Constant,

$$6.626\,075\,5 \times 10^{-34}\ \text{J}\cdot\text{s} \quad \text{or} \quad 4.135\,669\,2 \times 10^{-15}\ \text{eV}\cdot\text{s}$$

is a very small number and so $h\nu$, which has the units of energy, is itself a very small quantity. Accordingly, he set the value of the energy element equal to it: $\mathscr{E} = h\nu$.

This was a statistical analysis, and counting was central. Still, when the method was applied as Boltzmann intended, it naturally smoothed out energy, making it continuous as usual. Again, we needn't worry about the details. The amazing thing was that Planck had stumbled on a hidden mystery of nature: **energy is quantized**—it comes in tiny bursts, but he didn't realize it then.

Planck derived the following formula for the spectral exitance (or spectral irradiance)—which he had already arrived at by fitting curves to the data—and it's the answer to Kirchhoff's challenge:

$$I_\lambda = \frac{2\pi hc^2}{\lambda^5}\left[\frac{1}{e^{\frac{hc}{\lambda k_B T}} - 1}\right] \tag{13.4}$$

where k_B is Boltzmann's Constant. This is **Planck's Radiation Law**, and, of course, it fit blackbody data splendidly (Fig. 13.4). Notice how the expression contains the speed of light, Boltzmann's Constant, and Planck's Constant (h). It bridges Electromagnetic Theory to the domain of the atom.

Although Eq. (13.4) represents a great departure from previous ideas, Planck did not mean to break with classical theory. It would have been unthinkable for him even to suggest that radiant energy was anything but continuous. "That energy is forced, at the outset, to remain together in certain quanta ...," Planck later remarked, "was purely a formal assumption and I really did not give it much thought." It was only around 1905, at the hands of a much bolder thinker, Albert Einstein, that we learned that the atomic oscillators

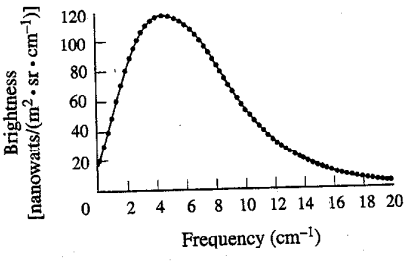

Figure 13.4 The cosmic background radiation of the Universe. Since the creation of the Universe with the Big Bang, it has expanded and cooled. The data points (measured in the microwave spectrum) were detected by the Cosmic Background Explorer (COBE) satellite. The solid line is the Planck blackbody curve for a temperature of 2.735 ± 0.06 K.

were real and that their energies were quantized. Each oscillator could only exist with an energy that was a whole-number multiple of $h\nu$ (a little like the *gravitational*-PE of someone walking up a flight of stairs). Moreover, **radiant energy itself is quantized**, existing in localized blasts of an amount $\mathscr{E} = h\nu$.

13.1.2 Stimulated Emission

The LAser accomplishes "light amplification" by making use of energetic atoms in a medium to reinforce the light field. Let's therefore examine the manner in which the energy states of a system of atoms at some arbitrary temperature is normally distributed. The problem is part of the broader discipline of Statistical Mechanics and is addressed specifically in terms of the Maxwell–Boltzmann distribution.

Population of Energy Levels

Imagine a chamber filled with a gas in equilibrium at some temperature T. If T is relatively low, as it is in a typical room, most of the atoms will be in their ground states, but a few will momentarily pick up enough energy to "rise" into an excited state. The classical Maxwell–Boltzmann distribution maintains that, on average, a number of atoms per unit volume, N_i, will be in any excited state of energy $\mathscr{E}_i$ such that

$$N_i = N_0 e^{-\mathscr{E}_i/k_B T}$$

where N_0 is a constant for a given temperature. *The higher the energy state, that is, the greater the value of* $\mathscr{E}$ *(the smaller is the exponential) and the fewer atoms there will be in that state.*

Since we will be interested in atomic transition between arbitrary states, consider the jth energy level where $\mathscr{E}_j > \mathscr{E}_i$. Then for it $N_j = N_0 e^{-\mathscr{E}_j/k_B T}$, and the ratio of the populations occupying these two states is

$$\frac{N_j}{N_i} = \frac{e^{-\mathscr{E}_j/k_B T}}{e^{-\mathscr{E}_i/k_B T}} \tag{13.5}$$

This is the *relative population*, and it follows that

$$N_j = N_i e^{-(\mathscr{E}_j - \mathscr{E}_i)/k_B T} = N_i e^{-h\nu_{ji}/k_B T} \tag{13.6}$$

where use was made of the fact that a transition for the jth-state to the ith-state corresponds to an energy change of $(\mathscr{E}_j - \mathscr{E}_i)$ and since such transitions are accompanied by the emission of a photon of frequency ν_{ji}, we can substitute $(\mathscr{E}_j - \mathscr{E}_i) = h\nu_{ji}$.

The Einstein *A* and *B* Coefficients

In 1916 Einstein devised an elegant and rather simple theoretical treatment of the dynamic equilibrium existing for a material medium bathed in electromagnetic radiation, absorbing and reemitting. The analysis was used to affirm Planck's Radiation Law, but more importantly it also created the theoretical foundation for the laser. The reader should already be familiar with the basic mechanism of *absorption* (see Fig. 3.35, p. 64). Suppose the atom is in its lowest energy or ground-state configuration. A photon having an adequate amount of energy interacts with the atom, imparting that energy to the atom, thereby causing the electron cloud to take on a new configuration. The atom jumps into a higher-energy excited state (Fig. 13.5). In a dense medium, the atom is likely to interact with its jiggling neighbors and pass off its bounty of energy via collisions.

Such an excess-energy configuration is usually (though not always) exceedingly short-lived, and in 10 ns or so, without the intercession of any external influence, the atom will emit its overload of energy as a photon. As it does, it reverts to a stable state in a process called *spontaneous emission* (Fig. 13.5*b*).

The remarkable thing is that there is a third alternative process, one first appreciated by Einstein and crucial to the operation of the laser—which wasn't invented until almost a half century later. For a medium inundated with EM-radiation, it's possible for a photon to interact with an excited atom while that atom is still in its higher-energy configuration. The atom can then dump its excess energy in-step with the incoming photon, in a process now called **stimulated emission** (Fig. 13.5).

In the case of absorption, the rate-of-change of the number of atoms in some initial state, as they leave to some higher state, must depend on the strength of the photon field inundating those atoms. In other words, it must depend on the energy density u, given by Eq. (3.34), but more specifically it must depend on the energy density in the frequency range driving the transition, that is, the spectral energy density u_ν, which has units of joules per meter-squared per inverse second ($\mathrm{J \cdot s/m^2}$). (Note that if we consider the radiation field as a photon gas, the spectral energy density can be thought of as the photon

density per unit frequency range.) The rate-of-change of the number of atoms, the **transition rate**, will also be proportional to the population, that is, the number density of atoms in that state (N_i); the more there are, the more can leave (via absorption) per second. Because the process is driven by the photon field, let's call it **stimulated absorption**, whereupon the transition rate is

[stimulated absorption]
$$\left(\frac{dN_i}{dt}\right)_{ab} = -B_{ij}N_i u_\nu \tag{13.7}$$

Here B_{ij} is a constant of proportionality, and the minus arises because N_i is decreasing. Similarly, for stimulated emission

[stimulated emission]
$$\left(\frac{dN_j}{dt}\right)_{st} = -B_{ji}N_j u_\nu \tag{13.8}$$

In the case of spontaneous emission, the process is independent of the field environment and

[spontaneous emission]
$$\left(\frac{dN_j}{dt}\right)_{sp} = -A_{ji}N_j \tag{13.9}$$

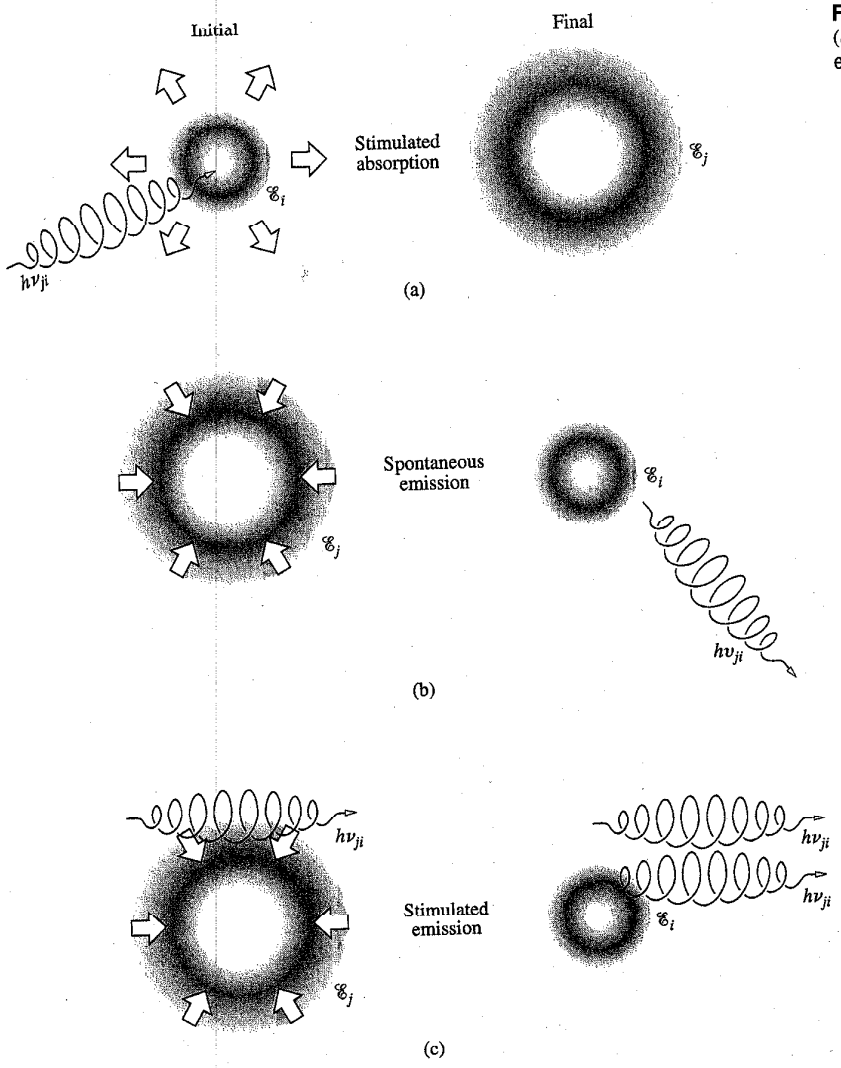

Figure 13.5 A schematic representation of (*a*) stimulated absorption, (*b*) spontaneous emission, and (*c*) stimulated emission.

Keep in mind that the transition rate, the number of atoms making transitions per second, divided by the number of atoms, is the probability of a transition occurring per second, $\mathcal{P}$. Consequently, the probability per second of spontaneous emission is $\mathcal{P}_{sp} = A_{ji}$.

For a single excited atom making a spontaneous transition to a lower state, the inverse of the transition probability per second is the **mean life** or **lifetime** of the excited state τ. Thus (operating under conditions that exclude any other mechanism but spontaneous emission), if N atoms are in that excited state, the total rate of transitions, that is, the number of emitted photons per second, is $N\mathcal{P}_{sp} = NA_{ji} = N/\tau$. A low-transition probability means a long lifetime.

The three constants A_{ji}, B_{ji}, and B_{ij} are Einstein's coefficients. Following his lead, we assume (1) that thermodynamic equilibrium exists between the radiation field and the atoms in it at any T; (2) that the energy density has the characteristics of a blackbody at T; and (3) that the number densities of the two states are in accord with the Maxwell–Boltzmann distribution.

Given that the system is in equilibrium, the rate of upward ($i \rightarrow j$) transitions must equal the rate of downward transitions ($j \rightarrow i$):

$$B_{ij}N_i u_\nu = B_{ji}N_j u_\nu + A_{ji}N_j$$

Dividing both sides by N_i and rearranging terms yields

$$\frac{N_j}{N_i} = \frac{B_{ij}u_\nu}{A_{ji} + B_{ji}u_\nu}$$

Making use of Eq. (13.6), that is, what we found from the application of the Maxwell–Boltzmann distribution, this becomes

$$e^{-h\nu_{ji}/k_B T} = \frac{B_{ij}u_\nu}{A_{ji} + B_{ji}u_\nu}$$

and solving for u_ν leads to

$$u_\nu = \frac{A_{ji}/B_{ji}}{(B_{ij}/B_{ji})e^{h\nu_{ji}/k_B T} - 1} \tag{13.10}$$

Here Einstein pointed out that as $T \rightarrow \infty$, the spectral energy density, that is, the spectral photon density, approaches infinity. Figure 13.3 shows that I_λ increases with T, and that implies that u_ν will behave in a like fashion. In fact, $I_\nu = \frac{c}{4}u_\nu$, a point we will address presently. In any event, since $e^0 = 1$, the only way u_ν will be large is if

$$B_{ij} = B_{ji} = B$$

for large T, but since these constants are temperature independent, they must be equal at all T. The probabilities of stimulated emission and absorption are $\mathcal{P}_{st} = B_{ji}u_\nu$ and $\mathcal{P}_{ab} = B_{ij}u_\nu$, respectively. Hence, **the probability of stimulated emission is identical to the probability of stimulated absorption**; an atom in the lower state is just as likely to make a stimulated transition up, as an excited atom is to make a stimulated transition down.

Simplifying the notation (let $A = A_{ji}$), Eq. (13.10) becomes

$$u_\nu = \frac{A}{B}\left[\frac{1}{e^{h\nu_{ji}/k_B T} - 1}\right] \tag{13.11}$$

The ratio A/B can be expressed via basic quantities by comparing this equation with

$$I_\lambda = \frac{2\pi hc^2}{\lambda^5}\left[\frac{1}{e^{\frac{hc}{\lambda k_B T}} - 1}\right] \tag{13.4}$$

But first transform I_λ into I_ν where these are expressions for exitance (which is irradiance going outward) per interval $d\lambda$ and $d\nu$, respectively. Using the fact that $\lambda = c/\nu$, differentiating yields $d\lambda = -cd\nu/\nu^2$. Because $I_\lambda d\lambda = I_\nu d\nu$, and dropping the sign (since it just says that one differential increases while the other decreases), we get $I_\lambda c/\nu^2 = I_\nu$; and so

$$I_\nu = \frac{2\pi h\nu^3}{c^2}\left[\frac{1}{e^{\frac{h\nu}{k_B T}} - 1}\right] \tag{13.12}$$

Now as a last step we need only to compare the spectral energy density u_ν in the chamber with the spectral exitance,

$$I_\nu = \frac{c}{4}u_\nu \tag{13.13}$$

emerging from it. Rather than burden the reader with a complete derivation of this relationship, let it suffice merely to justify it. Keep in mind that I_ν corresponds to a flow of energy across a unit normal area, in one side and out the other—a beam leaving the chamber. In Section 3.3.1 we saw that the instantaneous flow of power per unit normal area, the Poynting vector, was given by $S = cu$ and so on average $I = cu$ for a beam. Inside a chamber, however, with light traveling in every direction, not all the photons that contribute to u will contribute to the exitance in a particular direction. Presumably, inside the chamber a unit area held horizontally would have as much energy flowing up through it as down. Moreover, only the components perpendicular to the area contribute to S, so a factor of $1/4$ is not unreasonable.

From Eqs. (13.11), (13.12), and (13.14) it follows that

$$\frac{A}{B} = \frac{8\pi h\nu^3}{c^3} \tag{13.14}$$

The probability of spontaneous emission is proportional to the probability of stimulated emission; an atom susceptible to one mechanism is proportionately susceptible to the other. Lasers work by stimulated emission, and anything that enhances spontaneous emission (i.e., A) at the price of stimulated emission (i.e., B) can be expected to work to the detriment of the process. Because the ratio of A/B varies as ν^3, it would seem that X-ray lasers ought to be very difficult to build—they are!

Imagine a system of atoms in thermal equilibrium having only two possible states. Furthermore, require that the atoms have a long mean life so that we can ignore spontaneous emission. When the system is inundated by photons of the proper energy, stimulated absorption depopulates the lower i-level, while stimulated emission depopulates the upper j-level. The number of photons vanishing from the system per second via stimulated absorption is proportional to $\mathcal{P}_{ab}N_i$, and the number entering it via stimulated emission is proportional to $\mathcal{P}_{st}N_j$, but from the equality of the B-coefficients it follows that $\mathcal{P}_{st} = \mathcal{P}_{ab}$. Therefore $\mathcal{P}_{ab}N_j = \mathcal{P}_{st}N_j$. However, if the system is in thermal equilibrium, $N_i > N_j$, which means that the number of photons vanishing per second exceeds the number entering per second; there's a net absorption of photons by the lower state because there are more atoms in the lower state at any given temperature. The reverse would be true if we could create a situation—a *population inversion*—in which $N_i < N_j$; then stimulated emission would dominate over stimulated absorption.

13.1.3 The Laser

Consider an ordinary medium in which a few atoms are in some excited state; call it $|j\rangle$ to conform with quantum-mechanical notation. If a photon in an incident beam is to trigger one of these excited atoms into stimulated emission, it must have the frequency ν_{ji}, as per Fig. 13.5*c*. A remarkable feature of this process is that ***the emitted photon is in-phase with, has the polarization of, and propagates in the same direction as, the stimulating radiation***. The emitted photon is said to be in the same radiation mode as the incident wave and tends to add to it, increasing its flux density. However, since most atoms are ordinarily in the ground state, absorption is usually far more likely than stimulated emission.

This raises an intriguing point: What would happen if a substantial percentage of the atoms could somehow be excited into an upper state, leaving the lower state all but empty? For obvious reasons this is known as **population inversion**. An incident photon of the proper frequency could then trigger an avalanche of stimulated photons—*all in-phase*. The initial wave would continue to build, so long as there were no dominant competitive processes (such as scattering) and provided the population inversion could be maintained. In effect, energy (electrical, chemical, optical, etc.) would be pumped in to sustain the inversion, and a beam of light would be extracted after sweeping across the *active medium*.

The First (Pulsed Ruby) Laser

To see how all of this is accomplished in practice, let's take a look at Maiman's original device (Fig. 13.6). The first operative laser had as its active medium a small, cylindrical, synthetic, pale pink ruby, that is, an Al_2O_3 crystal containing about 0.05 percent (by weight) of Cr_2O_3. Ruby, which is still one of the most common of the crystalline laser media, had been used earlier in maser applications and was suggested for use in the laser by Schawlow. The rod's end faces were polished flat, parallel and normal to the axis. Then both were silvered (one only partially) to form a **resonant cavity**.

It was surrounded by a helical gaseous discharge flashtube, which provided broadband **optical pumping**. Ruby appears red because the chromium atoms have absorption bands in the blue and green regions of the spectrum (Fig. 13.7*a*). Firing the flashtube generates an intense burst of light lasting a few milliseconds. Much of this energy is lost in heat, but many of the Cr^{3+} ions are excited into the absorption bands. A simplified energy-level diagram appears in Fig. 13.7*b*. The excited ions rapidly relax (in about 100 ns), giving up energy to the crystal lattice and making nonradiative transitions. They preferentially drop "down" to a pair of closely spaced, especially long-lived, interim states. They remain in these so-called **metastable states** for up to several milliseconds ($\approx$3 ms at room temperature) before randomly, and in most cases spontaneously, dropping down to the ground state. This is accompanied by the emission of the characteristic red fluorescent radiation of ruby. The lower-level transition dominates, and the resulting emission occurs in a relatively broad spectral range centered about 694.3 nm; it emerges in all directions and is incoherent.

When the pumping rate is increased somewhat, a population inversion occurs, and the first few spontaneously emitted

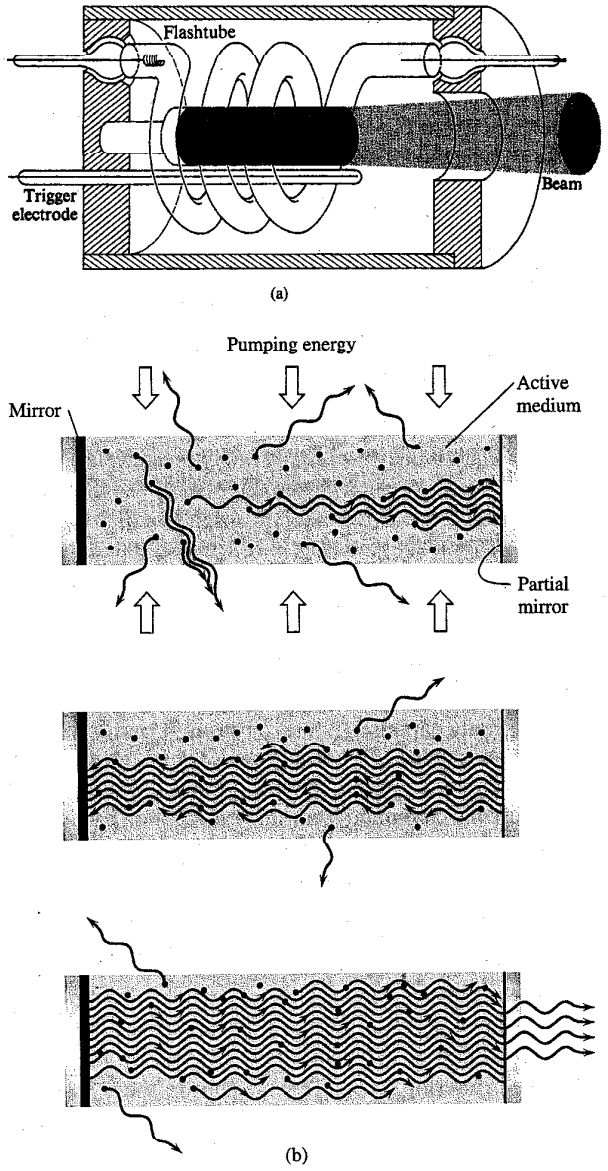

Figure 13.6 The first ruby-laser configuration, just about life-sized.

photons stimulate a chain reaction. One quantum triggers the rapid, in-phase emission of another, dumping energy from the metastable atoms into the evolving lightwave (Fig. 13.6*b*). The wave continues to grow as it sweeps back and forth across the active medium (provided enough energy is available to overcome losses at the mirrored ends). Since one of those reflecting surfaces was partially silvered, an intense pulse of red laser light (lasting about 0.5 ms and having a line-width of about 0.01 nm) emerges from that end of the ruby rod.

Notice how neatly everything works out. The broad absorption bands make the initial excitation rather easy, while the long lifetime of the metastable state facilitates the population inversion. The atomic system in effect consists of (1) the absorption bands, (2) the metastable state, and (3) the ground state. Accordingly, it is spoken of as a *three-level* laser.

Today's ruby laser is generally a high-power source of pulsed coherent radiation used extensively in work on interferometry, plasma diagnostics, holography, and so forth. Such devices operate with coherence lengths of from 0.1 m to 10 m.

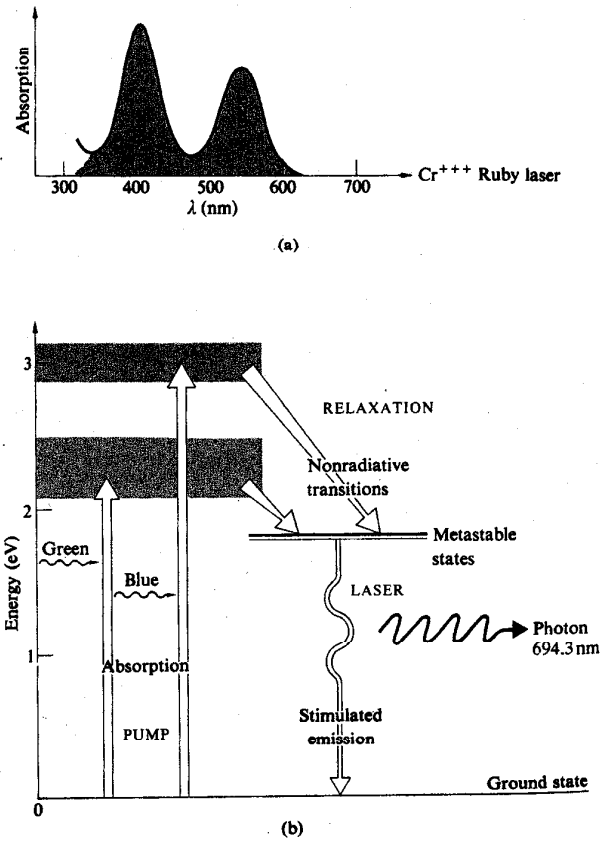

Figure 13.7 Ruby-laser energy levels.

Modern configurations usually have flat external mirrors, one totally and the other partially reflecting. As an oscillator, the ruby laser generates millisecond pulses in the energy range from around 50 J to 100 J, but by using a tandem oscillator-amplifier setup, energies well in excess of 100 J can be produced. The commercial ruby laser typically operates at a modest overall efficiency of less than 1%, producing a beam that has a diameter ranging from 1 mm to about 25 mm, with a divergence of from 0.25 mrad to about 7 mrad.

Optical Resonant Cavities

The resonant cavity, which in this case is of course a Fabry–Perot etalon, plays a significant role in the operation of the laser. In the early stages of the laser process, spontaneous photons are emitted in every direction, as are the stimulated photons. But all of these, with the singular exception of those propagating very nearly along the cavity axis, quickly pass out of the sides of the ruby. In contrast, the axial beam continues to build as it bounces back and forth across the active medium. This accounts for the amazing degree of collimation of the issuing laserbeam, which is then effectively a coherent plane wave. Although the medium acts to amplify the wave, the *optical feedback* provided by the cavity converts the system into an oscillator and hence into a light generator—the acronym is thus somewhat of a misnomer.

In addition, the disturbance propagating within the cavity takes on a standing-wave configuration determined by the separation (L) of the mirrors. The cavity resonates (i.e., standing waves exist within it) when there is an integer number (m) of half wavelengths spanning the region between the mirrors. The idea is simply that there must be a node at each mirror, and this can only happen when L equals a whole number multiple of $\lambda/2$ (where $\lambda = \lambda_0/n$). Thus

$$m = \frac{L}{\lambda/2}$$

and

$$\nu_m = \frac{mv}{2L} \tag{13.15}$$

There are therefore an infinite number of possible oscillatory **longitudinal cavity modes**, each with a distinctive frequency ν_m. Consecutive modes are separated by a constant difference,

$$\nu_{m+1} - \nu_m = \Delta\nu = \frac{v}{2L} \tag{13.16}$$

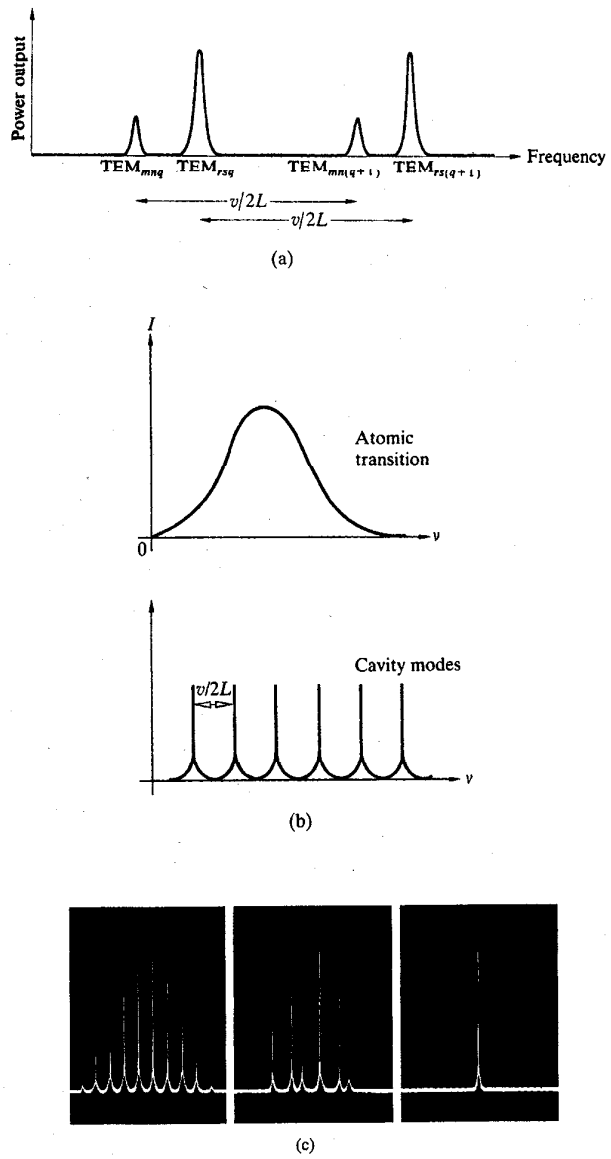

Figure 13.8 Laser modes: (*a*) illustrates the nomenclature; (*b*) compares the broad atomic emission with the narrow cavity modes; (*c*) depicts three operation configurations for a c-w gas laser, first showing several longitudinal modes under a roughly Gaussian envelope, then several longitudinal and transverse modes, and finally a single longitudinal mode.

which is the free spectral range of the etalon [Eq. (9.79)] and, incidentally, the inverse of the round-trip time. For a gas laser 1 m long, $\Delta\nu \approx 150$ MHz.

The resonant modes of the cavity are considerably narrower in frequency than the bandwidth of the normal spontaneous atomic transition. These modes, whether the device is constructed so that there is one or more, will be the ones that are sustained in the cavity, and hence the emerging beam is restricted to a region close to those frequencies (Fig. 13.8). In other words, the radiative transition makes available a relatively broad range of frequencies out of which the cavity will select and amplify only certain narrow bands and, if desired, even only one such band. This is the origin of the laser's extreme quasimonochromaticity. Thus, while the bandwidth of the ruby transition to the ground state is roughly a rather broad 0.53 nm (330 GHz)—because of interactions of the chromium ions with the lattice—the corresponding laser cavity bandwidth, the frequency spread of the radiation of a single resonant mode, is a much narrower 0.000 05 nm (30 MHz). This situation is depicted in Fig. 13.8*b*, which shows a typical transition lineshape and a series of corresponding cavity spikes—in this case each is separated by $v/2L$, and each is 30 MHz wide.

A possible way to generate only a single mode in the cavity would be to have the mode separation, as given by Eq. (13.16), exceed the transition bandwidth. Then only one mode would fit within the range of available frequencies provided by the transition. For a ruby laser (with an index of refraction of 1.76) a cavity length of a few centimeters will easily ensure single longitudinal mode operation. The drawback of this particular approach is that it limits the length of the active region contributing energy to the beam and so limits the output power of the laser.

In addition to the longitudinal or axial modes of oscillation, which correspond to standing waves set up along the cavity or z-axis, **transverse modes** can be sustained as well. Since the fields are very nearly normal to z, these are known as TEM_{mn} modes (transverse electric and magnetic). The m and n subscripts are the integer number of transverse nodal lines in the x- and y-directions across the emerging beam. That is to say, the beam is segmented in its cross section into one or more regions. Each such array is associated with a given TEM mode, as shown in Figs. 13.9 and 13.10. The lowest order or TEM_{00} transverse mode is perhaps the most widely used, and this for several compelling reasons: the flux density is ideally

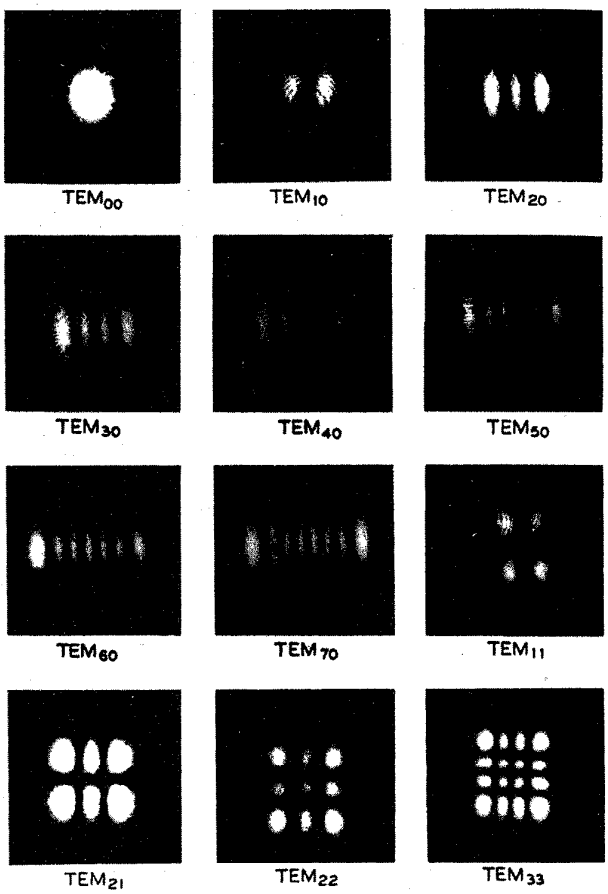

Figure 13.9 Mode patterns (without the faint interference fringes this is what the beam looks like in cross section). (Photos courtesy Bell Telephone Laboratories.)

Gaussian over the beam's cross section (Fig. 13.11); there are no phase shifts in the electric field across the beam, as there are in other modes, and so it is completely spatially coherent; the beam's angular divergence is the smallest; and it can be focused down to the smallest-sized spot. Note that the amplitude in this mode is actually not constant over the wavefront, and it is consequently an inhomogeneous wave.

A complete specification of each mode has the form TEM_{mnq}, where q is the longitudinal mode number. For each transverse mode (m, n) there can be many longitudinal modes

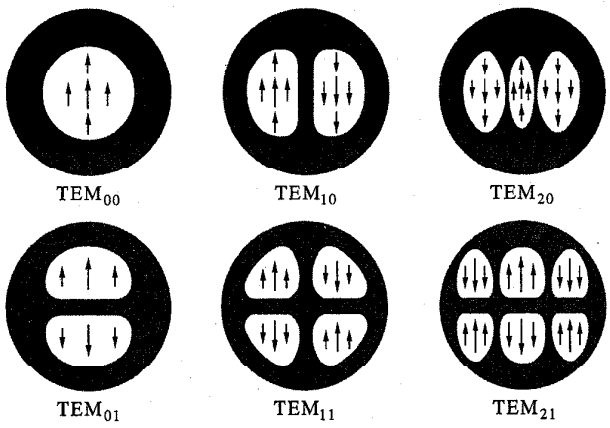

Figure 13.10 Mode configurations (rectangular symmetry). Circularly symmetric modes are also observable, but any slight asymmetry (such as Brewster windows) destroys them.

(i.e., values of q). Often, however, it's unnecessary to work with a particular longitudinal mode, and the q subscript is usually simply dropped.*

Several additional cavity arrangements are of considerably more practical significance than is the original plane-parallel setup (Fig. 13.12). For example, if the planar mirrors are replaced by identical concave spherical mirrors separated by a distance very nearly equal to their radius of curvature, we have the *confocal* resonator. Thus the focal points are almost coincident on the axis midway between mirrors—ergo the name confocal.

If one of the spherical mirrors is made planar, the cavity is termed a *hemispherical* or *hemiconcentric* resonator. Both of these configurations are considerably easier to align than is the plane-parallel form. Laser cavities are either *stable* or *unstable* to the degree that the beam tends to retrace itself and so remain relatively close to the optical axis (Fig. 13.13). A beam in an unstable cavity will "walk out," going farther from the axis on each reflection until it quickly leaves the cavity altogether. By contrast, in a stable configuration (with mirrors that are, say, 100% and 98% reflective) the beam might traverse the resonator 50 times or more. Unstable resonators are commonly used in high-power lasers, where the fact that the beam traces across a wide region of the active medium enhances the amplification and allows for more energy to be extracted. This approach will be especially useful for media (like carbon dioxide or argon) wherein the beam gains a good deal of energy on each sweep of the cavity. The needed number of sweeps is determined by the so-called *small-signal gain* of the active medium. The actual selection of a resonator configuration is governed by the specific requirements of the system—there is no universally best arrangement.

The decay of energy in a cavity is expressed in terms of the Q or **quality factor** of the resonator. The origin of the expression dates back to the early days of radio engineering, when it was used to describe the performance of an oscillating (tuning) circuit. A high-Q, low-loss circuit meant a narrow bandpass and a sharply tuned radio. If an optical cavity is somehow disrupted, as for example by the displacement or removal of one of the mirrors, the laser action generally ceases. When this is done deliberately in order to delay the onset of oscillation in the laser cavity, it's known as *Q-spoiling* or **Q-switching**. The power output of a laser is self-limited in the sense that the population inversion is continuously depleted through stimulated emission by the radiation field within the cavity. However, if oscillation is prevented, the number of atoms pumped into the (long-lived) metastable state can be considerably increased, thereby creating a very extensive population inversion. When the cavity is switched on at the proper moment, a tremendous-

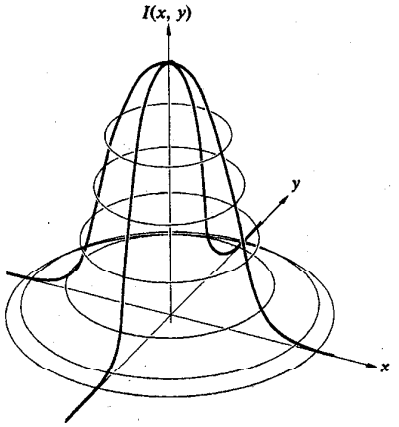

Figure 13.11 Gaussian irradiance distribution.

*Take a look at R. A. Phillips and R. D. Gehrz, "Laser Mode Structure Experiments for Undergraduate Laboratories," *Am. J. Phys.* **38**, 429 (1970).

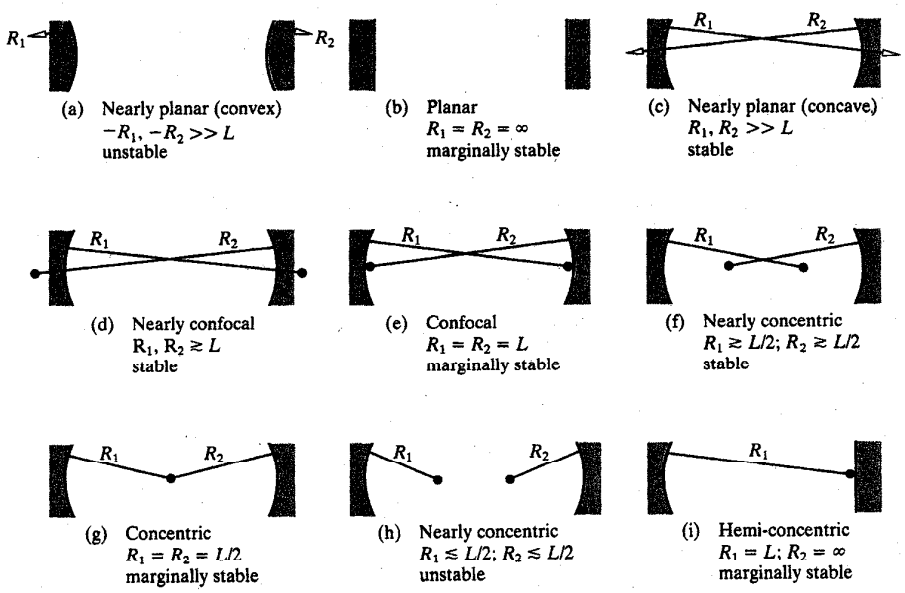

Figure 13.12 Laser cavity configurations. (Adapted from O'Shea, Callen, and Rhodes, *An Introduction to Lasers and Their Applications.*)

ly powerful *giant pulse* (perhaps up to several hundred megawatts) will emerge as the atoms drop down to the lower state almost in unison. A great many *Q-switching* arrangements utilizing various control schemes, for example, bleachable absorbers that become transparent under illumination, rotating prisms and mirrors, mechanical choppers, ultrasonic cells, or electro-optic shutters such as Kerr or Pockels cells, have all been used.

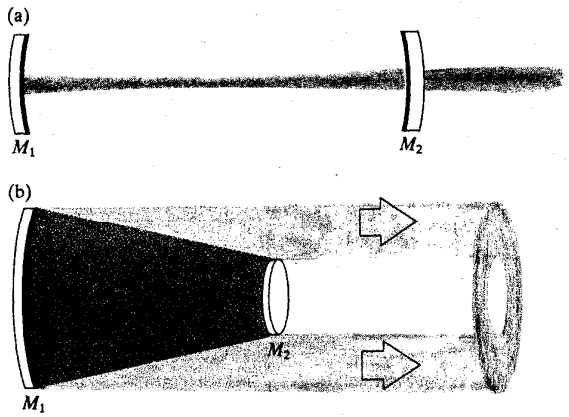

Figure 13.13 Stable and unstable laser resonators. (Adapted from O'Shea, Callen, and Rhodes, *An Introduction to Lasers and Their Applications.*)

Gaussian Laserbeams

The TEM_{00} mode that develops within a resonator has a Gaussian profile (Fig. 13.11); that is, the strength of the beam-like wave falls off transversely following a bell-shaped curve that's symmetrical around the central axis (Fig. 13.14*a*). Recall (p. 12) that a Gaussian is a negative exponential that's a function of the square of the variable, in this case, the distance (r) measured, in a transverse plane, from the central axis of propagation (z). Because the beam trails off radially it's useful to put an arbitrary boundary to its width. Accordingly, let $r = w$ be the *beam half-width*, the distance at which the electric field of the beam drops from its maximum axial value of E_0 to E_0/e or $37\%E_0$. At $r = w$ the beam's irradiance, which depends on the square of the amplitude, is then I_0/e^2 which is only $14\%I_0$. Most of the energy of the beam resides within this imaginary cylinder of radius w, where (Fig. 13.14*b*)

$$I = I_0 e^{-2r^2/w^2}$$

and $I = I_0 e^{-2}$, as it's supposed to, at $r = w$.

As can be seen in Fig. 13.13*a*, when curved mirrors form the laser cavity there is a tendency to "focus" the beam, giving it a minimum cross section or *waist* of radius w_0. Under such circumstances, the external divergence of the laserbeam is essentially a continuation of the divergence out from this waist (Fig. 13.15). In general, there will be a beam waist somewhere

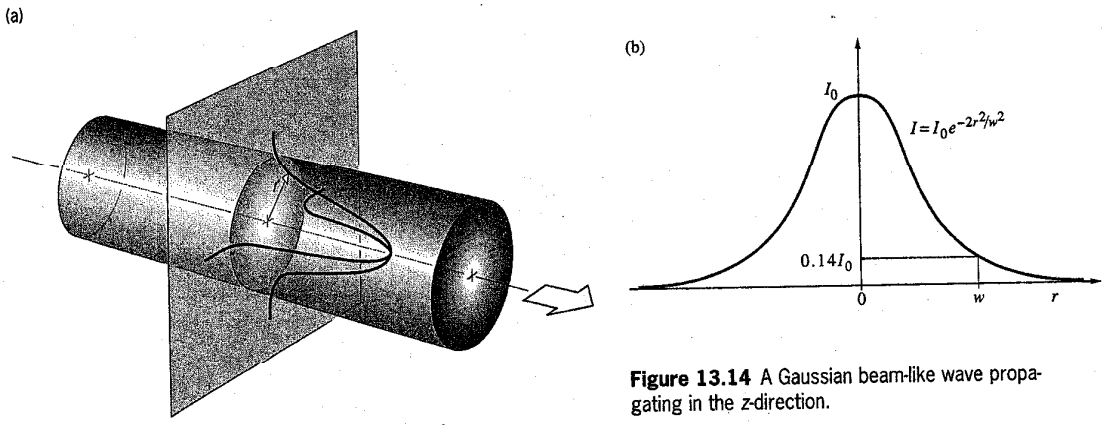

Figure 13.14 A Gaussian beam-like wave propagating in the z-direction.

between the mirrors of a laser resonator; its exact location depends on the specific design. For example, a confocal resonator (Fig. 13.12) has a waist halfway between the mirrors.

A more complete analysis of EM-waves in the cavity, setting $z = 0$ at the beam waist, yields the expression

$$w(z) = w_0\left[1 + \left(\frac{\lambda z}{\pi w_0^2}\right)^2\right]^{1/2} \tag{13.17}$$

for the half-width at any location z. The shape of the beam as specified by this expression for $w(z)$ is a hyperbola of revolution about the z-axis. A practical measure of the divergence of the beam is *the distance over which its cross-sectional area doubles*, or equivalently, the value of z for which $w(z) = \sqrt{2}w_0$. This special distance, z_R, is known as the **Rayleigh range**, and it follows from the above equation for $w(z)$ that

$$z_R = \frac{\pi w_0^2}{\lambda} \tag{13.18}$$

The smaller the waist (or equivalently, the smaller the minimum cross-sectional area), the smaller the Rayleigh range and the faster the beam diverges. At large distances from the waist ($z >> z_R$) the *full-angular width* of the beam (Θ, in radians) approaches $2w(z)/z$. In other words, as the line of length z, rotates through the angle Θ, its endpoint sweeps out a distance of $\approx 2w(z)$. Thus, when z is large and w_0 is small, the second term in the expression for $w(z)$ is much greater than 1 and

$$w(z) \approx w_0\left[\left(\frac{\lambda z}{\pi w_0^2}\right)^2\right]^{1/2} \approx \frac{\lambda z}{\pi w_0}$$

Since $\Theta \to 2w(z)/z$,

$$\Theta = \frac{2\lambda}{\pi w_0} = 0.637\,\frac{\lambda}{w_0}$$

Again, ***the smaller w_0 is, the larger will be*** Θ, the beam divergence. In part, that's why people used to use megaphones—waves emerging from a larger aperture diverge less.

While two plane mirrors forming a laser cavity will produce a beam that is aperture limited via diffraction, this will

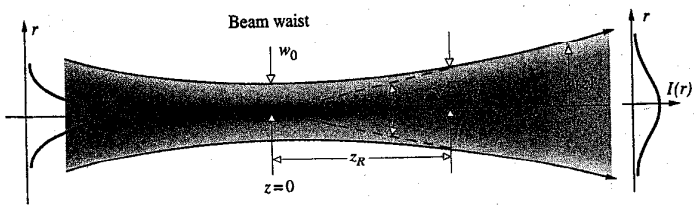

Figure 13.15 The spreading of a Gaussian beam.

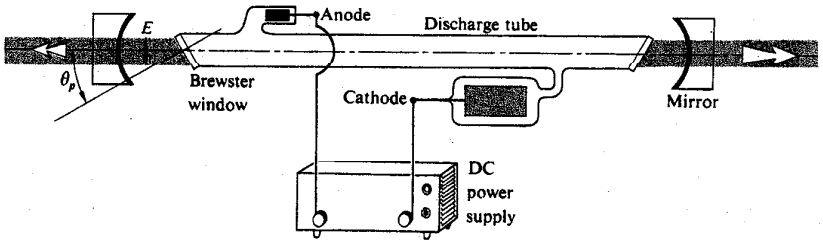

Figure 13.16 A simple, early He–Ne laser configuration.

not now be the case. Recall Eq. (10.58), which describes the radius of the Airy disk, and divide both sides by f to get the half-angular width of the diffracted circular beam of diameter D. Doubling this yields Θ, the *full-angular width* or **divergence of an aperture-limited laserbeam**:

$$\Theta \approx 2.44\lambda/D$$

By comparison, far from the region of minimum cross section, the full-angular width of a waisted laserbeam is

$$\Theta \approx 1.27\lambda/D_0 \qquad (13.19)$$

where $D_0 = 2w_0$, the beam-waist diameter, can be calculated from the particular cavity configuration.

The Helium–Neon Laser

Maiman's announcement of the first operative laser came at a New York news conference on July 7, 1960.* By February of 1961 Ali Javan and his associates W. R. Bennett, Jr., and D. R. Herriott had reported the successful operation of a *continuous-wave* (c-w) helium–neon, gas laser at 1152.3 nm. The He–Ne laser (Fig. 13.16) is still among the most popular devices of its kind, most often providing a few milliwatts of continuous power in the visible (632.8 nm). Its appeal arises primarily because it's easy to construct, relatively inexpensive, and fairly reliable and in most cases can be operated by a flick of a single switch. Pumping is usually accomplished by electrical discharge (via either dc, ac, or electrodeless rf excitation). Free electrons and ions are accelerated by an applied field and, as a result of collisions, cause further ionization and excitation of the gaseous medium (typically, a mixture of about 0.8 torr of He and about 0.1 torr of Ne). Many helium atoms, after dropping down from several upper levels, accumulate in the long-lived 2^1S- and 2^3S-states. These are metastable states (Fig. 13.17) from which there are no allowed radiative transitions. The excited He atoms inelastically collide with and transfer energy to ground-state Ne atoms, raising them in turn to the 5s- and 4s-states. These are the upper laser levels, and there then exists a population inversion with respect to the lower 4p- and 3p-states. Transitions between the 5s- and 4s-states are forbidden. Spontaneous photons initiate stimulated emission, and the chain reaction begins. The dominant laser transitions

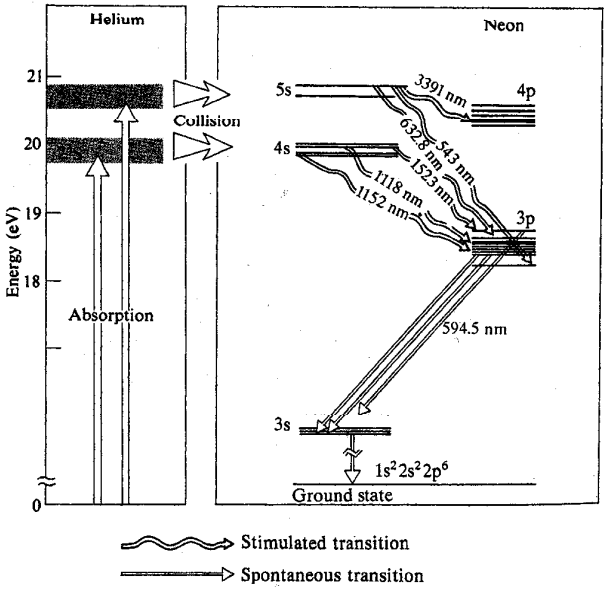

Figure 13.17 He–Ne laser energy levels.

*His initial paper, which would have made his findings known in a more traditional fashion, was rejected for publication by the editors of *Physical Review Letters*—this to their everlasting chagrin.

correspond to 1152.3 nm and 3391.2 nm in the infrared and, of course, the ever-popular 632.8 nm in the visible (bright red). The *p*-states drain off into the 3s-state, thus themselves remaining uncrowded and thereby continuously sustaining the inversion. The 3s-level is metastable, so that 3s-atoms return to the ground state after losing energy to the walls of the enclosure. This is why the plasma tube's diameter inversely affects the gain and is, accordingly, a significant design parameter. In contrast to the ruby, where the laser transition is down to the ground state, stimulated emission in the He–Ne laser occurs between two upper levels. The significance of this, for example, is that since the 3p-state is ordinarily only sparsely occupied, a population inversion is very easily obtained, and this without having to half empty the ground state.

Return to Fig. 13.16, which pictures the relevant features of a basic early He–Ne laser. The mirrors are coated with a multilayered dielectric film having a reflectance of over 99%. The laser output is made linearly polarized by the inclusion of Brewster end windows (i.e., plates tilted at the polarization angle) terminating the discharge tube. If these end faces were instead normal to the axis, reflection losses (4% at each interface) would become unbearable. By tilting them at the polarization angle, the windows presumably have 100% transmission for light whose electric-field component is parallel to the plane-of-incidence (the plane of the drawing). This polarization state rapidly becomes dominant, since the normal component is partially reflected off-axis at each transit of the windows. Linearly polarized light in the plane-of-incidence soon becomes the preponderant stimulating mechanism in the cavity, to the ultimate exclusion of the orthogonal polarization.*

Epoxying the windows to the ends of the laser tube and mounting the mirrors externally was a typical though dreadful approach used commercially until the mid-1970s. Inevitably, the epoxy leaked, allowing water vapor in and helium out. Today, such lasers are *hard sealed*; the glass is bonded directly to metal (Kovar) mounts, which support the mirrors within the tube. The mirrors (one of which is generally ≈100% reflective) have modern resistive coatings so they can tolerate the discharge environments within the tube. Operating lifetimes of 20 000 hours and more are now the rule (up from only a few hundred hours in the 1960s). Brewster windows are usually optional, and most commercial He–Ne lasers generate more or less "unpolarized" beams. The typical mass-produced He–Ne laser (with an output of from 0.5 mW to 5 mW) operates in the TEM_{00} mode, has a coherence length of around 25 cm, a beam diameter of approximately 1 mm, and a low overall efficiency of only 0.01% to about 0.1%. Although there are infrared He–Ne lasers, and even a new green (543.5 nm) He–Ne laser, the bright red 632.8-nm version remains the most popular.

A Survey of Laser Developments

Laser technology is so dynamic a field that what was a laboratory breakthrough a year or two ago may be a commonplace off-the-shelf item today. The whirlwind will certainly not pause to allow descriptive terms like "the smallest," "the largest," "the most powerful," and so on to be applicable for very long. With this in mind, we briefly survey the existing scene without trying to anticipate the wonders that will surely come after this type is set. Laserbeams have already been bounced off the Moon; they have spot welded detached retinas, generated fusion neutrons, stimulated seed growth, served as communications links, read CD discs, guided milling machines, missiles, ships, and grating engines, carried color television pictures, drilled holes in diamonds, levitated tiny objects,* and intrigued countless among the curious.

Along with ruby there are a great many other **solid-state lasers** whose outputs range in wavelength from roughly 170 nm to 3900 nm. For example, the trivalent rare earths Nd^{3+}, Ho^{3+}, Gd^{3+}, Tm^{3+}, Er^{3+}, Pr^{3+}, and Eu^{3+} undergo laser action in a host of hosts, such as $CaWO_4$, Y_2O_3, $SrMoO_4$, LaF_3, yttrium aluminum garnet (YAG for short), and glass, to name only a few. Of these, neodymium-doped glass and neodymium-doped YAG are of particular importance. Both constitute high-powered laser media operating at approximately 1060 nm. Nd: YAG lasers generating in excess of a kilowatt of continuous power have been constructed. Tremendous power outputs in pulsed systems have been obtained by operating several lasers in tandem. The first laser in the train serves as a

*Half of the output power of the laser is *not* lost in reflections at the Brewster windows when the transverse $\mathcal{P}$-state light is scattered. Energy simply isn't continuously channeled into that polarization component by the cavity. If it's reflected out of the plasma tube, it's not present to stimulate further emission.

*See M. Lubin and A. Fraas, "Fusion by Laser," *Sci. Am.* 224, **21** (June 1971); R. S. Craxton, R. L. McCrory, and J. M. Soures, "Progress in Laser Fusion," *Sci. Am.* 255, **69** (August 1986); and A. Ashkin, "The Pressure of Laser Light," *Sci. Am.* 226, **63** (February 1972).

Q-switched oscillator that fires into the next stage, which functions as an amplifier; and there may be one or more such amplifiers in the system. By reducing the feedback of the cavity, a laser will no longer be self-oscillatory, but it will amplify an incident wave that has triggered stimulated emission. Thus the amplifier is, in effect, an active medium, which is pumped, but for which the end faces are only partially reflecting or even nonreflecting. Ruby systems of this kind, delivering a few GW (gigawatts, i.e., 10^9 W) in the form of pulses lasting several nanoseconds, are available commercially.

On December 19, 1984, the largest laser then in existence, the Nova at the Lawrence Livermore National Laboratory in California, fired all 10 of its beams at once for the first time, producing a warm-up shot of a mere 18 kJ of 350-nm radiation in a 1-ns pulse (see photo). This immense neodymium-doped glass laser was designed to focus up to 120 TW onto a fusion pellet—that's roughly 500 times more power than all the electrical generating stations in the United States—albeit only for about 10^{-9} s. In the late 1900s, the last years of its operation, using just one beamline of the Nova, LLNL researchers were able to produce 1.25-PW pulses, each lasting 490 fs and carrying 580 J.

Nova's successor, which came on line in 1980 (with 24 solid state lasers), is housed in the Laboratory for Laser Energetics at the (LLE) University of Rochester. At present, LLE operates the 30 to 45 kJ Omega laser and is the world's premier laser fusion research facility. Upgraded in 1995, Omega is a 60-beam neodymium-doped phosphate glass laser that can concentrate 60×10^{12} W of radiant power onto a pinhead sized target. To accomplish that feat, the initial laser output is split repeatedly and each beam is subsequently amplified using Nd:glass discs and rods. Just before reaching the target the several beams are frequency-tripled to 351 nm using KDP crystals (p. 641). Supplying the needs of a variety of researchers, Omega is being operated at its maximum rate of one shoot per hour.

Still under construction, Omega's successor will be housed in the Department of Energy's immense National Ignition Facility. Its 192 beams will deliver 50 times more energy to its targets than does Omega.

A large group of **gas lasers** operate across the spectrum from the far IR to the UV (1 mm to 150 nm). Primary among these are helium–neon, argon, and krypton, as well as several molecular gas systems, such as carbon dioxide, hydrogen fluoride, and molecular nitrogen (N_2). Argon lases mainly in the green, blue-green, and violet (predominantly at 488.0 and 514.5 nm) in either pulsed or continuous operation. Although

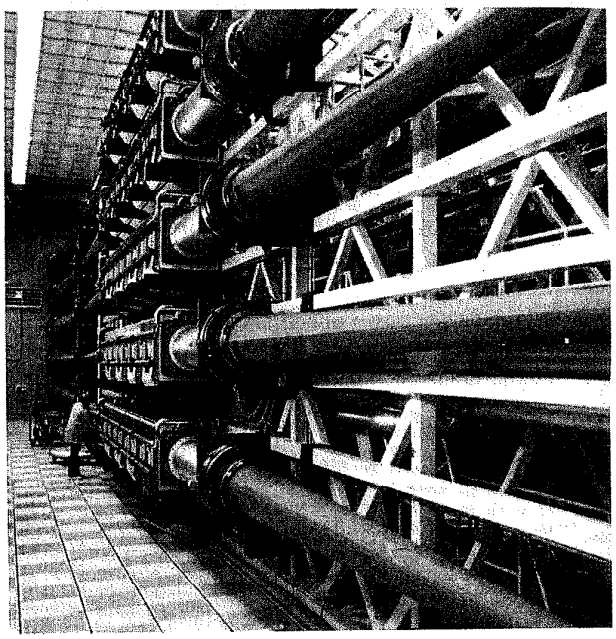

The Nova laser. (Photo courtesy Lawrence Livermore National Laboratory.)

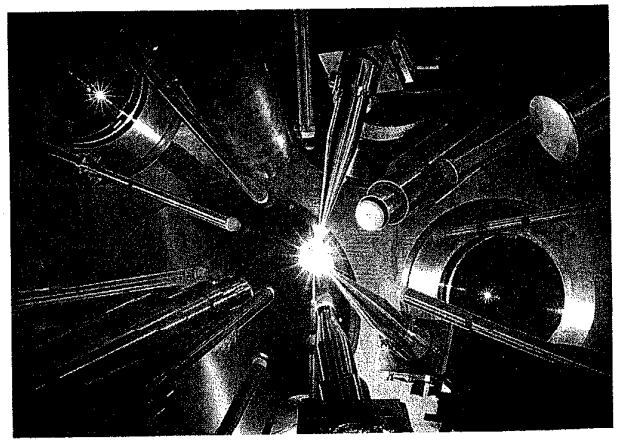

Inside the target chamber of the LLE laser-fusion device. Fusion reactions take place in a tiny target-sphere filled with deuterium/tritium and irradiated by the 30-kJ Omega laser. (Photo courtesy University of Rochester's Laboratory for Laser Energetics, Eugene Kowaluk, Image specialist.)

its output is usually several watts c-w, it has gone as high as 150 W c-w. The argon ion laser is similar in some respects to the He–Ne laser, although it evidently differs in its usually greater power, shorter wavelength, broader linewidth, and higher price. All of the noble gases (He, Ne, Ar, Kr, Xe) have been made to lase individually, as have the gaseous ions of many other elements, but the former grouping has been studied most extensively.

The CO_2 molecule, which lases between vibrational modes, emits in the IR at 10.6 μm, with typical c-w power levels of from a few watts to several kilowatts. Its efficiency can be an unusually high 15% when aided by additions of N_2 and He. While it once took a discharge tube nearly 200 m long to generate 10 kW c-w, considerably smaller "table models" are now available commercially. For a while in the 1970s, the record output belonged to an experimental gas-dynamic laser utilizing thermal pumping on a mixture of CO_2, N_2, and H_2O to generate 60 kW c-w at 10.6 μm in multimode operation.

The pulsed nitrogen laser operates at 337.1 nm in the UV, as does the c-w helium–cadmium laser. A number of metal vapors (e.g., Zn, Hg, Sn, Pb) have displayed laser transitions in the visible, but problems such as maintaining uniformity of the vapor in the discharge region have handicapped their exploitation. The He–Cd laser emits at 325.0 nm and 441.6 nm. These are transitions of the cadmium ion arising after excitation resulting from collisions with metastable helium atoms.

The **semiconductor laser**—alternatively known as the junction or diode laser—was invented in 1962, soon after the development of the light-emitting diode (LED). Today it serves a central role in electro-optics, primarily because of its spectral purity, high efficiency ($\approx$100%), ruggedness, ability to be modulated at extremely rapid rates, long lifetimes, and moderate power (as much as 200 mW) despite its pinhead size. Junction lasers have already been used in the millions in fiberoptic communications, laser disk audio systems, and so forth.

The first such lasers were made of one material, gallium arsenide, appropriately doped to form a *p–n* junction. The associated high lasing threshold of these so-called homostructures limited them to pulsed mode operation and cryogenic temperatures; otherwise the heat developed in their small structures would destroy them. The first tunable lead–salt diode laser was developed in 1964, but it was not until almost a dozen years later that it became commercially available. It operates at liquid nitrogen temperatures, which is certainly inconvenient, but it can scan from 2 μm to 30 μm.

Later advances have since allowed a reduction in the threshold and resulted in the advent of the continuous-wave (c-w), room temperature diode laser. Transitions occur between the conduction and valence bands, and stimulated emission results in the immediate vicinity of the *p–n* junction (Fig. 13.18). Quite generally, as a current flows in the forward direction through a semiconductor diode, electrons from the *n*-layer conduction band will recombine with *p*-layer holes, thereupon emitting energy in the form of photons. This radiative process, which competes for energy with the existing absorption mechanisms (such as phonon production), comes to predominate when the recombination layer is small and the current is large. To make the system lase, the light emitted

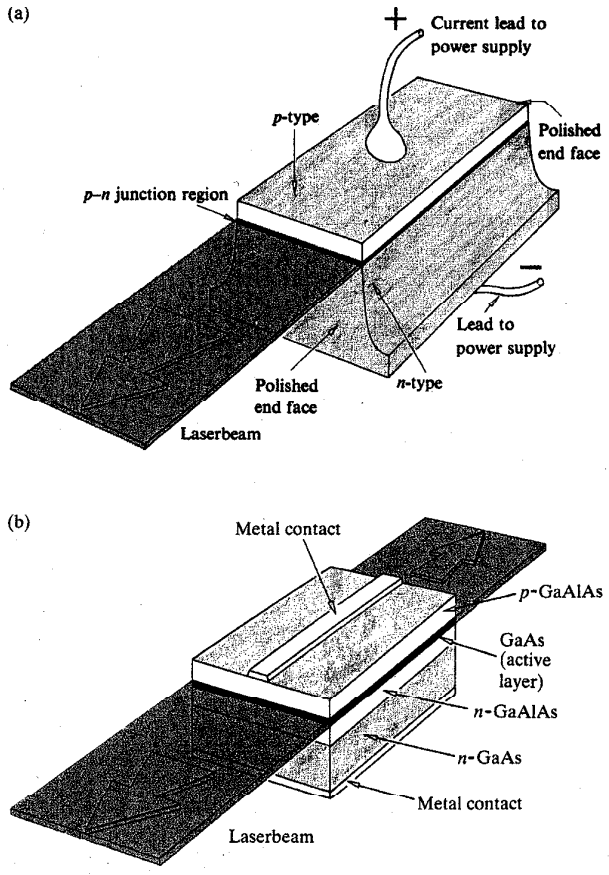

Figure 13.18 (a) An early GaAs *p–n* junction laser. (b) A more modern diode laser.

from the diode is retained within a resonant cavity, and that's usually accomplished by simply polishing the end faces perpendicular to the junction channel.

Nowadays semiconductor lasers are created to meet specific needs, and there are many designs producing wavelengths ranging from around 700 nm to about 30 μm. The early 1970s saw the introduction of the c-w GaAs/GaAlAs laser. Operating at room temperature in the 750-nm to 900-nm region (depending on the relative amounts of aluminum and gallium), the tiny diode chip is usually about a sixteenth of a cubic centimeter in volume. Figure 13.18*b* shows a typical heterostructure (a device formed of different materials) diode laser of this kind. Here the beam emerges in two directions from the 0.2-μm thick active layer of GaAs. These little lasers usually produce upward of 20 mW of continuous wave power. To take advantage of the low-loss region ($\lambda \approx 1.3$ μm) in fiberoptic glass the GaInAsP/InP laser was devised in the mid-1970s with an output of 1.2 μm to 1.6 μm.

The cleaved–coupled-cavity laser is shown in the accompanying photo. In it the number of axial modes is controlled in order to produce very-narrow-bandwidth tunable radiation. Two cavities coupled together across a small gap restrict the radiation to the extremely narrow bandwidth that can be sustained in both resonant chambers.*

The first **liquid laser** was operated in January of 1963.† All of the early devices of this sort were exclusively *chelates* (i.e., metallo-organic compounds formed of a metal ion with organic radicals). That original liquid laser contained an alcohol solution of europium benzoylacetonate emitting at 613.1 nm. The discovery of laser action in nonchelate organic liquids was made in 1966. It came with the fortuitous lasing (at 755.5 nm) of a chloroaluminum phthalocyanine solution during a search for stimulated Raman emission in that substance.‡

A great many fluorescent dye solutions of such families as the fluoresceins, coumarins, and rhodamines have since been made to lase at frequencies from the IR into the UV. These have usually been pulsed, although c-w operation has been obtained. There are so many organic dyes that it would seem possible to build such a laser at any frequency in the visible. Moreover, these devices are distinctive in that they inherently can be tuned continuously over a range of wavelengths (of perhaps 70 nm or so, although a pulsed system tunable over 170 nm exists). Indeed, there are other arrangements that will vary the frequency of a primary laserbeam (i.e., the beam enters with one color and emerges with another, Section 13.4), but in the case of the dye laser, the primary beam itself is tuned internally. This is accomplished, for example, by changing the concentration or the length of the dye cell or by adjusting a diffraction grating reflector at the end of the cavity. Multicolor dye laser systems, which can easily be switched from one dye to another and thereby operate over a very broad frequency range, are available commercially.

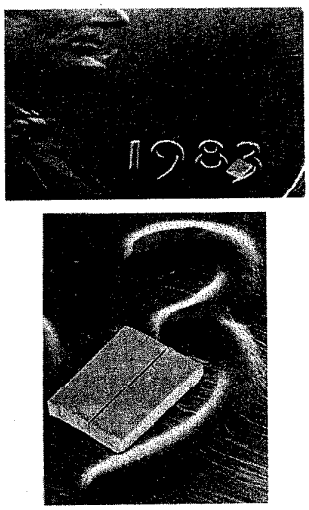

The cleaved–coupled-cavity laser. (Photos courtesy of Bell Laboratories.)

A **chemical laser** is one that is pumped with energy released via a chemical reaction. The first of this kind was operated in 1964, but it was not until 1969 that a continuous-wave chemical laser was developed. One of the most promising of these is the deuterium fluoride–carbon dioxide ($DF–CO_2$) laser. It is self-sustaining in that it requires no external power source. In brief, the reaction $F_2 + D_2 \rightarrow 2DF$, which occurs on the mixing of these two fairly common gases, generates enough energy to pump a CO_2 laser.

There are solid-state, gaseous, liquid, and vapor (e.g., H_2O) lasers; there are semiconductor lasers, free electron (600 nm to

*See Y. Suematsu, "Advances in Semiconductor Lasers," *Phys. Today*, **32** (May 1985). For a discussion of heterostructure diode lasers, refer to M. B. Panish and I. Hayashi, "A New Class of Diode Lasers," *Sci. Am.* 225, **32** (July 1971).

†See Adam Heller, "Laser Action in Liquids," *Phys. Today* (November 1967), p. 35, for a more detailed account.

‡P. Sorokin, "Organic Lasers," *Sci. Amer.* 220, **30** (February 1969).

3 mm) lasers, X-ray lasers, and lasers with very special properties, such as those that generate extremely short pulses, or those that have extraordinary frequency stability. These latter devices are very useful in the field of high-resolution spectroscopy, but there is a growing need for them in other research areas as well (e.g., in the interferometers used to attempt to detect gravity waves). In any event, these lasers must have precisely controlled cavity configurations despite the disturbing influences of temperature variations, vibrations, and even sound waves. To date, the record is held by a laser at the Joint Institute for Laboratory Astrophysics in Boulder, Colorado, which maintains a frequency stability (p. 315) of nearly one part in 10^{14}.

13.1.4 The Light Fantastic

Laserbeams differ somewhat from one type of laser to another; yet there are several remarkable features that are displayed, to varying degrees, by all of them. Quite apparent is the fact that most laserbeams are exceedingly directional, or if you will, highly collimated. One need only blow some smoke into the otherwise invisible, visible laserbeam to see (via scattering) a fantastic thread of light stretched across a room. A He–Ne beam in the TEM_{00} mode generally has a divergence of only about one minute of arc or less. Recall that in that mode the emission closely approximates a Gaussian irradiance distribution; that is, the flux density drops off from a maximum at the central axis of the beam and has no side lobes. The typical laserbeam is quite narrow, usually issuing at no more than a few millimeters in diameter. Since the beam resembles a truncated plane wave, it is of course highly *spatially coherent*. In fact, its directionality may be thought of as a manifestation of that coherence. Laserlight is quasimonochromatic, generally having an exceedingly narrow frequency bandwidth (p. 313). In other words, it is highly *temporally coherent*.

Another attribute is the large flux or *radiant power* that can be delivered in that narrow frequency band. As we've seen, the laser is distinctive in that it emits all its energy in the form of a narrow beam. In contrast, a 100-W incandescent lightbulb may pour out considerably more radiant energy in toto than a lower-power c-w laser, but the emission is incoherent, spread over a large solid angle, and it has a broad bandwidth as well. A good lens* can totally intercept a laserbeam and focus essentially all of its energy into a minute spot (whose diameter varies directly with λ and the focal length and inversely with the beam diameter). Spot diameters of just a few thousandths of an inch can readily be attained with lenses that have a conveniently short focal length. And a spot diameter of a few hundred-millionths of an inch is possible in principle. Thus flux densities can readily be generated in a focused laserbeam of over 10^{17} W/cm^2, in contrast to, say, an oxyacetylene flame having roughly 10^3 W/cm^2. To get a better feel for these power levels, note that a focused CO_2 laserbeam of a few kilowatts c-w can burn a hole through a quarter-inch stainless steel plate in about 10 seconds. By comparison, a pinhole and filter positioned in front of an ordinary source will certainly produce spatially and temporally coherent light, but only at a minute fraction of the total power output.

*Spherical aberration is usually the main problem, since laserbeams are, as a rule, both quasimonochromatic and incident along the axis of the lens.

Femtosecond Optical Pulses

The advent of the mode-locked dye laser in the early part of the 1970s gave a great boost to the efforts then being made at generating extremely short pulses of light.* Indeed, by 1974 subpicosecond (1 ps = 10^{-12} s) optical pulses were already being produced, although the remainder of the decade saw little significant progress. In 1981 two separate advances resulted in the creation of femtosecond laser pulses (i.e., <0.1 ps or <100 fs)—a group at Bell Labs developed a colliding-pulse ring dye laser, and a team at IBM devised a new pulse-compression scheme.

Above and beyond the implications in the practical domain of electro-optical communications, these accomplishments have firmly established a new field of research known as *ultrafast phenomena*. The most effective way to study the progression of a process that occurs exceedingly rapidly (e.g., carrier dynamics in semiconductors, fluorescence, photochemical biological processes, and molecular configuration changes) is to examine it on a time scale that is comparatively short with respect to what's happening. Pulses lasting ≈ 10 fs allow an entirely new access into previously obscure areas in the study of matter.

Pulses lasting a mere 8 fs (10^{-15} s), which corresponds to wavetrains only about 4 wavelengths of red light in length, can now be produced routinely. One of the techniques that makes these femtosecond wavegroups possible is based on an idea

*See Chandrashekhar Joshi and Paul Corkum, "Interactions of Ultra-Intense Laser Light with Matter," *Phys. Today* **36** (January 1995).

used in radar work in the 1950s called *pulse compression*. Here an initial laser pulse has its frequency spectrum broadened, thereby allowing the inverse or temporal pulse width to be shortened—remember that $\Delta\nu$ and Δt are conjugate Fourier quantities [Eq. (7.63)]. The input pulse (several picoseconds long) is passed into a nonlinear dispersive medium, namely, a single-mode optical fiber. When the light intensity is high enough, the index of refraction has an appreciable nonlinear term (Section 13.4), and the carrier frequency of the pulse experiences a time-dependent shift. On traversing perhaps 30 m of fiber, the frequency of the pulse is drawn out or "chirped." That is, a spread occurs in the spectrum of the pulse, with the low frequencies leading and the high frequencies trailing. Next the spectrally broadened pulse is passed through another dispersive system (a delay line), such as a pair of diffraction gratings. By traveling different paths, the blue-shifted trailing edge of the pulse is made to catch up to the red-shifted leading edge, creating a time-compressed output pulse.

The Speckle Effect

A rather striking and easily observable manifestation of the spatial coherence of laserlight is its granular appearance on reflection from a diffuse surface. Using a He–Ne laser (632.8 nm), expand the beam a bit by passing it through a simple lens and project it onto a wall or a piece of paper. The illuminated disk appears speckled with bright and dark regions that sparkle and shimmer in a dazzling psychedelic dance. Squint and the grains grow in size; step toward the screen and they shrink; take off your eyeglasses and the pattern stays in perfect focus. In fact, if you are nearsighted, the diffraction fringes caused by dust on the lens blur out and disappear, but the speckles do not. Hold a pencil at varying distances from your eye so that the disk appears just above it. At each position, focus on the pencil; wherever you focus, the granular display is crystal clear. Indeed, look at the pattern through a telescope; as you adjust the scope from one extreme to the other, the ubiquitous granules remain perfectly distinct, even though the wall is completely blurred.

The spatially coherent light scattered from a diffuse surface fills the surrounding region with a *stationary* interference pattern (just as in the case of the wavefront-splitting arrangements of Section 9.3). At the surface the granules are exceedingly small, and they increase in size with distance. At any location in space, the resultant field is the superposition of many contributing scattered wavelets. These must have a constant relative phase determined by the optical path length from each scatterer to the point in question, if the interference pattern is to be sustained. The accompanying photo illustrates this point rather nicely. It shows a cement block illuminated in one case by laserlight and in the other by collimated light from a Hg arc lamp, both of about the same spatial coherence. While the laser's coherence length is much greater than the height of the surface features, the coherence length of the Hg light is not. In the former case, the speckles in the photograph are large, and they obscure the surface structure; in the latter, despite its spatial coherence, the speckle pattern is not observable in the photograph, and the surface features predominate. Because of the rough texture, the optical path length difference between two wavelets arriving at a point in space, scattered from different surface bumps, is generally greater than the coherence length of the mercury light. This means that the relative phases of the overlapping wavetrains change rapidly and randomly in time, washing out the large-scale interference pattern.

A real system of fringes is formed of the scattered waves that converge in front of the screen. The fringes can be viewed by intersecting the interference pattern with a sheet of paper at a convenient location. After forming the real image in space, the rays proceed to diverge, and any region of the image can therefore be viewed directly with the eye appropriately focused. In contrast, rays that initially diverge appear to the eye as if they had originated behind the scattering screen and thus form a virtual image.

It seems that as a result of chromatic aberration, normal and farsighted eyes tend to focus red light behind the screen. Contrarily, a nearsighted person observes the real field in

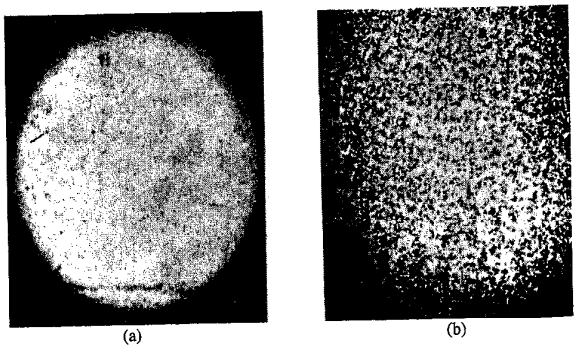

Speckle patterns. (*a*) A cement block illuminated by a mercury arc and (*b*) a He–Ne laser. [From B. J. Thompson, *J. Soc. Phot. Inst. Engr.* **4**, 7 (1965).]

front of the screen (regardless of wavelength). Thus if the viewer moves her head to the right, the pattern will move to the right in the first instance (where the focus is beyond the screen) and to the left in the second (focus in front). The pattern will follow the motion of your head, if you're viewing it very close to the surface. The same apparent parallax motion can be seen by looking through a window; outside objects will seem to move with your head, inside ones opposite to it. The brilliant, narrow-bandwidth, spatially coherent laserbeam is ideally suited for observing the granular effect, although other means are certainly possible.* In unfiltered sunlight the grains are minute, on the surface, and multicolored. The effect is easy to observe on a smooth, flat-black material (e.g., poster-painted paper), but you can see it on a fingernail or a worn coin as well.

Although it provides a marvelous demonstration, both aesthetically and pedagogically, the granular effect can be a real practical nuisance in coherently illuminated systems. For example, in holographic imagery the speckle pattern corresponds to troublesome background noise. Incidentally, very much the same kind of thing is observable when listening to a mobile radio where the signal strength fluctuates from one location to the next, depending on the environment and the resulting interference pattern.

The Spontaneous Raman Effect

It is possible that an excited atom will not return to its initial state after the emission of a photon. This kind of behavior had been observed and studied extensively by George Stokes prior to the advent of quantum theory. Since the atom drops down to an interim state, it emits a photon of lower energy than the incident primary photon, in what is usually referred to as a ***Stokes transition***. If the process takes place rapidly (roughly 10^{-7} s), it is called **fluorescence**, whereas if there is an appreciable delay (in some cases seconds, minutes, or even many hours), it is known as **phosphorescence**. Using ultraviolet quanta to generate a fluorescent emission of visible light has become an accepted occurrence in our everyday lives. Any number of commonplace materials (e.g., detergents, organic dyes, and tooth enamel) will emit characteristic visible photons so that they appear to glow under ultraviolet illumination; ergo the widespread use of the phenomenon for commercial display purposes and for "whitening" cloths.

If quasimonochromatic light is scattered from a substance, it will thereafter consist mainly of light of the same frequency. Yet it is possible to observe very weak additional components having higher and lower frequencies (side bands). Moreover, the difference between the side bands and the incident frequency ν_i is found to be characteristic of the material and therefore suggests an application to spectroscopy. The **Spontaneous Raman Effect**, as it is now called, was predicted in 1923 by Adolf Smekal and observed experimentally in 1928 by Sir Chandrasekhara Vankata Raman (1888–1970), then professor of physics at the University of Calcutta. The effect was difficult to put to actual use, because one needed strong sources (usually Hg discharges were used) and large samples. Often the ultraviolet from the source would further complicate matters by decomposing the specimen. And so it is not surprising that little sustained interest was aroused by the promising practical aspects of the Raman Effect. The situation was changed dramatically when the laser became a reality. ***Raman spectroscopy*** is now a unique and powerful analytical tool.

To appreciate how the phenomenon operates, let's review the germane features of molecular spectra. A molecule can absorb radiant energy in the far-infrared and microwave regions, converting it to rotational kinetic energy. Furthermore, it can absorb infrared photons (i.e., ones within a wavelength range from roughly 10^{-2} mm down to about 700 nm), transforming that energy into vibrational motion of the molecule. Finally, a molecule can absorb energy in the visible and ultraviolet regions through the mechanism of electron transitions, much like those of an atom. Suppose then that we have a molecule in some vibrational state, which, using quantum-mechanical notation, we call $|b\rangle$, as indicated diagrammatically in Fig. 13.19*a*. This need not necessarily be an excited state. An incident photon of energy $h\nu_i$ is absorbed, raising the system to some intermediate or virtual state, whereupon it immediately makes a Stokes transition, emitting a (scattered) photon of energy $h\nu_s < h\nu_i$. In conserving energy, the difference $h\nu_i - h\nu_s = h\nu_{cb}$ goes into exciting the molecule to a higher vibrational energy level $|c\rangle$. It is possible that electronic or rotational excitation results as well.

Alternatively, if the initial state is an excited one (just heat the sample), the molecule, after absorbing and emitting a photon, may drop back to an even lower state (Fig. 13.19*b*), thereby making an ***anti-Stokes transition***. In this instance $h\nu_s > h\nu_i$, which means that some vibrational energy of the molecule

*For further reading on this effect, see L. I. Goldfischer, *J. Opt. Soc. Am.* **55**, 247 (1965); D. C. Sinclair, *J. Opt. Soc. Am.* **55**, 575 (1965); J. D. Rigden and E. I. Gordon, *Proc. IRE* **50**, 2367 (1962); B. M. Oliver, *Proc. IEEE* **51**, 220 (1963).

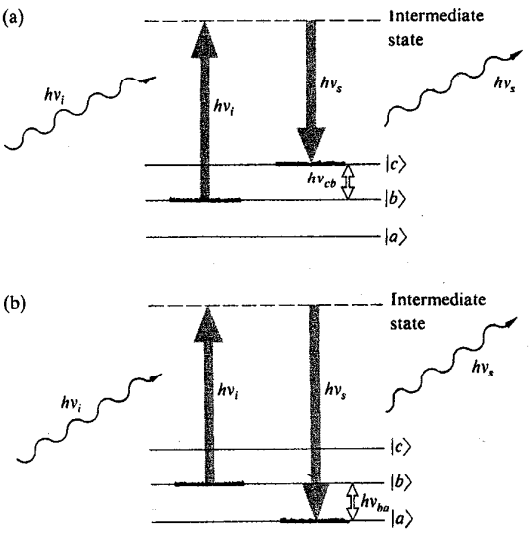

Figure 13.19 Spontaneous Raman Scattering.

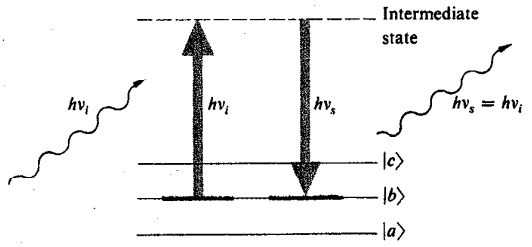

Figure 13.20 Rayleigh Scattering.

$(h\nu_{ba} = h\nu_s - h\nu_i)$ has been converted into radiant energy. In either case, the resulting differences between ν_s and ν_i correspond to specific energy-level differences for the substance under study and as such yield insights into its molecular structure. Figure 13.20, for comparison's sake, depicts Rayleigh Scattering where $\nu_s = \nu_i$.

The laser is an ideal source for spontaneous Raman Scattering. It is bright, quasimonochromatic, and available in a wide range of frequencies. Figure 13.21 illustrates a typical laser-Raman system. Complete research instruments of this sort are commercially available, including the laser (usually helium–neon, argon, or krypton), focusing lens systems, and photon-counting electronics. The double scanning monochromator provides the needed discrimination between ν_i and ν_s, since unshifted laserlight (ν_i) is scattered along with the

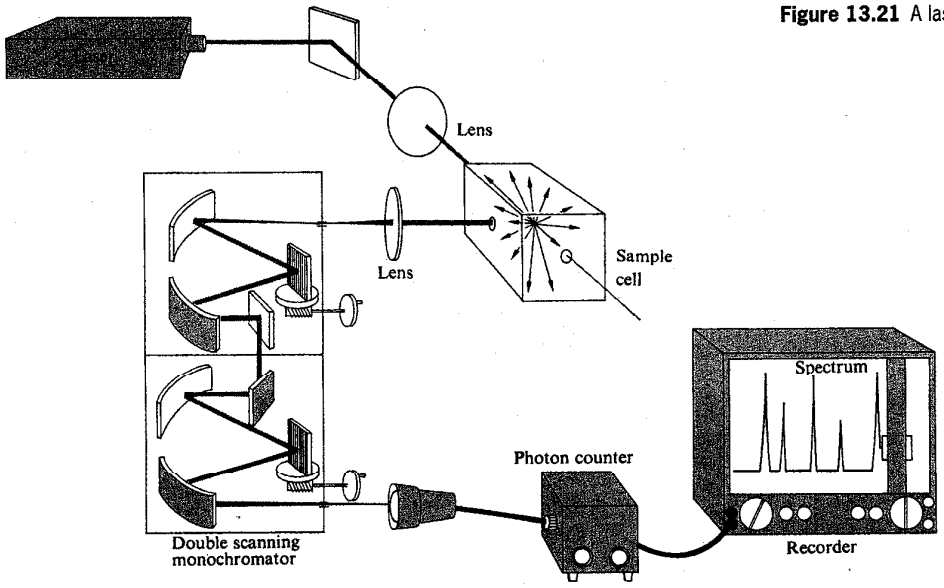

Figure 13.21 A laser–Raman system.

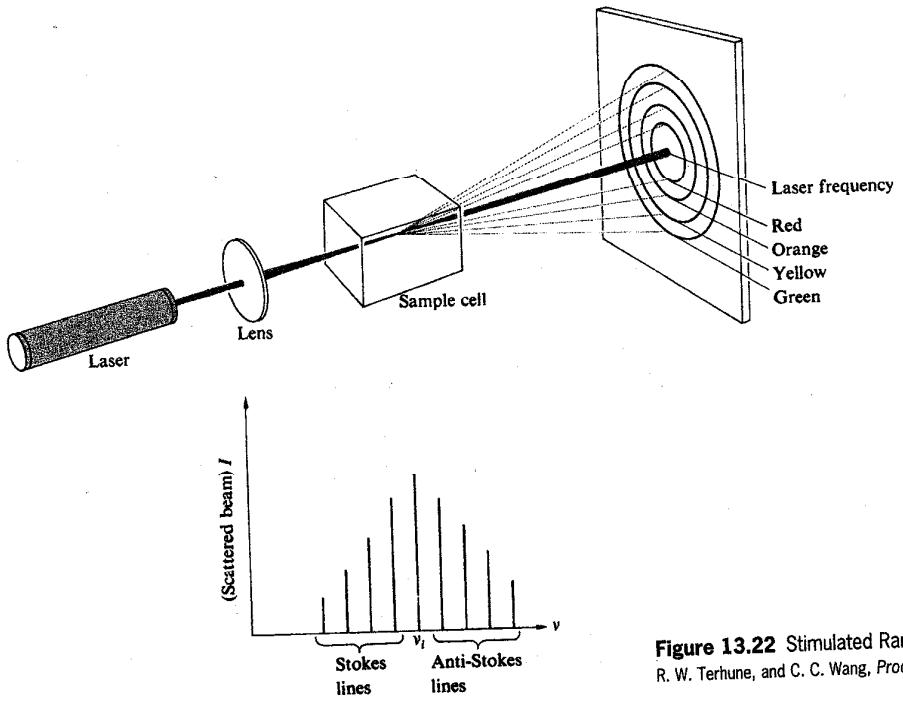

Figure 13.22 Stimulated Raman Scattering. [See R. W. Minck, R. W. Terhune, and C. C. Wang, *Proc. IEEE* **54**, 1357 (1966).]

Raman spectra (ν_s). Although Raman Scattering associated with molecular rotation was observed prior to the use of the laser, the increased sensitivity now available makes the process easier and allows even the effects of electron motion to be examined.

The Stimulated Raman Effect

In 1962 Eric J. Woodbury and Won K. Ng rather fortuitously discovered a remarkable related effect known as *Stimulated Raman Scattering*. They had been working with a million-watt pulsed ruby laser incorporating a nitrobenzene Kerr cell shutter (see Section 8.11.3). They found that about 10% of the incident energy at 694.3 nm was shifted in wavelength and appeared as a *coherent* scattered beam at 766.0 nm. It was subsequently determined that the corresponding frequency shift of about 40 THz was characteristic of one of the vibrational modes of the nitrobenzene molecule, as were other new frequencies also present in the scattered beam. Stimulated Raman Scattering can occur in solids, liquids, or dense gases under the influence of focused high-energy laser pulses (Fig. 13.22). The effect is schematically depicted in Fig. 13.23. Here two photon beams are simultaneously incident on a molecule, one corresponding to the laser frequency ν_i, the other having the scattered frequency ν_s. In the original setup, the scattered beam was reflected back and forth through the specimen, but the effect can occur without a resonator. The laserbeam loses a photon $h\nu_i$, while the scattered beam gains a photon $h\nu_s$ and

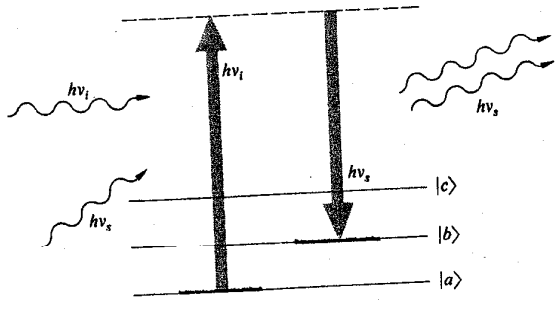

Figure 13.23 Energy-level diagram of Stimulated Raman Scattering.

is subsequently *amplified*. The remaining energy ($h\nu_i - h\nu_s = h\nu_{ba}$) is transmitted to the sample. The chain reaction in which a large portion of the incident beam is converted into stimulated Raman light can only occur above a certain high-threshold flux density of the exciting laserbeam.

Stimulated Raman Scattering provides a whole new range of high flux-density coherent sources extending from the infrared to the ultraviolet. It should be mentioned that in principle each spontaneous scattering mechanism (e.g., Rayleigh and Brillouin Scattering) has its stimulated counterpart.*

13.2 Imagery — The Spatial Distribution of Optical Information

The manipulation of all sorts of data via optical techniques has already become a technological *fait accompli*. The literature since the 1960s reflects, in a diversity of areas, this far-reaching interest in the methodology of optical data processing. Practical applications have been made in the fields of television and photographic image enhancement, radar and sonar signal processing (phased and synthetic array antenna analysis), as well as in pattern recognition (e.g., aerial photointerpretation and fingerprint studies), to list only a very few.

Our concern here is to develop the nomenclature and some of the ideas necessary for an appreciation of this contemporary thrust in Optics.

13.2.1 Spatial Frequencies

In electrical processes one is most frequently concerned with signal variations in time, that is, the moment-by-moment alteration in voltage that might appear across a pair of terminals at some fixed location in space. By comparison, in Optics we are most often concerned with information spread across a region of space at a fixed location in time. For example, we

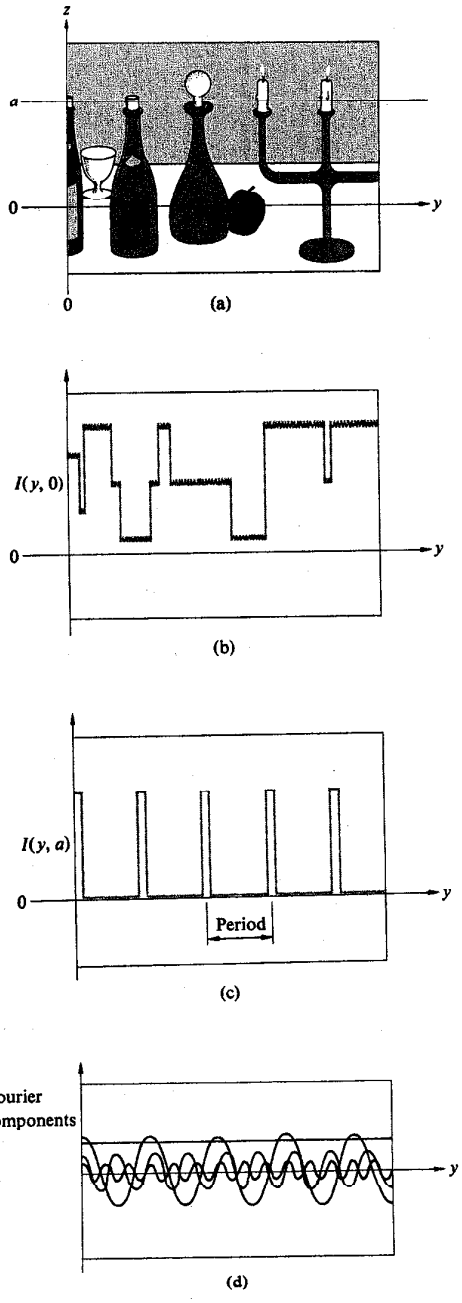

Figure 13.24 A two-dimensional irradiance distribution.

*For further reading on these subjects you might try the review-tutorial paper by Nicolaas Bloembergen, "The Stimulated Raman Effect," *Am. J. Phys.* **35**, 989 (1967). It contains a fairly good bibliography as well as a historical appendix. Many of the papers in *Lasers and Light* also deal with this material and are highly recommended reading.