LASERS

1. Introduction

Once dubbed as "a solution without a problem", lasers have found numerous practical uses ranging from forensic science to home CD and DVD players. An enormous number of materials have been discovered that will produce lasing action including gases, liquids, and solids covering a range of wavelengths spanning from extreme-ultraviolet (uv) to far-infrared (ir).

By examining the structure of a typical laser the purpose of each element shall become evident. For a working example the helium–neon (HeNe) laser is used. The most common gas laser, the HeNe laser, was one of the first lasers, discovered in the early 1960s. Its red output has found uses in everything from holography to bar code scanning. A few years ago it was the most common laser, but its dominance has been replaced by cheaper semiconductor lasers. Still, it makes an excellent example of how a typical laser operates.

The basic laser itself consists of a glass tube filled with helium and neon gases in a ratio of ∼10 parts helium to 1 part neon. The internal pressure is low (∼1.8 Torr, where 760 Torr is 1 atm), allowing a sustained electrical discharge. The electrical energy required is supplied by a high voltage power supply, which is often encapsulated in a small block of epoxy material. When energized, the tube glows a bright pink color. Most common HeNe lasers have a red output at 632.8 nm; however, the quantum mechanics of the neon atom (which is the active lasing species) also allows lasing transitions in the orange, yellow, and green. The discharge itself takes place between a small anode and a much larger cathode through a tiny capillary tube called a *plasma tube*. At either end of the laser are cavity mirrors. One mirror is fully reflecting, the other is partially reflecting. The small portion of light (typically ∼1%) that is transmitted through the front mirror is the actual laser beam itself. The structure of a typical HeNe laser is depicted in Figure 1.

2. Theoretical Aspects

2.1. Laser Light. Before examining the mechanism of the laser the nature of laser light itself is briefly discussed. The vast majority of light sources produce incoherent light such as the light normally produced by gas discharges and blackbody sources. The basic quantum mechanics behind these light sources is the same as that governing the operation of lasers, but laser light is unlike any other source found in nature. It has three special properties that lead to its usefulness in many applications: coherence, monochromaticity, and collimation (directionality).

Coherence is the most interesting property of laser light. This property states that all photons emitted from a laser are at exactly the same phase; as waves they all "crest" and "valley" at the same time. Coherence is brought about by the mechanism of the laser itself (stimulated emission), in which photons are essentially copied. To stay in phase, it is required that all emitted photons are at exactly the same wavelength (or extremely close). If some photons are at a different wavelength, the phase of those photons relative to others will

be different and the light will not be coherent. They must also be highly directional, all moving in the same direction. The property of coherence, then, gives rise to the other two properties of laser light.

Coherence is not trivial and is brought about by the amplification mechanism of the laser. Unlike a laser, a beam of incoherent light is thought of as having a large number of photons, all acting as waves, with random phases relative to each other. As with any other wave, constructive and destructive interference will occur. The beam is a jumble of all sorts of photons at various wavelengths and phases. In a laser beam, all photons travel in lockstep with each other, cresting at the same time. The photons complement each other, rendering high amplitudes (high intensities). Indeed, although a 1-mW laser beam does not sound like much power, this intensity is packed into a spot 1mm in diameter. As a density this represents 1.2 kW/m^2, a level that surpasses the intensity of the sun on the surface of the Earth!

Monochromaticity is the ability of the laser to produce light that is at one well-defined wavelength. It is a requirement for coherence since photons of varying wavelengths cannot be coherent (ie, in phase). When white light is dispersed through a prism, it is composed of an infinite number of wavelengths of Light covering the entire visible spectrum as well as into the uv and ir regions. Emission lines from a gas discharge are much narrower when viewed on a spectroscope. The range of wavelengths spanned by such a line depends on many factors, such as gas pressure, but suffice it to say that it is finite and could well be 0.1 nm in width. The light produced by most lasers is much narrower, spectrally. A standard HeNe laser produces a red light centered at 632.8 nm with a linewidth of only 0.002 nm. That is much narrower than any incoherent source can produce. Some lasers can produce multiple output lines; the argon-ion laser produces up to 10 lines in the blue-green region of the visible spectrum. Each of those lines is also very narrow spectrally. The argon laser has many allowable lasing transitions. (Quantum mechanics shows that argon has intricate and complex energy level structures). Each of those transitions lases, and many can lase simultaneously, but each produces a spectrally narrow output.

Collimation is the property of laser light that allows it to stay as a tight, confined beam for large distances. It can be thought of as the spread in a beam of light (called *divergence*). This property of laser light makes it possible to use the laser as a level in construction or to pinpoint speeders on a highway. The simplest explanation for the highly directional output of the laser is in the mechanism of the laser itself. In the example of the HeNe laser, photons of light inside the tube must be reflected by the mirrors hundreds of times before exiting through the front of the laser tube. To do this, photons must be very well aligned to the axis of the tube: Photons emitted at even a slight angle to the tube axis will be bounced into the walls of the tube and will not contribute to the output beam. All laser beams have some amount of divergence. Some lasers with very high gains, such as nitrogen lasers, have wider divergences, whereas others, such as the HeNe laser, have low divergences. The divergence of a laser beam can never be made zero, though, since diffraction limiting (from the laser aperture) is always an overriding principle.

2.2. The Basic Mechanism of the Laser. Like any other light source, the emission of photons in a laser is generated by a jump or transition of a lasing

species from a high to a low energy state (an atom in the case of the HeNe laser or a molecule in some lasers). However, unlike other sources of light in which transitions occur at random, transitions in a laser are coaxed to occur at the same phase and wavelength as incident photons.

In any given system without an external source of energy, atomic populations are governed solely by temperature according to a Boltzmann distribution (Fig. 2). This distribution predicts an exponential drop in atomic populations of higher energies. A central concept of quantum mechanics is the condition that atomic energies are confined to certain allowed, discrete, states: it is between these states that transitions occur and lead to photon emission. Applying Boltzmann's law, the population of atoms at each discrete energy state may be predicted.

For mercury atoms, the first energy level is at 4.66 eV above ground. There are essentially no atoms with an energy this high at room temperature, so no light will be emitted from the mercury, as no transitions are possible. If the temperature is increased to, 2000 K, eg, the population of atoms at this particular energy level increases considerably [about one part in 10^{10}], a leading to emission of light when these atoms lose their energy by emitting photons and falling back to ground state.

Such a system is said to be at *thermal equilibrium*. The population of atoms at any given energy level is governed solely by temperature, since there is no external source of energy. If the temperature of the entire system were raised, the distribution would shift and more atoms would reach higher energies. However, the population of a lower level will always exceed that of a higher level. Had an incandescent lamp been chosen (a blackbody radiator), the situation would be the same. Electrical energy heats a tungsten filament. Many atoms at high energies drop to a lower energy state. In doing so, they emit a photon of light. The difference in energy between the upper and lower energy states for the transition is manifested as the photon energy, hence the wavelength of the emitted light. In such a system, there are a large number of upper and lower levels. These levels may span a range of energies so that the output from such a source is *broadband*. It is not at a single discrete wavelength like the line spectra emission from a gas, but rather, spans a range of wavelengths. No matter how the example above is considered, the situation still arises where there is a higher population of atoms at lower than at higher energy states.

Now, consider a situation where energy is injected into a system at thermal equilibrium to cause a population of atoms at a higher energy level to be greater than that of a lower level: called a *population inversion*. Such a nonequilibrium condition is indeed required for lasing action in almost all lasers with the exception being one special class of lasers (1). Energy may be injected selectively to pump an upper energy level from which transitions occur to a lower level, causing photons of light to be emitted. It is not necessary to excite all upper energy levels in the gas: The requirement is that a higher population of atoms exists at the upper level of a particular transition than at the lower level of that same transition. The reasons for this will become clear when considering rates of stimulated emission.

Pumping is the process of supplying energy to the laser medium to excite the upper energy levels and may be accomplished by any number of means,

including electrical, optical, chemical, or nuclear. Most common lasers utilize electrical pumping (as seen in the example with the HeNe laser) or optical pumping, as shown in Figure 3. In the case of optical pumping, light from a flashlamp or an arc lamp is focused onto a rod containing the lasing atoms. Common laser rods include ruby (chromium ions in an aluminum oxide host glass) and YAG (neodymium ions in a yttrium–aluminum–garnet host glass). The lasing atoms (chromium or neodymium in this case) absorb photons of incident pump light. Regardless of the method of pumping, the end goal of pumping is to excite high energy states within the lasing medium so that the population of atoms at a high energy state is greater than the population of atoms at a low energy state for the lasing transition.

Pumping can be effected to ensure that the upper level is filled without filling the lower level (ie, creating an inversion), by designing the pumping source or method to ensure that the upper level is pumped selectively while the lower level is not. For the HeNe gas laser, which consists of a glass tube containing both helium and neon gases in which an electrical discharge occurs the discharge, similar to that in the tubes of a neon sign, excites helium atoms by electron collision. The vast majority of energy is wasted in the process, some of it in the pink glow given off by the tube, still more in the form of heat. Some excited helium atoms (with a specific energy of 20.61 eV) can collide with neon atoms, and in doing so transfer energy to them at an almost identical energy level (the target is the neon level at 20.66 eV). The energy levels involved are depicted in Figure 4. This energy level for neon, which accepts energy from helium, is the upper level for the transitions that produce the visible HeNe laser output. As expected for a multielectron atom such as neon, quantum mechanics predicts a splitting of energy levels into a series of closely spaced levels. In the case of neon, the actual upper level at 20.66 eV splits into four discrete levels.

Through this process neon atoms are then pumped, quite selectively, to an upper energy level from which lasing transitions can occur (helium lacks a level close to the lower lasing level of neon, so it does not receive energy from collisions with helium). This scheme ensures that the upper energy level for the transition, which produces laser light, has a higher population of neon atoms than the lower level. The situation in the HeNe laser can be contrasted to the Boltzmann distribution of neon at thermal equilibrium in Figure 5 to see the inversion of the levels.

In Figure 5, the populations of neon atoms are compared both at thermal equilibrium and in an operating laser. The two levels of interest are the $2p^5 5s^1$ and $2p^5 3p^1$ levels. (This is the configuration of the excited neon atom that describes the state of electrons in the atom. Neon at ground state has a configuration $2p^6$ or, more fully expanded, $1s^2 2s^2 2p^6$.) These are the upper and lower lasing levels for the 632.8-nm red transition, the common red HeNe laser. Numerous other lasing transitions are possible given neon's many energy levels (eg, electron spin leads to a hyperfine structure, where the $2p^5 3p^1$ level is actually 10 closely spaced levels, giving us 10 possible laser transitions!). Using Boltzmann's equation, it is possible to calculate the populations of atoms at these two levels (or at any other level). The energy of the upper level is 20.66 eV (or 3.31×10^{-18}J). Knowing that the photon is at 632.8 nm, the lower levels calculated to be at 2.99×10^{-18}J (18.69 eV). The rest is trivial. As expected,

the population of atoms at the 20.66-eV level for neon in thermal equilibrium is extremely low, with $N/N_0 < 10^{-100}$ for room temperature neon. Even if there were a few atoms at these two levels involved in the lasing transition, there would still be many more at the lower level than at the upper level. Because excited helium atoms transfer energy to the 20.66-eV and not the 18.69-eV neon level, the lower level will not fill from collisions the way the upper level does (the level for helium is not at all close to the 18.69-eV level). This situation produces the inversion desired.

It is no accident that helium is used in a HeNe laser: Helium has an excited energy level very close to the level of neon used as the upper level for the lasing transition. Direct electrical discharge in pure neon does not excite this particular upper level enough to cause an inversion since the lower levels would fill as well in that case [although a fast electrical discharge in pure neon can produce lasing action on a different set of transitions than those described here (2)].

If a photon of light at an energy of E is fired into a cell containing atoms at ground state, it is possible that it will be absorbed and, in the process, pump the atom to a higher energy state. The energy of the incident photon must be at least equal to that of the upward transition. This, of course, is absorption. If the atom is already in a high energy state, though (the upper energy level we refer to continually), and a photon of the correct wavelength comes along (the wavelength corresponding to a transition from the upper state to a lower state in the atom), it can stimulate the excited atom to emit a photon of exactly the same wavelength and phase as the incident photon, leaving two photons exiting this process going in exactly the same direction. The fact that these two photons are identical makes the emitted radiation coherent and monochromatic, which are two key properties of laser light. The fact that they are emitted in the same direction will play a role (along with a well-aligned cavity) in making the light collimated, which is the third key property of laser light. In essence, the original photon is *amplified* by this process, which is called *stimulated emission*. The word *laser* is an acronym for "light amplification by stimulated emissionof radiation"; stimulated emission is the *key* process here. The process is diagrammed in Figure 6.

Of course, the atom that emits the photon loses its energy in the process and must be pumped to an excited state again or it will reabsorb another photon. This is why population inversion is (generally speaking) required. If inversion is not maintained, atoms will absorb rather that emit photons of light. This outcome is not trivial. When the population of atoms at the lower state exceeds that of the upper state, emitted photons of laser light are actually absorbed and laser action is not possible. This effect can actually be demonstrated by passing light from a gas laser (eg, a HeNe laser) through a cell containing the same, but unexcited, gas as the laser. In the process of passing through the cell, the power of the laser will actually be measurably reduced.

At this point, the analysis of the situation becomes statistical. Statistically speaking, a higher rate of stimulated emission than absorption is required. Each of these rates can be computed. Another consideration, is spontaneous emission. This creates an alternative, unwanted, pathway, allowing atoms in the upper energy state to lose their energy by emitting a photon of light spontaneously. In the case of the HeNe laser, this spontaneous emission shows up as the pink

glow from the tube (which is, of course, composed of discrete lines like any other low pressure gas discharge). Spontaneous emission robs the upper level of atoms that would otherwise be available for stimulated emission. Finally, the attenuation of emitted laser light by the gain medium must be considered as well as the mechanisms of the laser itself (windows, mirrors, etc). Generally, all mechanisms causing a reduction in laser gain (absorption in the lasing medium and optics, etc) are referred to as *losses*, which are considered later.

3. Practical Aspects

3.1. Criteria for Lasing. The general criteria for a net photon gain to occur (ie, a *laser gain*) is that the rate of stimulated emission must exceed that of spontaneous emission plus that of all the losses. The three processes involved in a laser are absorption, stimulated emission, and spontaneous emission. The rate of absorption of photons depends on the number of atoms in the lower state (ie, the number of atoms available to absorb photons) as well as the energy density of incident photons. The second parameter should be obvious. More photons to absorb leads to a higher rate of absorption. Mathematically, the rate of absorption may be stated as

$$r_{\text{absorption}} = B_{12} N_{1\rho} \tag{1}$$

where B_{12} is a proportionality constant called *Einstein's coefficient*, N_1 is the number of atoms at the lower energy state, and ρ is the energy density. The energy density in this case is specific. It represents the number of photons that have the exact energy for the transition between energy levels E_1 and E_2. Similarly, the rate of stimulated emission depends on the number of atoms at the upper state and can be written as

$$r_{\text{stimulated}} = B_{21} N_{2\rho} \tag{2}$$

where B_{21} is Einstein's coefficient, N_2 is the number of atoms at the upper energy state, and ρ is the energy density. Finally, the rate of spontaneous emission must be calculated. This rate does not depend on incident energy density (atoms emit photons spontaneously regardless of external conditions), but solely on the number of atoms at the upper energy state available to emit a photon:

$$r_{\text{spontaneous}} = A_{21} N_2 \tag{3}$$

where A_{21} is Einstein's coefficient for spontaneous emission and is simply the inverse of the spontaneous lifetime of the species. Level lifetime is an important consideration in determining the possibility of laser action and especially the ability of a species to operate in CW (continuous) mode.

Rates are equated at thermal equilibrium. Assume, then, that there is no external energy input to the system. The rate of upward transitions (absorption)

can be equated with the total rate of downward transitions (spontaneous and stimulated emission):

$$r_{\text{absorption}} = r_{\text{stimulated}} + r_{\text{spontaneous}} \tag{4}$$

or

$$B_{12}N_{1\rho} = B_{21}N_{2\rho} + A_{21}N_2$$

Now Planck's blackbody radiation law as well as Boltzmann's law are used to solve for the energy density ρ. Use Planck's radiation law and substitute for the energy difference between levels E_2 and E_1 via Boltzmann's law, so that

$$\rho = \frac{8\pi h\nu^3}{c^3 N_2/N_1 - 1} \tag{5}$$

Realizing that the two Einstein B coefficients represent the probability of an atom making the upward transition E_1 and E_2 and the downward transition E_2 to E_1, respectively, and that these are equal, the original rate equation 4 is solved to yield

$$\frac{8\pi h\nu^3}{c^3} = \frac{A_{21}}{B} \tag{6}$$

Now that there is a solution for the Einstein rate equations, some mathematical conclusions can be drawn. Consider the ratio of stimulated-to-spontaneous emission. As already discussed, the rate of stimulated emission must exceed that of spontaneous emission for laser action to occur (in order to have amplification occur):

$$\frac{r_{\text{stimulated}}}{r_{\text{spontaneous}}} = \frac{B_{21}N_2\rho}{A_{21}N_2} \tag{7}$$

Substituting for the Einstein coefficients (eq. 6), this ratio simplifies to

$$\frac{r_{\text{stimulated}}}{r_{\text{spontaneous}}} = \frac{c^3\rho}{8\pi h\nu^3} \tag{8}$$

It is further possible to solve the ratio outright by substituting for ρ using Planck's law from equation 5. Since the energy levels involved are known, the frequency of emitted photons ν is known.

Now, an important conclusion is that for this ratio to be large (ie, stimulated emission rate exceeds spontaneous rate), the energy density of incident photons must be high. This is the actual flux of photons within the laser cavity itself. Unless the gain of the medium is extremely high in order to create a huge flux of photons as they pass down the tube, cavity mirrors will be required to contain

photons within the cavity to create further amplification. Indeed, the vast majority of lasers require cavity mirrors for oscillation. Furthermore, these mirrors are usually very efficient (>95%) reflectors. Note that the ratio of these rates also depends on the inverse of the frequency, cubed. In other words, higher frequencies (shorter wavelengths) have lower ratios and require higher energy densities to operate. As an aside, consider that the first laser operated at microwave frequencies, where much lower energy densities are required to favor this ratio over optical frequencies.

Finally, the concern is with the rate of absorption of photons. It is necessary that the photons emitted are not absorbed within the medium itself. If the medium absorbs more photons than were emitted by stimulated emission, the laser cannot work. This ratio is calculated directly from equations 1 and 2 as

$$\frac{r_{\text{stimulated}}}{r_{\text{absorption}}} = \frac{N_2}{N_1} \tag{9}$$

But in this case, there is a population inversion where $N_2 > N_1$, so the rate of stimulated emission will indeed exceed the absorption rate. This relation proves the necessity of population inversion for laser action (for which only a qualitative argument has been given so far). This is a simple introduction to the rate equations, which can become considerably more complex, especially when multiple levels other than just the two lasing levels are involved, which they are in a real laser. Rate equations and the types of energy level systems involved in real lasers are examined in much greater detail in Chapter 5 Ref. 3.

3.2. Gain and Loss in a Laser. Assuming that a population inversion has been generated (by pumping energy into the lasing medium, in this example a gas, to excite upper lasing levels), laser gain may now occur. Laser gain (or optical gain) is a measure of how well a medium amplifies photons by stimulated emission. Consider a single photon traveling down a laser tube and stimulating an excited atom to emit a photon of the same frequency and phase: This is *laser gain*. The stream of photons travels down the tube as an electromagnetic wave, so the *power* of the wave increases if the rate of stimulated emissions exceeds that of spontaneous emissions. As the wave travels farther down the tube, the power increases as a function of length. Mathematically, the power increase is

$$\exp(gx) \tag{10}$$

where g is the optical gain coefficient of the laser medium and x is the distance down the tube. The power increases exponentially as the wave travels down the tube. Consider where one photon stimulates another, producing two photons; these two then stimulate two more, so these are then four, and so on. The exponential nature of this gain is evident. The gain coefficient, g, represents the optical power gain per unit length. Mathematically, it can be expressed as $\Delta P/\Delta x$, the change in power for a given length of lasing medium, or more correctly,

$$g = \Delta P/(P \times \Delta x) \tag{11}$$

which yields a quantity in units of length^{-1} (eg, m^{-1}). Gain is proportional to the net rate of stimulated emission, which is in turn is equal to the rate of stimulated emission minus the rate of absorption. The ramifications for practical laser construction are simple. The longer the lasing medium, the more power will increase per pass down the tube. When attempting to lase a weak transition (as many gases have), a longer tube will allow a larger power increase per pass through the tube. This may be important when considering losses in the laser.

When the laser is operating at steady-state conditions (ie, constant optical power output), the net gain in the laser must be 1. If the net gain were greater than 1, the output power would increase. Net gains below 1 cause the output power to drop until the laser ceases to operate. So consider, a round trip by a stream of photons in the tube as in Figure 7. The power after the round trip must equal the power before the round trip. The photon stream gains energy (through laser gain), but just as much power as is gained is lost through absorption in the medium, in addition to the portion extracted as the output beam.

The power gained during the round trip is $\exp(g2x)$, where $2x$ is the total path length of the laser (x is the length of the gain medium through which the stream of photons passes twice). The power lost during the same trip is $\exp(-\gamma\ 2x)$, where γ is a new term describing all losses in the cavity due to the lasing mechanisms, with the exception of the mirrors themselves. The losses from the mirrors themselves are seen as the reflectivities of the two mirrors, labeled R_1 and R_2. A perfect mirror, with 100% reflection, has an R value of 1.

Equating these parameters for a round trip through the tube yields

$$\begin{aligned}\text{Net gain} = &\ \text{Laser gain} \times \text{Loss} \times \text{Loss at morror 1} \times \text{Loss at mirror 2} \\ = &\ \exp(2gx) \times \exp(-\gamma 2x) \times R_1 \times R_2\end{aligned} \tag{12}$$

If knowing $R_1$and R_2 and g are known gain, can be determined, which must equal 1 for an operating laser. This gain will be the *threshold gain* or the gain required to allow the laser to operate in this steady-state condition:

$$g_{\text{threshold}} = \gamma + \frac{1}{2x}\ln\left(\frac{1}{R_1R_2}\right) \tag{13}$$

where x is the length of the gain medium. In many lasers, the gain of the laser medium is proportional to the pump energy (but not without limits). A minimum pump energy will be required to generate a gain of at least the threshold value in order for lasing to begin.

While it is convenient to consider gain as a finite quantity occurring at a specific wavelength, this is not an entirely accurate model since gain is "spread-out" across a series of wavelengths as shown in Figure 8, which represents the gain curve of a typical gas laser.

Mechanisms such as Doppler broadening, in a gas laser, as well as other mechanisms such as crystal lattice vibrations, in a solid-state laser, serve to distribute the gain over a series of wavelengths. At all points on the gain curve where the gain is sufficient to overcome losses in the laser (and the laser cavity is resonant), the laser may oscillate and have output. It is not that a narrow gain

profile (and hence narrow output linewidth) is always desirable, and it is worth mentioning that sometimes a wide gain profile is desirable in order to generate ultrashort pulses using a technique called modelocking.

There is a further complication in this scenario. The cavity itself, which acts as an interferometer, is resonant only at certain wavelengths. The configuration of the cavity leads to longitudinal modes: a series of discrete frequencies close together. Collectively, they resemble a broad line much wider, spectrally, than a single line.

In a laser medium, gain is not a constant value, but rather varies with incident (or, in a practical laser, circulating) power. Imagining the laser gain medium as an amplifier, the gain is quite large. This is termed the *small-signal value*, in which the ULL is well populated and is replenished continually by pumping, so that the population does not change appreciably (in other words, the rate of pumping is large enough to keep the ULL populated and to keep inversion large). As the input signal to the amplifier reaches a large value, the photon flux is large enough to depopulate the ULL, which then lowers the overall gain in the device since fewer excited atoms will be available at that level to contribute to the stimulated emission process. The result is a saturated gain figure in which the gain is reduced by large photon fluxes in the cavity. Whereas the gain of an unsaturated amplifier is exponential with length, in a completely saturated amplifier, power increase is a linear function of length. In an operating laser, gain is "burned down" until it matches losses in the system so that gain equals loss. Extra pump energy is manifested as increased power output in the system.

4. Atomic Systems in Lasers

Lasers are classed by the number of energy levels involved in the actual lasing process as three- or four-level lasers. In a three-level system (the simplest), energy injected into the gain medium excites atoms to a pump level above the upper lasing level. From there, atoms decay to the upper lasing level. This decay usually occurs by emitting heat, not photons. It is rapid and quickly populates the upper energy level, which often has a long lifetime, so a healthy population of atoms builds in that level. Lasing transitions now occur between the upper level and the ground state, emitting laser light in the process. This system is characterized by the lack of a discrete lower lasing level; the ground state serves that purpose.

Four-level systems feature a discrete lower lasing level between the upper and ground states. Atoms making a laser transition to the lower state decay further to the ground state, in some cases by emitting a photon (in the case of an argon-ion laser, a 74-nm photon is emitted during this decay to the ground state). Four-level lasers are by far the most common. In a helium–neon laser, eg, the pump level is in helium (it need not be in the same species as the lasing atom), while the upper and lower lasing levels are in neon. A carbon dioxide laser has a similar situation, where the pump level is in nitrogen gas, which like helium in a helium–neon laser, is in larger proportions than the lasing species. It is also possible in some lasers, such as metal-vapor lasers, to pump directly to

the upper lasing level (ie, there is no pump level). In this case, atoms are excited to the upper lasing level, where they make a transition to the lower lasing level, finally decaying to the ground state. These are particularly efficient systems since they lack an energy loss when atoms in the pump level decay to the upper lasing level (which, in most cases, shows up as heat in the laser). Although the metal–vapor system has only three levels, it has more in common with a four-level system than does a classic three- level system.

In a four-level laser, laser gain is realized as soon as pump energy is applied to the system. Pump energy is injected into the pump level, where it decays, in most cases almost instantaneously, to the upper lasing level. Assuming that the upper level has a longer lifetime than the lower level (in most four-level lasers it does), a population inversion occurs almost immediately after pump energy is injected. Although there may not be a usable output beam (since any small gain produced will be lost to absorption and other losses within the system until threshold is reached), there is a population inversion, and hence a gain. Injecting a little more energy into the system will raise the gain to a level where it exceeds lasing threshold and a usable output beam appears.

So the nature of a four-level laser is such that a population inversion is easy to achieve, but what happens in a three-level laser where the lower lasing level is the ground level? Boltzmann statistics show that at a given operating temperature there will be an enormous population of lasing atoms at the ground state, and that to achieve a laser gain over one- half of these must be raised to the upper lasing level, which in many cases is far above ground. Consider the dynamics of the ruby laser system in Figure 9, in which the upper lasing level is $\sim$1.8 eV above ground. The thermal population of this level at 300 K is negligible. Also evident in the figure are two pump bands in the blue and green regions of the spectrum.

Strong pumping is required for such a laser to achieve a population inversion, and these lasers exhibit a high threshold of pump energy. A considerable amount of pump energy is required just to achieve inversion, let alone lasing threshold. There is also a time delay between the onset of pumping power and lasing output. As pump energy is injected into the ruby, one must wait until a significant population of atoms (certainly, over one-half of the chromium ions in the rod) reaches the upper lasing level before lasing can begin. So why does ruby work quite well as a laser material? The reason lies in the two broad pump bands, which readily absorb energy from a flashlamp (xenon flashlamps emit a good portion of their radiation in the violet-green region, so are commonly used with ruby lasers). As well as being broad, the pump bands have incredibly short lifetimes (on the order of 1 μs), which causes energy absorbed into these pump bands to relax almost immediately to the upper lasing level. Finally, the lifetime of the upper level, 3 ms, is quite long, allowing the excited ions to remain there long enough to have a good chance of emitting light by stimulated (as opposed to spontaneous) emission. Although most ruby lasers are pulsed, it is possible to operate the material as a CW laser. Special conditions are required, however, such as extreme cooling to depopulate a portion of the ground level; this is, not surprisingly, a collection of closely spaced levels. Doing so allows the system to operate more like a four-level system, as described next.

Now, in contrast to the three-level ruby laser, consider a four-level laser such as the YAG (also a solid-state laser), with energy levels as depicted in Figure 10. The YAG system is characterized by a cluster of pump levels from which excited atoms decay rapidly to the upper laser level. There are multiple pump levels, allowing the system to absorb energy at a variety of wavelengths, including the important band at 790–810 nm, useful for pumping via a semiconductor laser. All pump levels have short lifetimes, ~100 ns, and decay rapidly to the upper lasing level. The upper level has a very long lifetime of 1.2 ms compared to the lower level, which decays to ground level 30 ns. Many lasing transitions are possible in this system, including the most powerful, at 1064 nm. The lower lasing level of the transition at 946 nm is extremely close to the ion ground state. Because of the large thermal populations of the lower lasing level, lasing on this transition occurs in a mode resembling a three-level laser, and strong pumping (or extreme cooling) is required to allow this transition to oscillate. For other transitions such as 1064 nm, the situation is more favorable and a population inversion will occur quickly.

Perhaps the most dramatic effect of the type of energy level system is the amount of pump energy required to generate laser action. In any laser, output can begin as soon as population inversion occurs, followed by stimulated emission. Gain increases until it exceeds total losses in the system and laser output occurs. In a three-level laser, it takes a large amount of pump energy just to generate an inversion, since the population of the ULL must exceed that of the ground state. In a four-level laser, inversion can be achieved with much lower pump energies, since the LLL is a discrete level that is not populated to nearly the level of the ground state. Assuming that the LLL for a four-level laser is a few electron-volts above ground state, it will not be thermally populated to any large extent—unlike the three-level laser, in which the ground state has a large population, making inversion difficult to achieve. The vast majority of lasers, including HeNe, argon ion, and carbon dioxide are all four-level lasers.

5. Cavity Optics

When the rate equations of a laser were considered, one important result is that for the rate of stimulated emission to exceed the rate of spontaneous emission (a required condition for laser action), the energy density of incident photons (denoted ρ) must be large. Unless the gain of the medium is extremely high in order to create a huge flux of photons as they pass down the tube, or the gain medium is very long, cavity mirrors will be required to contain photons within the cavity, keeping ρ high and creating further amplification. Indeed, the vast majority of lasers require cavity mirrors for oscillation. Furthermore, these mirrors are usually very efficient reflectors where typical values for a gas laser are 99.99% reflectivity for the HR and 99–98% reflectivity for the OC. The resonator must trap photons completely within the cavity since any photons that escape the resonator represent a loss in the laser, which can drastically affect output power.

Assume that a laser operates with an output of 1 mW and has an output coupler with a reflectivity of 99.0% (and hence a transmission of 1.0%). This

implies that the power circulating inside the cavity (the *circulating power*) is 100 mW. Of course, such intracavity power level is not usable. It is required to keep the rate of stimulated emission large enough, and hence sustain laser action. The output beam (defined by the transmissivity of the output coupler) represents a loss in a laser cavity. Other losses include absorption and scattering in the laser medium itself and losses at windows terminating the lasing medium (eg, windows on a gas laser plasma tube). Most laser gain media have low gains, so it is essential that losses be kept to a minimum. Since the output coupler represents a loss, the amount of transmissivity depends on the gain of the laser. When total losses in the laser exceed the gain of the medium, the laser will fail to oscillate.

In an experiment illustrating the effect of loss in a laser, a helium–neon gas laser was setup in the lab in which a variable loss (a glass slide at a tilted angle) is inserted into the resonator cavity of the laser. By changing the angle of the slide in the cavity with respect to the optical axis of the laser, the reflectivity of the glass–air interface is changed. The reflectivity of the slide, as computed using a Fresnel equation, represents a loss inserted into the laser. When output power (which is proportional to circulating power) is plotted against loss in the cavity (as shown in Fig. 11) the effect is obvious.

Circulating power is easily measured simply by measuring laser output and multiplying by the transmission of the OC. In this experiment, the OC has a transmission of 1.0%, so that for a measured output of 370 μW, the intracavity circulating power is 37 mW. The maximum inserted loss that may be tolerated in this laser is found to be 1.2%. When this loss is summed with other losses in the laser, such as absorption in the medium itself and the loss of the OC (1.0%), the gain of the laser medium may be calculated.

It is obvious from Figure 11 that even small losses have a large effect on the output power of a laser: losses in a real laser must be kept as low as possible. Where possible, such as in a HeNe laser, mirrors are often sealed directly into the ends of the tube so that there are no windows in the optical path to increase loss. In many lasers, though, cavity mirrors must be isolated from the discharge to prevent attack from the plasma. Use of external optics is also required when a laser uses a wavelength selector. In these cases, the ends of the plasma tube must be capped with an optical window. To minimize losses, optical windows are often angled at the Brewster angle, which polarizes the output of the laser since loss is essentially zero in one polarization and very significant in the other. The optimal angle (Brewster's angle) can be determined using Fresnel equations to determine the angle at which reflection (R_p) is reduced to zero. The other alternative to a Brewster window is to use a plane window with an antireflective coating to reduce loss due to reflections at the air–glass surfaces.

In general, a laser cavity acts as an interferometer into which an integral number of waves must fit. The cavity is resonant at wavelengths such that the number of waves inside the cavity is an integer: These are standing waves inside the cavity. At all other wavelengths, destructive interference causes any wave inside the cavity to be extinguished. These resonant wavelengths are called *longitudinal modes* and are spaced apart at regular intervals of frequency. The spacing of modes is called the *free spectral range* (FSR) of the interferometer and is outlined in Figure 12.

Assuming that the distance between the cavity mirrors is L, the condition for a standing wave in the cavity is

$$m(\lambda/2) = L \tag{14}$$

where m is an integer, so the spacing (FSR) of resonant modes in the cavity is

$$\Delta v = m(c/2L) \tag{15}$$

where Δv is the spacing in hertz.

As mentioned earlier, the linewidth of laser transitions is not extremely narrow and that effects such as Doppler broadening leads to the existence of a gain curve in which the gain of the laser (and hence the output as well) peaks at a center wavelength. At all points on the gain curve where the gain is sufficient to overcome losses in the laser (and the laser cavity is resonant), the laser may oscillate and have output. The cavity itself is an interferometer and is resonant only at wavelengths spaced apart by the FSR of the arrangement. The actual output of the laser will not be a smooth curve, then, but rather a series of closely spaced wavelengths that follow the general envelope of the curve and exist at points where the gain exceeds losses in the cavity. The development of the output is seen in Figure 13, where the top diagram shows the gain curve of the laser medium along with the lasing threshold (the point at which gain equals losses). In a simple model, one would expect an output spectra resembling the shaded area on the top diagram. The response of the cavity, resonant at frequencies spaced apart by the FSR, is shown in the middle diagram. The result is depicted on the lower diagram, in which the output is seen to be several modes (11 in this case), with the strongest at the center of the curve (where the gain is highest) and the weakest near the edges, where gain just slightly exceeds losses.

Since the laser cavity is an interferometer, the resonant modes of the cavity are spaced apart by FSR, which is defined as

$$FSR = c/2nl \tag{16}$$

where n is the refractive index of the medium inside the cavity and l is the spacing of the cavity mirrors. In this case, the FSR, is in terms of frequency (Hz). The number of modes that will oscillate simultaneously may be determined by dividing the spectral width of the laser by the FSR.

As well as longitudinal modes that define the span of wavelengths over which laser oscillation occurs, the output beam of a laser may be defined spatially by transverse modes. Transverse modes describe the manner in which energy is stored, spatially, throughout the entire lasing volume and how that volume is utilized. The nature of the cavity gives rise to electromagnetic modes in which standing waves are set up not only in the longitudinal direction (ie, the length of the cavity; these modes vary slightly in frequency), but also in the transverse direction. These represent alternative solutions to the wave equation that defines the beam (which, in a "pure" beam—designated TEM_{00} mode—yield a Gaussian solution). With energy stored in various areas of the lasing medium, the patterns

formed are also manifested in the output beam, which can assume shapes such as those shown in Figure 14. Six common modes are shown.

TEM_{00}, the Gaussian mode, is shown in the upper left corner of the figure, which represents the "purest" beam possible. The beam has circular symmetry where intensity is highest in the center and tapers off equally in all directions. In many cases, the highest power output from a particular laser can be obtained on a high-order mode in which the volume of the gain medium is utilized more effectively than is possible with the TEM_{00} mode. This finding is evident by the fact that the TEM_{11} and some other higher order modes are considerably brighter than lower order modes shown in Figure 14 (all photographed with the same exposure time for comparison). This fact can be exploited (in reverse) to limit a laser to operation in TEM_{00} mode by using an intracavity aperture to limit the use of the lasing volume, and hence restrict the mode.

6. Fast Pulse Production

Pulsed solid-state lasers are among the most important and widely deployed commercial laser systems. The Nd:YAG laser, eg, is the most popular commercial laser for applications such as marking, cutting, drilling, range finders, and retinal surgery. For many of these applications, the shorter the pulse, the better. Fast, powerful pulses tend to ablate material quickly without heating (and potentially altering) surrounding material or tissue. As an example, consider the precision trimming of thick-film resistors. Such resistors are fabricated directly on a ceramic substrate using screen-printing techniques, then trimmed to the correct value using a YAG laser. Fast pulses will quickly vaporize resistor material in a controlled manner, increasing the resistance of the structure by a controlled amount. A CW laser beam, or one with a long pulse, would tend to heat the material surrounding the target area and either damage the substrate or alter the resistor in an uncontrolled way by changing the properties of the resistor film. For this reason, a technique called Q-switching is used to produce very short pulses for these types of applications.

In a *Q-switching technique*, the laser output is switched by controlling loss within the laser cavity as outlined in Figure 15. More correctly, Q-switching is *loss switching* in which a loss is inserted into the cavity, thus spoiling it for laser action. In the simplest manner, a Q-switch can be though of as an optical gate blocking the optical path to one cavity mirror, and hence causing laser action to cease. It is not necessary to block the optical path completely. Simply inserting a loss high enough to raise the lasing threshold beyond the maximum gain of the laser is sufficient. When the Q-switch is on, it blocks the intracavity beam. This state is called a *low Q-state*, meaning that the quality factor, or Q, of the cavity (which measures the ability of a laser cavity to act as a resonator) is ruined. With the switch off, losses in the cavity are reduced, and the cavity is resonant: a high Q-state. During the time when the Q factor of the cavity is low, the laser is not oscillating, and hence has no usable output beam, but the pump energy continues to drive the laser medium. In this state, population inversion continues to build and an inversion much larger than normally possible builds in the laser gain medium. When the Q of the cavity is restored in such

a condition, the laser begins to oscillate immediately and the large population inversion releases energy in the laser medium in one enormous pulse. This technique allows energy to be stored in the laser medium during the low Q-state and is released in a single massive pulse. The peak power of the pulse is much larger than is possible with gain switching, which for a CW laser, yields a peak power equal to the CW output power of the laser.

The entire concept of Q-switching relies on the fact that energy can be stored in the lasing medium itself in the form of an excited atomic population at the ULL. A large Q-factor represents a low loss resonator that can store a large amount of energy. In Q-switching, the Q of the cavity is spoiled (the Q-factor is purposely made low), so that it is not resonant, and hence lasing is not possible. Energy storage takes place, but rather than within the cavity as optical energy, energy is stored in the atomic population.

In a Q-switched laser, the laser medium itself is used as a sort of capacitor, storing energy gradually and releasing it in a single burst. Not surprisingly, the capacity of the medium to store energy depends on the lifetime of the upper lasing level (ULL). A long lifetime implies that the lasing species can absorb energy over a long period without losing it to spontaneous emission; hence it has a large storage capacity and is a good candidate for Q-switching.

If a laser cavity is blocked, so that laser action cannot occur, the inversion is now free to build. This is unlike CW lasers, where the level of population inversion grows to reach an equilibrium level. The primary limit on this inversion is the lifetime of the upper lasing level, which serves to deplete the upper level through spontaneous emission. With continuous pumping the inversion builds until reaching a level equal to the rate of pumping times the lifetime of the upper level. Excess pump energy after the maximum inversion is reached is essentially wasted. Still, in many lasers, especially those driven by continuous sources such as CW arc lamps, the time during which the laser is pumped may be very long. In this case, the Q-switch is also used to allow the laser to fire on command and so when not actively firing, the pump source (the lamp) is kept on, so that the laser is always in a ready state. Implications are that any laser medium may be Q-switched, but it should be evident that the lifetime of the ULL must be much longer than the opening time for the Q-switch. For most gas lasers with ULL lifetimes of under 1 ns, Q-switching simply will not work, since no practical switch exists that can open in considerably less than this time. One of the few exceptions is the carbon dioxide gas laser, which has an exceptionally long ULL for a gas laser, although in practice, Q-switching of a CO_2 laser is rarely done. Solid-state lasers, on the other hand, invariably feature long ULL lifetimes. The ruby laser, eg, has a lifetime of 3 ms, and the YAG laser has a lifetime of 1.2 ms. These are ideal candidates for Q-switching, which is a standard option on most solid-state lasers.

Most common Q-switches are either of the electrooptic (EO) or acoustooptic (AO) type. The EO modulators rely on the phenomenon of *birefringence*, as induced in a crystal by the application of an external electric field. Birefringence is an effect exhibited in certain crystalline materials in which an incident light ray separates into two rays that may travels in different directions. The direction traveled is dependent on that ray's polarization. For each of the two rectangular states of polarization (perpendicular and parallel), the light travels in a different

direction as it exits the crystal. This effect is also called *double refraction*, since light entering the crystal is refracted into two different directions. Calcite is a natural crystal that exhibits birefringence. This crystal separates an incident beam of light into two beams with different polarizations.

There are a number of crystals that exhibit birefringence only when an external electric field is applied. This phenomenon, called the *electrooptic effect*, may be used in the case of a Q-switch to block an intracavity light beam. An electrooptic modulator is shown schematically in Figure 16. Incident light is polarized and then passed through the electrooptic material (this polarization is not needed if the incident light is already polarized). Light exiting the crystal is passed through a second polarizing filter, called an *analyzer*. Since the crystal causes no change to the light passing through it when not energized, the analyzer, oriented at 90° to the polarizer, prevents any light from being transmitted through the device. In essence, when the crystal is not energized, the device acts as a shutter. When an external electric field is applied to the crystal, the direction of the polarization of light passing through it is rotated by 90°. Then the light will pass through the analyzer; effectively the shutter is open.

In a modulator, the two components (with different polarizations) travel through the crystal in the same direction; they do not separate from one another. However, the relative phase of each component changes. Since the index of refraction is different for each polarization, they race through the crystal at different velocities, and so emerge from the crystal with different phases. As they continue through the crystal, the phase difference increases. This depends on the thickness of the crystal as well as on the electric field applied. When the beams emerge from the crystal, the polarization of the combined single beam depends on the accumulated phase difference between the two individual polarized components. A phase difference of one-half wavelength results in the overall polarization of the beam being rotated by 90° from its original direction. When this polarization of light falls on the analyzer, it passes through.

Two EO effects are possible, classified according to how the refractive index changes relative to the applied electric field: the Pockels effect and the Kerr effect. Most EO modulators utilize the *Pockels effect* in which the change in refractive index is a linear function of an applied electric field E. Crystals such as KDP (potassium dihydrogen phosphate) exhibit this effect.

Perhaps the simplest (and most common) modulator is the acoustooptic modulator. An acoustic wave, at radio frequencies, originates from a piezoelectric crystal that generates a surface acoustic wave in a crystal. The waves generated by the piezoelectric crystal form almost entirely flat plane wavefronts in the crystal. This wave is transmitted through a quartz crystal and induces a strain in the crystal that changes its refractive index on a localized scale. Like a sound wave propogating through air, the acoustic wave creates rarefactions and compressions that change the density of the crystal. Regions of high and low index of refraction will similarly exist in the crystal at these points, as outlined on Figure 17. This photoelastic effect, induces strain that distorts bonds between atoms of the crystal. The end of the crystal is either cut at an angle or terminated with an acoustic absorber so that a standing wave is not generated in the crystal.

Because these layers of varying n are stacked together, diffraction will occur and an incoming beam of light will change deflection on passing through

the crystal. Two types of diffraction may occur here, Bragg and Raman–Nath diffraction, depending on the angle of incidence of the beam to be diffracted. In the Bragg regime, a light beam is deflected by the alternating layers of differing refractive indexes. The acoustic wave inside a crystal sets up parallel planes of varying indexes of refraction, which generate Bragg diffraction of an incident light beam. An incident light beam on an AO modulator crystal is deflected only when an acoustic wave is present in the crystal (ie, it is driven by a radio frequency, rf, source). Most AO Q-switches are setup to use Bragg diffraction since this results in a large loss when in the "on" state, which is desired to spoil the Q of the cavity and completely foil laser action. As well as being useful as a modulator (by deflecting the intracavity beam out of the cavity, hence preventing amplification when the modulator is on), this type of diffraction may be used to deflect a beam, and hence may be used as a beam scanner.

The final method to Q-switch a laser is to use a saturable dye absorber. An incident pulse inside the laser cavity bleaches (saturates) the dye, which opens the switch, allowing the pulse through it.

Q-switching can generate fast pulses, however, the shortest pulses of light are produced using a technique called modelocking. In a modelocked laser, energy in the laser itself is compressed into a single packet of light that traverses the laser, reflecting from cavity mirrors and through the gain medium. The pulse inside the cavity is much shorter than the cavity itself: If the cavity was 1 m long, the pulse might typically be 10 cm within this cavity. If the cavity contains a partially reflecting mirror as an output coupler, a short output pulse is transmitted each time the pulse is reflected from that mirror.

As depicted in Figure 18, the modelocked laser consists of a normal Q-switched laser in which the Q-switch is opened at regular intervals corresponding to the transit time of the pulse within the cavity ($c/2L$). Once per round trip through the laser cavity the Q-switch is opened (Fig. 18c) to allow the pulse to pass; at all other times, the switch is closed to prevent any other light from oscillating in the cavity except for this modelocked pulse. With the modulator in the center of the cavity, it is required to open twice for each round trip of the pulse. If it is placed near one mirror in the cavity (as in the figure), only one opening of the switch is required per round trip.

The output of a modelocked laser with the configuration described is a continuous series of short pulses. In the case of a laser with mirrors 1 m apart, the pulses will appear at a frequency of $c/2L$ or 150 MHz. Pulse duration depends, among other factors, on the time for which the Q-switch is open as well as the gain bandwidth of the lasing species. Analysis of the bandwidth argument requires modelocking to be viewed in the frequency domain that is beyond the scope of this article.

Q-switches for a modelocked laser must open and close in a very short time period. Regular Q-switches, such as the AO modulators used for a Q-switched laser, are generally not fast enough for these purposes, so EO modulators are usually employed. One possibility for an AO modulator, however, is to set up a standing wave in the modulator. An acoustic wave can be generated that bounces back and forth through the crystal. At two points in the period of the wave there is a point where light is not diffracted (ie, a node in the standing wave where the electric field is zero), and hence the switch is open at those

points. In our example of a 1-m cavity, such a modulator opening at a rate of 150 MHz will modelock the laser. Note that the frequency required of the modelocker increases as the cavity size decreases.

7. Nonlinear (Harmonic) Light Generation

In 1961, it was discovered that certain materials could literally double the frequency of laser light passing through them (4). In the demonstration outlined in Figure 19, a powerful laser beam from a ruby laser at a wavelength of 694.3 nm in the red was focused on a quartz crystal. The beam exiting the crystal was analyzed by a spectroscope that recorded spectral components of the exiting beam on a photographic plate. Most of the source light passed through the crystal unchanged, but a small portion (1 part in 10^8) was converted to light at exactly double the frequency (one-half of the wavelength) or 347.1 nm in the uv. Since this discovery, much study has been made of nonlinear materials and techniques have been devised to exploit this phenomenon in an efficient way. Today, laser pointers are commonly available that use this technique to generate green light by doubling the ir output from a tiny diode-pumped neodymium solid-state laser, which normally has an output in the infrared.

Many materials, such as a spring, behave in a linear manner in which, eg, the deflection of the spring is linear to the force applied (eg $F = kx$). When large forces are applied, though, nonlinearity may be seen in which high order terms such as x^2 or x^3 enter into the equation. Many optical materials behave in a similar manner; when the force applied (ie, light intensity) is small, they are quite linear, but when a large force is applied, nonlinear behavior is exhibited. The structure of a quartz crystal is a lattice with atoms placed at regular patterns. Electrons in the crystal are held to the nucleus of each atom by a force similar to a spring. When low and moderate light intensities are incident on such a crystal, electrons behave in a linear manner. Bonds stretch in response to an electric field in a linear manner. However, when an intense electric field strikes an atom in such a crystal, valence electrons move to one side of the atom, and the electron shell becomes eccentric and offset. Instead of a spherical shell (in which the shell is centered around the positive nucleus, and hence there is no electric dipole), electrons literally move toward one end of the atom and the entire atom becomes polarized. A positive and negative end develop on the atom, and an electric dipole with a dipole moment develops in the direction of the electric field applied.

In the presence of an applied electric field, atoms in the crystal become polarized in this manner, and the entire material becomes polarized on a localized level (around the intense beam). This is a *macroscopic charge polarization* of the material, and the amount of charge polarization depends on the intensity of the applied electric field. When the applied force is very large, nonlinear effects are exhibited for the same reason that a spring exhibits nonlinear behavior. The charge polarization becomes nonlinear according to a geometric series:

$$P = a_1 E + A_2 E^2 + a_3 E^3 + \cdots \tag{17}$$

The nonlinearity of the charge polarization becomes apparent, and high order terms such as a_2 and a_3 contribute to the polarization.

To accomplish this nonlinear polarization, an intense electric field is required. The electron is held to the positive nucleus by an internal force of $\sim 10^9$ V/cm. Normal light sources with electric fields <100 V/cm exhibit only linear optical effects. Few sources exist with a high enough electric field (comparable to the internal force) to generate nonlinear effects. A focused laser beam, eg, can impress an electric field of 10^6–10^7V/cm on a crystal. Even higher intensities, with even higher electric fields, can be found inside the cavities of most lasers, where the intracavity intensities are often 100 times the external intensity.

To demonstrate mathematically how second harmonic light is produced under such conditions (eg, how a ruby laser at 694.3 nm can be doubled to 347.2 nm), consider the electric field applied to the crystal as $E = E_0 \cos(\omega t)$, where ϖ is the angular frequency of the light ($2\pi f$). Equation 17 then becomes

$$P = a_1 E_0 \cos(wt) + a_2 E_0^2 \cos^2(wt) + \cdots \tag{18}$$

Trigonometric substitution can be made for the $\cos^2$ term in terms of $\cos(2\omega)$ so that the equation becomes

$$P = a_1 E_0 \cos(wt) + \frac{1}{2} a_2 E_0^2 - \frac{1}{2} a_2 E_0^2 \cos(2wt) + \cdots \tag{19}$$

This leads to an interesting conclusion since the third term in equation 19 contains a term with order 2ϖ; one component of the charge polarization equation is at *twice* the fundamental frequency of the original laser beam. As the incident beam falls on the crystal, oscillating electric dipoles absorb light at a frequency ϖ and reradiate it at the same frequency ϖ and at twice the original frequency 2ϖ. The total energy of the beam is not altered much, but is now split between two components. This is the second harmonic of the original beam and the effect is called *second-harmonic generation* (SHG).

Nonlinear harmonic generation is possible only if the nonlinear crystal has no internal symmetry, so that electron shells are relocated more effectively in one direction than in the other by the electric field applied. About 10% of all crystals have such a lack of symmetry, and hence can be used as nonlinear media. The most common nonlinear materials are ADP (ammonium dihydrogen phosphate), KDP (potassium dihydrogen phosphate), KTP (potassium titanyl phosphate), and niobates such as lithium niobate ($LiNbO_3$).

Now, considering that there are two frequencies of light in the crystal (the fundamental and the harmonic), it is quite likely that each will have a slightly different index of refraction. This effect, of course, is responsible for dispersion of light of various wavelengths by a prism. In most crystals, the index of refraction will be higher for the higher frequency component (2ϖ) than for the lower frequency component. A higher index of refraction means that light travels more slowly. When the second harmonic light becomes out of phase with the fundamental component as it passes through the crystal, destructive interference is

caused, which leads to extinction of the second harmonic component. If the crystal is thin enough, both components remain relatively in phase inside the crystal and can coexist. This is somewhat impractical and severely limits the efficiency of the harmonic light generation: a thick crystal is required for high conversion efficiency.

To solve the problem of phase mismatch between the components, the phenomenon of birefringence is exploited. Birefringence is an effect in which the index of refraction depends on the polarization of light passing through the crystal, or in this case, the individual light component passing through the crystal. By tilting the crystal such that the index of refraction for both the fundamental (with one specific polarization) and the harmonic component (with a different polarization) are the same, phase matching can be accomplished. When the direction of propagation of light in the crystal is chosen so that the indexes of refraction for both the fundamental and second harmonic are exactly the same, the crystal is said to be *phase matched*. Phase matching is often accomplished by simply tilting the crystal relative to the optical axis. The second harmonic component produced in the crystal now exits in phase with the fundamental component, so that constructive interference occurs. This has implications for efficiency of generation of second- harmonic light.

Nonlinear interaction is a general term describing the mixing of two light beams to produce a sum or difference of frequencies. The case of second-harmonic generation, the most common application of nonlinear phenomena, is a specific case in which the input components are identical (with the same frequency) and interact to produce a resulting beam with a doubled frequency. Other inte actions are possible, of course, such as the mixing of a fundamental and second-harmonic beam to produce a third-harmonic output, the most common application being the mixing of a 1064-nm fundamental from a YAG laser with a 532-nm frequency- doubled beam to produce a third-harmonic output at 355 nm or even a fourth-harmonic output at 266 nm.

8. Gas Lasers

8.1. Helium–Neon. The HeNe laser, the first gas laser and still one of the most common, normally operates with a familiar red (632.8 nm) beam, but multiple wavelegths are possible, allowing the laser to operate (with suitable optics) at wavelengths in the infrared, orange, yellow, and green. Power output for commercially available HeNe tubes ranges from under 1 mW for a small HeNe tube to just over 100 mW for a large behemoth unit. These lasers typically feature excellent spectral and coherence characteristics, so they are a popular choice for use in holography, where they are still the laser of choice, preferred over semiconductor types. They are also the *de facto* standard laboratory laser, used for a myriad of odd jobs around the lab, including alignment and testing of optical components.

The HeNe laser uses a mixture of very pure helium and neon gases in the approximate ratio of 10:1. Helium, the gas used to furnish the pump level in this four-level system, is in greater abundance. Tube pressures are typically between 1 and 3 Torr, being dependent primarily on the diameter of the plasma tube.

Being a four-level laser with favorable dynamics, HeNe lasers have low thresholds and operate in CW mode. The pump level is supplied from the helium atom, which absorbs energy from the discharge through electron collisions and becomes excited to a level 20.61 eV above the ground state (for all visible transitions; ir transitions at 1.15 and 3.39 μm use different levels). The ULLs in neon (denoted as the $1s^2 2s^2 2p^5 5s^1$ electronic state) are at almost exactly the same energy (20.66 eV) as the excited helium energy level, so that a collision with an excited helium atom will result in the transfer of energy to neon atoms, raising them to an excited state, with the small difference between the two energies being made up by thermal and kinetic energy in the system. From the ULL, they can emit a photon of coherent light (via stimulated emission) to decay to the LLLs (the $1s^2 2s^2 2p^5 3p^1$ electronic state).

Quantum mechanics dictates that this electronic state is actually composed of numerous, closely spaced levels, hence multiple transitions (and wavelengths) are possible in this system. Neon atoms at the LLLs then fall to ground through a depopulation process involving both radiative decay and collisions with the tube walls. This also explains why HeNe lasers invariably use small- diameter (<2-mm) plasma tubes. The entire process, including excitation of the helium atoms, transfer to the ULL in neon, visible lasing transitions, and decay from the LLL is shown in Figure 20.

In the visible spectrum, there are eight lines, ranging from green to red. Of these, several are too close to more powerful lines to be useful (eg, the output at 635 nm is too close to 633 nm to separate with a normal wavelength selector). Commercially, four visible wavelengths of HeNe laser are commonly available, as outlined in Table , along with gain relative to the 632.8-nm red line. Because of the low gain of the green HeNe, special low loss optics are required that are much more critical (and expensive) than red HeNe optics. Where a small red HeNe tube might have an output of up to 20 mW, a comparable green tube would be limited to under 3 mW.

Most commercial HeNe laser tubes are of a one-piece design featuring integral mirrors resembling that shown in Figure 1. This compact design has prealigned optics providing trouble-free operation for most users. Optics never require cleaning or realignment since they are sealed into the tube. Most one-piece tubes are randomly polarized, although some tubes have integral Brewster windows within the tube itself, designed solely to polarize the output beam.

Some lasers, primarily research and lab lasers, feature external optics. Tubes often have Brewster windows to reduce loss, but an alternative is a plane window with antireflective coatings. These coatings are required since losses of up to 8% at an uncoated window would certainly prevent oscillation. Some long (up to 1 m) research lasers also feature wavelength selectors allowing tuning of the laser over multiple wavelengths. These selectors are prisms with an affixed HR and allow tuning by changing the angle of the assembly with respect to the tube axis. Shorter lasers lack the necessary gain to overcome losses imposed by the wavelength selector, so they usually lack these features.

The HeNe tubes can vary in length from <12 cm to >1 m. The smallest tubes easily fit, along with a compact power supply, into handheld bar-code scanners and produce ~1 mW (although for this application semiconductor lasers are now often used). The largest units sport tubes >1 m long and are used for

research and holography purposes, producing ~120 mW of red output. Long tubes also have higher gains, allowing the use of wavelength selectors to lase one of several possible wavelengths (although the gain is still generally too low to allow multiline operation as is commonly done with argon lasers).

A HeNe laser produces an excellent-quality output beam, so it is the laser of choice for holography applications. With a spectral width of 1.5 GHz, common HeNe lasers exhibit a coherence length of ~20cm. This is much better than the average semiconductor laser (which is why semiconductor lasers have been slow to move into holographic applications). When configured for single-mode operation, coherence lengths of >300 m can be achieved (although this is no longer a common "garden variety" HeNe laser). Most HeNe lasers operate in TEM_{00} mode and yield excellent beam quality. Although most standard HeNe lasers are randomly polarized, some have integral Brewster windows (in many cases within the tube itself), giving a polarized output beam.

Until the 1990s, the HeNe laser was the only option for many pointing, alignment, and scanning applications. Not long ago, many supermarket bar code scanners and even some high end handheld scanners used HeNe lasers, as did the original (early 1980s) CD and laser disk players. Today, most applications have migrated toward more compact (and inexpensive) semiconductor laser diodes. However, in some applications where beam quality, beam collimation, and coherence length are important, HeNe lasers are still used extensively. For holography, the HeNe laser is still the laser of choice because of its good coherence length compared to other types of lasers. The same is true for many precision measurement applications that utilize interferometry to measure tiny distances. Access to the gain medium and the ability to use unique cavity configurations (eg, a ring) allows the HeNe laser to be used as the basis for instruments such as laser gyroscopes for aircraft. Finally, the HeNe laser is indeed the de facto standard laboratory laser, with uses ranging from alignment of optics to testing of optical components.

8.2. Ion Lasers. The term *ion laser* refers to a laser in which the lasing energy levels exist in the ionized atom of the species. Although ions are the active species in many solid- state lasers (eg, Cr^{3+} ions in ruby), the term *ion laser* is accepted to apply to gas lasers employing argon or krypton gas that is ionized. Ion lasers are generally high powered lasers (much higher powered than a HeNe laser) emitting in the green-blue region of the spectrum (for argon) or on many lines across the entire spectrum (for krypton), and even in the uv.

As an example of a typical ion laser, consider singly ionized argon (Ar^+), the active species of the most common ion laser by far. Ions are created by discharging a current of up to 40 A through low pressure (0.1 Torr) argon gas. Discharges may be pulsed, as the earliest lasers were, but most ion lasers are CW, so a continuous current of up to 40 A is required, which leads to plasma tube and power supply designs much more complex than the HeNe-ion laser. The Ar^+ ion has numerous energy levels in two bands: nine ULLs centered ~36 eV and two LLLs ~33 eV. Many transitions share an upper or lower energy level. Overall, 10 laser transitions exist in the violet-to-green region of the spectrum.

In the case of argon, the neutral (nonionized) configuration of the atom is $1s^2\ 2s^2\ 2p^6\ 3s^2\ 3p^6$, and when ionized, which requires 15.76 eV of energy, the

ground state for the ion (Ar^+) becomes $1s^2\ 2s^2\ 2p^6\ 3s^2\ 3p^5$. The argon ion can now be promoted to a higher energy state in which one electron assumes a 4*s* state. There are two such states possible with different spins (1/2 and 3/2), and these serve as lower lasing levels. Even higher energy states are possible, in which an electron enters a 4*p* state: there are numerous states here that serve collectively as upper lasing levels (in the visible argon laser spectrum seven tightly clustered ULLs are involved) (5). Four distinct pathways for pump energy have been identified that serve to excite ions to the ULL, including the direct pumping of ions from the neutral ground state to the ULL, a two-step process in which argon atoms are ionized to the ion ground state and later pumped to the ULL from there, and two processes in which the ULL is excited via decay from levels above the ULL. Having made a transition from one of the ULLs to a LLL, the ion decays almost immediately to the ion ground state (still 15.76 eV above the ground state of the neutral argon atom) by a radiative process in which a 74-nm extreme uv photon is emitted. This decay process is fast (a requirement to maintain a large population inversion), but the extreme uv photon created in the process poses problems, since it can damage optical windows and degrade tube materials. Windows must be fabricated from special crystalline quartz to withstand the constant bombardment of extreme-uv radiation. The energy levels involved, transitions, and depopulation process are shown in Figure 21.

Krypton gas may be used in an ion laser as well with various wavelengths, covering the entire visible spectrum from violet to red. Physically, krypton laser tubes are similar to argon tubes with the exception that krypton lasers require a large gas ballast volume since gas is consumed (or rather buried) as the laser operates. Krypton, however, is less efficient than argon, so output powers are lower than those for a comparable argon laser (the most powerful lines of the krypton laser are only one-fifth the power of the most powerful argon lines). The wider range of visible wavelengths available, however (including a powerful red line and a yellow line, both lacking in the argon laser spectrum), make this laser a popular choice for entertainment purposes. Table 2 lists the common visible wavelengths of argon and krypton ion lasers and typical output power for a comparably sized single-line (wavelength-selected) laser using each gas.

It is also possible to doubly or triply ionize a species. Doubly ionized argon (Ar^{2+}), produced by a higher current density than required for Ar^+, produces several powerful lines in the uv region between 275 and 364 nm, so this laser is an important source of continuous uv output. For doubly ionized argon, however, the ULLs are located $\sim$27 eV above the ion (Ar^{2+}) ground state, which is far above the Ar^+ or neutral ground state. The ULLs for the species are hence 71 eV above ground, so enormous amounts of pump power are required for this laser, with tube currents of 70 A common. Also, efficiency is much lower than for the singly ionized species with power output on the uv lines 20-fold less than the expected multiline visible output. Krypton can be doubly ionized as well, but is of even lower efficiency than doubly ionized argon and not commonly available. Other noble (inert) gases can be utilized in an ion laser as well although most are not found commercially. Xenon ion lasers, emitting on several lines in the blue, green, and yellow regions of the spectrum, have been demonstrated to operate in pulsed mode. Neon, another noble gas, is more efficient than doubly ionized argon at producing uv output, but few manufacturers offer lasers

of this type. Many other gases will operate as ion lasers, such as oxygen and nitrogen. Most were discovered by accident when they were present as impurities in laser tubes intended to lase argon or krypton (in many cases they were left over from an incomplete cleaning of the tubes and electrodes). Note, however, that the fact that the argon ion lases at all was discovered by accident in 1964 when argon was tried as a buffer gas for a mercury-vapor ion laser (the first ion laser discovered) (6).

Even many metals operate as an ion laser when vaporized. Mercury vapor was the first ion to lase and produces pulsed-laser output at 615 nm (in the red) and 568 nm (in the green). Other metals that lase as an ion include selenium, cadmium, silver, gold, nickel, and copper although many require elevated temperatures to increase the vapor pressure, and hence the number of atoms available in the plasma. Of the 41 elements that will lase as an ion, only cadmium (other than the noble gases argon, krypton) became a commercially available laser. The cadmium-ion laser, called the helium–cadmium (HeCd) laser, operates much more like a HeNe laser with much lower currents and temperatures than argon or krypton ion lasers. It is thus generally not classed as an ion laser (since it is so different in structure from other ion lasers) despite the fact that the active lasing species is indeed ionic. The HeCd lasers produce output in the uv at 325 nm, blue at 441.6 nm (the strongest output), and on four other lines in the green and red. These are commercially available from a number of manufacturers and are an important source of continuous uv light.

In any ion laser, a high current density in the gas discharge is required to excite the ions to the level required. High currents (up to 40 A for an Ar^+ laser) and a small-diameter bore are used in plasma tubes to increase current density since the output power increases as the square of current density. Doubly ionized species (Ar^{2+}) require even current densities. These high current densities create a host of problems for designers of ion laser tubes since the plasma runs at an extremely hot temperature (over 5000 K) and can easily erode tube structures such as electrodes and tube materials. Glass melts well below this point, so there are a limited number of materials available from which a plasma tube can be constructed. These materials include beryllium oxide (a ceramic) and a few high melting point (refractory) metals, including tungsten and graphite. Even with such exotic materials and construction techniques (such as inserted disks of refractory metal), the energetic plasma of a large ion laser (one with a discharge current of perhaps 30–40 A) can easily erode and destroy the tube material on contact. For this reason, magnetic confinement is invariably employed with large plasma tubes. The magnet is coaxial with the laser tube and is water cooled along with the plasma tube itself. Magnetic fields of $\sim$1200 G are employed with visible lasers, which serves to confine the discharge to the center of the plasma tube, away from the tube walls (uv lasers, with higher currents, require even higher magnetic fields). By squeezing the plasma to the center of the tube, current density is also increased, increasing the probability of pumping an ion to a high energy state, and hence increasing gain (and ultimately, output power).

To prevent damage due to ion bombardment, heated cathodes are required. By heating the cathode, electrons are emitted from the surface, which serves to reduce the voltage drop associated with the energy required to pull electrons off

the surface of the cathode. With a lower voltage drop, the power dissipated at the cathode drops; the tube runs at lower voltage and there is less damage to the cathode during operation. The cathode must be preheated before initiating the discharge and most lasers have a built-in delay, which activates the heater for 45 s before initiating the discharge.

As an ion laser operates, it consumes gas through a process called sputtering in which energetic (and heavy) ions strike the internal tube material (both the tube and the electrodes) and are buried into the surface, effectively reducing gas pressure in the tube. As tube pressure is reduced, tube voltage drops, damage to the cathode increases, and laser output drops until it is finally extinguished. Ion lasers hence require a gas replenishment system to keep the tube pressure within a usable range. For a simple ion laser tube such as a small air-cooled type, this gas is usually replenished by means of a large ballast volume connected to the tube. This serves as a buffer to supply extra gas to the tube as required. An external ballast volume (in the form of a tank) is almost always found on krypton laser tubes since the output of these lasers is considerably more sensitive to pressure than is that of argon lasers. On many small argon laser tubes the volume of the cathode and anode is sufficient to permit them to serve as ballast volumes, but with larger ion lasers, an active-replenishment system is used that injects gas into the tube as required. The system monitors the tube pressure either by means of a vacuum gauge directly attached to the tube or by monitoring the voltage across the tube (which drops as the pressure drops). When gas is required, a system of electromagnetic gas valves releases a small, controlled amount of gas from a high pressure reservoir into the tube.

Some small ion lasers have internal optics (notably, Cyonics/Uniphase argon tubes) and resemble overgrown HeNe tubes (with extensive cooling fins), but most lasers, even many small air-cooled tubes, feature external optics. In almost all cases the laser tube has two Brewster windows protruding from the ends of the tube (on quartz stems sealed to the laser tube), so most ion lasers have a polarized output. Like a HeNe, ion lasers have very low gain, so low loss windows are necessary for operation. Cavities are frequently of the plane-spherical type, the OC being spherical and having a radius of curvature slightly longer than the cavity length. This arrangement allows the use of interchangeable flat optics at the HR end. For multiline use, a broadband reflector can be installed in the HR position. Selective reflectors may also be used to allow only certain wavelengths to oscillate, as is frequently done with krypton lasers to select only the red line and argon to select the powerful green or blue line. Wavelength selectors using a prism and an HR are also an option for single-line operation, and most tunable lasers also allow the addition of an intracavity etalon, allowing single-frequency, narrow-spectral-width operation. To reduce losses at the mirrors, mirrors are made from multiple layers of thin dielectric films. Since the output of argon lasers frequently exhibit a spectral width of 5 GHz, intracavity etalons are often employed allowing single-mode operation.

Practical ion lasers range in size from small, air-cooled units to large, water-cooled lasers with outputs in the tens of watts range. Small units feature a ceramic plasma tube (beryllium oxide) shrouded in a mass of cooling fins cooled with forced air from a large blower. Most commercially available air-cooled lasers are packaged in a small, compact housing with an umbilical cable connecting the

laser head (containing the plasma tube, optics, and some electronics) to the main power supply. Large-frame ion lasers usually feature a water-cooled ceramic plasma tube and magnet. With the large magnet coaxial with the plasma tube the only parts of the tube that are visible are the heated cathode and the Brewster windows mounted on quartz stems protruding from the tube. Frequently, an aperture (for mode limiting) is mounted inside the cavity in front of the OC. The power supply, along with the laser tube–magnet combination, is usually water cooled: power supplies for ion lasers contain current regulation circuitry, often a series passbank, which must be cooled adequately.

Ion lasers, especially the argon laser, have numerous applications where higher power than that obtainable from HeNe lasers, which are limited to ~100 mW, or green-blue light is required. It has become a workhorse lab laser used for all sorts of measurement and excitation purposes, including chemical fluorescence studies and pumping CW dye and solid-state lasers. In the forensics lab, fingerprints are found to fluoresce when illuminated by argon laser light. In medical applications, the wavelength of the argon laser is absorbed readily by red blood cells, but passes easily through the liquid filling the eyeball, so it is ideal for certain types of opthalmological surgery, including welding detached retinas. Other medical lab applications include flow cytometry and DNA sequencing, where the blue and green wavelengths are important to the application.

The powerful beams are also useful for high speed copiers and printers, where they can be used in the same manner as an ir semiconductor laser is for small laser printers. The more energetic photons and more powerful beam allow faster printing than a semiconductor laser does (many small air-cooled argon lasers are designed precisely for this purpose). In the uv region, ion lasers can be used to expose tiny features on high resolution printing plates and films (useful since beam quality is excellent and uv wavelengths can be focused to a much smaller spot than visible wavelengths). Other industrial purposes include the mastering of CDs and DVDs.

The krypton laser, which can be designed for white-light output, is a favorite for laser light shows. Multiple wavelengths of the laser can be split using a prism and directed independently, or a single wavelength can be selected from the output by a high speed AO modulator called a PCAOM. Beams are then scanned using mirrors attached to precision galvanometer movements to create patterns. As in a television set, fast scanning ensures that the eye renders the image continuous.

For years the argon laser remained the only powerful CW source of green and blue laser light available, but at present the market is turning toward simpler (and smaller) frequency-doubled YAG lasers, which feature an output at 532 nm in the green, or frequency-tripled YAGs, which have a uv output at 355 nm. For specific applications requiring the 488-nm blue wavelength, newly developed solid-state lasers with precisely this wavelength threaten the dominance of the argon laser. Regardless of the appeal of solid-state lasers, the argon ion will remain an important source of blue-green light for years to come, especially since it features a multiwavelength output that cannot be achieved (easily, at least) with a solid-state laser.

8.3. Molecular Nitrogen Lasers. Most transitions from gas discharges (eg, those in a HeNe or argon ion gas laser) have very well-defined wavelengths.

These are *electronic transitions*, which involve an energy change in only electrons of the species. Precise and well-defined energy states characterize these transitions, which are responsible for most visible gas laser transitions as well as many of those in the ultraviolet. Other energy levels possible are *vibrational* and *rotational levels*, brought on, in molecules, by various supported modes of movements of individual atoms relative to each other. In the simplest case, a diatomic molecule such as nitrogen (N_2) or hydrogen (H_2) is composed of two atoms, bonded covalently, which are free to vibrate only in certain allowed ways. This molecular bond is not rigid, but rather flexible and may be stretched in various ways as the atoms move. For modeling a diatomic molecule as two weights attached by a spring (the covalent chemical bond between the two atoms), kinematics may be used to predict the behavior of this system, which resembles a simple harmonic oscillator. Transitions can take place between two of these vibrational levels, resulting in a purely vibrational transition with energies corresponding to transitions in the ir region.

It is also possible to combine vibrational states such as those described with electronic states to form hybrid *vibronic levels* and transitions. It is these type of transitions (from one vibronic energy level to another) that are responsible for the molecular nitrogen laser output at 337.1 nm in the ultraviolet. In nitrogen, the major separation of energy states is caused by a change in the electron energy of the atoms themselves and minor states caused by vibrations of the N_2 molecule itself. The combined effect is that essentially all energy levels consist of a series of closely spaced levels essentially forming a band. As expected, the lasing transition is actually a series of transitions closely spaced to render a primary output centered at 337.1 nm and having a bandwidth of 0.1 nm. These transitions are the result of a change in both electronic and vibrational states of the diatomic molecule.

Excitation of the nitrogen molecule is accomplished by collision of high energy electrons in a gas discharge. A massive electrical pulse, with currents on the order of thousands of amperes, initiates the process by creating a large flux of electrons with high energies in the laser tube. Electrons strike the nitrogen molecules and excite them to a high vibronic energy state. In the excitation process for this laser, nitrogen molecules do not dissociate into an atomic species, but rather remain as an excited molecular species.

The energy levels of the nitrogen molecule as they apply to this laser are outlined in Figure 22. Each level shown is actually a series of vibrational levels dependent on internuclear separation. The laser begins when a nitrogen molecule is excited by direct collision with electrons in the discharge to enter the ULL, termed the $C^3\ \Pi$ energy band of the molecule. From the ULL, the molecule falls to the LLL (termed the $B^3\ \Pi$ band), emitting a photon of uv light in the process. Transitions in a normal nitrogen laser operating at 337.1 nm are in the 0–0 band, in which the only levels involved are those with the lowest vibrational state ($\nu = 0$). Other transitions are possible in the 0–1 band, where the upper lasing level is the same lowest state ($\nu = 0$), but the lower lasing level involved is the $\nu = 1$ vibrational state. The transition is a smaller jump than the 337.1-nm transition, and hence has lower energy with a wavelength of 357.7 nm, the difference between the $\nu = 1$ and $\nu = 0$ levels being $\sim$0.2 eV.

After emitting a photon in the lasing transition, the molecule then falls to a metastable state and finally, to the ground state. Although not a four-level laser in the proper sense (the system lacks a pump level since the ULL is pumped by direct electron collision), it is certainly not a three-level system since it does feature a distinct LLL.

The energy levels for molecular nitrogen are far from favorable for lasing, with the upper lasing level having a much shorter lifetime than the lower lasing level, but to obtain lasing action from molecular nitrogen there must be a higher population of molecules at the ULL than at the lower level. The level lifetime situation makes CW laser action impossible with this species, but pulsed-laser action is still possible provided that the laser mechanism can pump energy preferentially and quickly into the upper lasing level to generate an inversion. In a nitrogen laser, this is done by a massive electrical pulse where electron collisions cause the preferential population of the upper energy band first. (Were it not for this effect of being able to pump the upper energy level first, this laser would not work at all.) After $\sim$20 ns, this energy will decay to the lower level, population inversion is lost, and lasing action will quickly cease. To make matters worse, nitrogen molecules at the lower state absorb uv light strongly, so gas in the laser channel quickly absorbs laser energy. On average, the pulse length of the nitrogen laser is limited to $<$20 ns (the point where population inversion is no longer possible since one-half of the molecules in the upper energy state have decayed to the lower state).

Nitrogen lasers operate without mirrors, a type of laser called *superradiant*. The gain medium exhibits such high gain that spontaneous emission from nitrogen molecules at the end of the laser are amplified by the same group of nitrogen molecules in the tube, producing a usable output pulse in a relatively short path length. Gain for a nitrogen laser is on the order of 40–50 dB/m or more (7), depending on the specific laser. Even at the lower figure of 40 dB/m, light is amplified by a factor of 10,000 for every meter of travel through the laser tube! This is an enormously high gain for a laser that also serves to prohibit tuning of the laser through any of the 61 lines in the gain band, since they will lase without any cavity mirrors, and hence a wavelength selector is ineffective. Although not required for laser action, a single high gain reflector is frequently installed in a nitrogen laser tube as well as a 4% reflective (one uncoated surface of a window) output coupler.

The basic requirement for a practical nitrogen laser is to supply a massive electrical current (ie, a huge quantity of electrons) with a fast rise time and short pulse length to excite the gas. To achieve this, most nitrogen lasers use an electrical configuration called a *Blumlein configuration*, which generates a massive overvoltage of the laser channel (and subsequent large current through the lasing gas) with a rise time of nanoseconds. A Blumlein configuration uses two capacitors essentially in parallel separated by the laser channel itself. Both capacitors charge simultaneously and the laser is then fired by discharging one capacitor through a spark gap or thyratron. By designing the entire laser as a transmission line, very high electrical current, in many cases thousands of amperes, can be generated through the laser channel in nanoseconds.

Despite the fact that the output of the nitrogen laser is optically poor with a spectral width of 0.1 nm this laser remains as a useful source of uv light for

biology and chemistry labs. Given their intense uv output and short pulse length, nitrogen lasers make excellent pump sources for pumping dye lasers: The short pulse length does not allow triplet production in the dye, so the conversion efficiency of the laser-pumped dye laser is very high. Nitrogen lasers are also useful for exciting fluorescence in substances other than laser dyes, allowing studies of these molecules. As an excitation source, this laser is commonly used to produce ions for time-of-flight spectroscopy. In the MALDI (matrix-assisted laser desorption–ionization) technique, the laser vaporizes and ionizes nonvolatile biological samples, which are then analyzed by a time-of-flight mass spectrometer and detected based on their mass-to-charge ratio. Nitrogen lasers may also be used for small microcutting procedures on individual biological cells or for trimming thin films for the semiconductor industry. With a typical pulse energy of <1 mJ and a repetition rate <100 Hz, nitrogen lasers offer a low cost source of intense uv light.

8.4. Excimer Lasers. Excimer lasers (8),(9) produce intense pulsed output in the uv. The excimer is unique because the lasing molecule consists of a halogen and an inert gas. This molecule is transient and exists for only an instant. Like the nitrogen laser, a fast, high current discharge is required to produce the excimer molecule, but excimer lasers are considerably more complex since they operate at high pressures and one of the active gases is highly toxic and corrosive. Modern excimer lasers produce uv pulses with energy ranging from 0.1 to 1 J and can (for a large industrial laser) produce these pulses at a rate of >300/s. Pulses are fast, with a full width at half-maximum (FWHM) of 10–30 ns. Fast pulses combined with high peak powers (a 0.5-J pulse with a 20-ns width corresponds to a power of 25 MW) serve to ablate target material without heating the surrounding material, a desirable effect for many processing purposes.

Energy levels in an excimer laser are defined by the state of the atomic components. When unbound, the energy of the system depends purely on the separation between the individual atoms; as the atoms move closer together, energy rises. This is illustrated by the lower curve in Figure 23. The lower energy level in an excimer system is defined as a separation of the halogen and inert gas atoms. This is the normal state for an inert species such as argon, krypton, or xenon, which will not, under ordinary circumstances, form compounds with any other atom. The upper energy state is formed when the inert atom and the halogen form an excited dimer (or *excimer*) molecule. This is a temporary state that forms only for an instant when the excited atoms combine to form a molecule. The energy of the excimer molecule is much higher than that of the unbound individual atoms and also depends on the interatomic spacing in a similar manner to that of any other molecule, such as nitrogen. As with the nitrogen molecule example, numerous vibrational levels are possible. Collectively, these closely spaced levels form a band serving as the upper lasing level.

The laser tube is filled with a mixture of an inert gas and a halogen such as fluorine or chlorine. When an enormous, fast-rising electric current (generated by a mechanism similar to that of a nitrogen laser) passes through the mixture, gas atoms become extremely excited and excimer molecules form. This forms the ULL for the species, which has a short lifetime, on the order of tens of nanoseconds. In this respect, the excimer laser resembles that of a nitrogen laser and a

fast discharge is required to generate the excimer molecules effectively and ensure an inversion before the lifetime of the species has elapsed.

Various excimer species are outlined in Table 3, which also lists the relative power output relative to KrF, the most powerful excimer laser. Although KrF produces the most powerful output, other gas mixtures, such as XeCl, are popular for use in excimer lasers. Shortcomings of the KrF laser include the output wavelength of 249 nm, which is absorbed readily by air, and the extremely corrosive nature of fluorine, which shortens the useful life of the gas mix in this laser. Worse yet from a beam-management perspective is ArF, which produces a wavelength so short that it produces ozone gas from atmospheric oxygen as it passes through air. When using ArF, beam paths must be enclosed and flushed with dry nitrogen, helium, or argon. On the other hand, XeCl has a longer wavelength, allowing better transmission in air and the use of considerably cheaper optics. The gas mixture also has a much longer useful lifetime (up to 10 times longer). The useful lifetime of lasing gases may also be extended by using a cryogenic gas processor in which the lasing gas mixture is passed through coils immersed in liquid nitrogen to trap impurities (some of which adsorb from internal components and others produced by the hostile environment inside the laser) in the gas mixture. A cryogenic gas processor is a standard option on most commercial excimer lasers and most large excimer lasers (ie, intended for prolonged use) have gas ports allowing connection of a gas processor. The cost of a gas fill is quite pricey (for many lasers, well over $100/fill), so gas processors pay back quickly for many industrial users.

The actual gas mixture employed consists of a small quantity of fluorine ($< 0.5\%$) combined with a moderate amount of inert gas, with the balance helium or neon. Pure fluorine is such a corrosive and toxic gas to handle that it is not available in pure form but only in diluted form, as 5% fluorine and 95% neon gas (or possibly, helium).

With high average powers (commonly >100 W for many commercial lasers) and an output in the uv region of the spectrum, excimer lasers are useful for many applications, ranging from dye laser pumping to cutting and materials processing applications. The largest commercial application for excimers are use in eye surgery to correct the shape of the cornea to reduce the need for corrective lenses (commonly known as *lasik surgery*). For these applications, ArF with an output wavelength of 193 nm and a pulse energy of <0.5 J is used. The 193-nm wavelength is used since it is readily absorbed by tissue on the surface of the cornea, which is ablated without producing significant heat to damage surrounding tissue. Since the laser light is absorbed readily by tissue on the front of the eye, no laser energy is transmitted to the retina, sparing it damage. By 2003, over two million such corrective surgeries had been performed using the excimer laser.

The largest current industrial use for excimer lasers is photolithography for semiconductor manufacturing. The excimer laser replaces the standard mercury lamps used for years for this purpose, allowing finer features to be resolved. The KrF laser has become the workhorse of the semiconductor industry for this purpose. The ArF excimers are also used industrially for glass- marking applications, since the beam is readily absorbed by glass (which is transparent to longer wavelengths). Along these lines, ArF (and sometimes KrF) is used to manufacture fiber Bragg gratings for optical fiber communications. During manufacture,

a pattern is cut into a single-mode fiber using a phase grating. Other applications for excimer lasers include wire stripping (especially for ultrafine wires used in the microelectronics industry), surface-mount component marking, drilling inkjet printer nozzle holes, and marking wires.

8.5. Carbon Dioxide Lasers. Like the nitrogen and excimer, the carbon dioxide laser is a molecular laser, however, the more complex carbon dioxide molecule, as depicted in the model of Figure 24, has many more degrees of freedom than a simple diatomic molecule. This molecule features a single carbon atom chemically bonded to two oxygen atoms with double covalent bonds acting as springs leading to vibrational modes each having unique levels. Several modes are shown in the figure, including an asymmetric stretch (where one bonds stretches while the other compresses), symmetric stretching (where both bonds stretch or compress simultaneously), and a bending motion. Energies for such vibrational levels are quite low (ie, the difference between allowed energy states is quite minute), and hence transitions between allowed vibrational modes correspond to the ir and far-ir regions of the spectrum (eg, the carbon dioxide laser output, generated primarily by transitions between vibrational modes, is in the ir at 10.6 μm, which corresponds to an energy of 0.117 eV; unlike the nitrogen laser, electronic levels are not involved here).

Each molecular bond can be considered as a spring between two weights, so it is easy to imagine the oscillator (in a classical sense) as being confined to oscillating on certain discrete frequencies. It could oscillate at its natural resonant frequency ν_0 or at twice this frequency, but not at a noninteger multiple such as 1.5 times this frequency, since this would not allow a standing wave in the spring (and would hence quickly be dampened). This is the quantized nature of atomic energy levels and applies to these types of levels as well. Oscillating at $2\nu_0$, we expect the energy level for this vibration to be twice as large as that of ν_0. In Figure 25, the energy levels for carbon dioxide as well as the corresponding vibrational modes are outlined. Lasing transitions (those responsible for the actual laser output) are shown in the figure between a higher energy allowed vibrational state and two lower energy allowed vibrational states.

The energy levels for the CO_2 molecule are assigned a three-number designation that describes the vibrational mode as well as the frequency. In the case of the bending mode, two modes (010 and 020) appear in the figure. Mode 010 corresponds to a frequency of ν_0, while 020 corresponds to a frequency of $2\ \nu_0$, which has twice the energy of mode 010. Energy is provided to the CO_2 molecule from nitrogen in a four-level scheme. The *pump energy level* in nitrogen is also a vibrational energy level (logical since electronic levels would correspond to a much higher energy). Nitrogen (N_2) is composed of two nitrogen atoms held together by a bond and vibrates only in certain allowed modes. The energy of one of these purely vibrational modes (not to be confused with the vibronic modes mentioned earlier) corresponds closely with the upper lasing level for the CO_2 laser, so nitrogen acts as a pump.

As well as vibrational energies, which usually correspond to transitions in the ir, there are *rotational energies* involved in this system. These energies originate from various modes in which an outer atom rotates about a central atom: In the case of CO_2, an oxygen atom literally rotates around the axis of the carbon and other oxygen. As one might imagine, there is a great deal of freedom to rotate

in a given molecule, so there are many more possibilities for rotational modes in a large molecule than there are vibrational (although it, too, is constrained and even rotational levels are quantized). With many possibilities, rotational energy levels appear in bands of tightly clumped energy levels. Purely rotational energy levels are possible as well; these correspond to transitions in the far-ir and are seen in many organic molecules. In the CO_2 laser, vibrational energies combine with rotational energies to turn each vibrational level into a tight group of energy levels, as depicted in Figure 25. The result is that the two transitions of the CO_2 laser centered at 9.6 and 10.6 μm are actually a series of closely spaced spectral lines spanning the region from ~9.2 μm to almost 11 μm. One specific line can be isolated using wavelength-selective feedback (such as a diffraction grating) in such a laser making this laser useful as a source for ir spectroscopy.

Water cooling is required for most CO_2 lasers not just to remove discharge heat, but also to reduce the thermal population of the lower energy levels, which are very close to ground level. The output power of most CO_2 lasers is quite sensitive to plasma temperature, and a blocked or restricted cooling water line can easily result in a decrease in output power. To this end, many lasers have thermal sensors on the water cooling jacket of the tube, designed to shut down the laser should the temperature reach 40 to 50°C. While the plasma tube would probably tolerate much higher temperatures, laser output would drop drastically at these temperatures.

The relatively long ir wavelength of the CO_2 laser is readily absorbed by most organic materials (such as plastics), as well as glass, water molecules, and many common materials. For cutting materials such as cotton (used in making jeans), the CO_2 laser is ideal. It is also used in surgical applications since the wavelength is readily absorbed by flesh vaporizing it; the heat also serves to cauterize the cut, thus reducing bleeding.

Most metals are somewhat reflective at this long wavelength, so higher powers are required to ensure that enough energy is absorbed to cause vaporization of the material (easily provided, though, by these types of lasers). Cuts made into metals are hence somewhat rougher than those made with other lasers (eg, the YAG), but the high average power available from the CO_2 laser, which no other commercially available laser can generate, usually means that higher speed operations are possible, so for many large-scale cutting applications, this is the dominant laser. This laser is used extensively to cut stainless steel and titanium, which are difficult to cut by any other means. For drilling applications, the YAG is usually preferred for controllability.

9. Other Lasers

9.1. Solid-State Lasers. The oldest technology, but one reborn recently and becoming increasingly important, is that of the optically pumped solid-state laser. Solid-state lasers (not to be confused with semiconductor lasers) consist of a crystal of glass-like material doped with a small concentration of a lasing ion such as chromium (in the case of ruby) or neodymium (in the case of YAG). The ruby laser, the first ever, was a reasonably simple structure with integral mirrors on a rod of synthetic ruby (chromium ions embedded into a host crystal

of Al_2O_3) pumped with a helical flashlamp. For many solid-state lasers, the technology has not changed much, however, in recent years more efficient materials with lower pumping thresholds are used and compact solid-state lasers have been developed that are pumped by semiconductor laser diodes instead of lamps. Many solid-state lasers have integral harmonic generator crystals to produce visible and even uv light. While ruby lasers continue to be used in a few niche markets, most modern solid-state lasers use more efficient neodymium-doped crystals such as Nd:YAG or Nd:YVO_4. There are other solid-state materials of importance as well such as erbium, which is used in fiber amplifiers for communications systems, holmium, and titanium (useful in a tunable solid-state laser).

The very first laser discovered (by Theodore Maiman in 1960), the ruby laser, is a three-level system having a high pumping threshold and requiring high pump energies. With the development of more efficient lasers, many of which are capable of CW operation, the demand for ruby lasers has diminished steadily, however, it is still a useful source of intense pulsed red light for applications such as holography: The ability to Q-switch this laser results in fast, intense pulses of red light that cannot be generated by other means.

Ruby is synthetically grown aluminum oxide (Al_2O_3) doped with chromium ions (Cr^{3+}) at a concentration of $\sim$0.05%. The ruby laser is a three-level system and as such exhibits high pumping thresholds. The dynamics of ruby are poor for lasing, however, broad absorption bands and a relatively long upper lasing level lifetime allow ruby to operate in pulsed mode in which inversion is achieved only temporarily. The energy levels of ruby are outlined in Figure 26.

Broad absorption bands in the violet and green portions of the spectrum absorb light, usually from a xenon flashlamp, pumping chromium ions (Cr^{3+}) to the pump levels. Pump levels have very short lifetimes ($\sim$1 μs) and a fast decay occurs from those levels to the upper lasing level that has a much longer lifetime of 4.3 ms. From there, ions decay to ground state emitting a photon of 694.3 nm light in the process. The long lifetime of the upper lasing level allows ruby to store energy in that level making lasing possible (as a pulse) and also allowing Q-switching of the laser to produce massive pulses.

Most of the ruby lasers built nowadays are purpose-built and so many have special optics for a specific purpose. Double-pulse lasers used for holography, eg, often take extreme measures to ensure beam quality. Such lasers may feature a "standard" HR with an etalon for an OC that reflects only certain wavelengths, those wavelengths separated by the FSR of the etalon. This scheme allows single-frequency operation of the otherwise spectrally wide ruby output, and hence a longer coherence length required for holography.

Often a MOPA (Master Oscillator Power Amplifier) configuration is used with two optically pumped rods in which one rod is used as an oscillator and a second as an amplifier. The oscillator is a complete laser with HR, OC, Q-switch, and intracavity optics as required. The oscillator usually produces as "clean" a beam as possible, which then passes through an amplifier to increase the output power by up to 10 times that of the oscillator output. Many times the amplifier has a longer rod and more pump power than the oscillator. The optical train of a relatively complex ruby laser is depicted in Figure 27.

Once the "only game in town", the dominance of the ruby laser has diminished in years following the discovery of more efficient lasers such as the YAG.

The primary applications for ruby lasers nowadays are research lasers and as sources for holography. The double-pulse ruby laser, eg, is used to record deformation of a test material by using each of the two closely spaced pulses to record a holographic image. Any deformation or movement is recorded as an interference pattern between the two images.

Similar in structure to the ruby laser is the more common *neodymium* laser in which the active lasing ion is the rare-earth metal Nd^{3+}. Originally used as a replacement for the ruby in many applications, YAG lasers feature much higher efficiencies than ruby (it is a four-level system), have lower pumping thresholds, and can oscillate in CW mode. Traditionally pumped by either a flashlamp or, more commonly, a CW arc lamp, these lasers are most efficiently pumped by semiconductor diode lasers. Having separate cavity optics, SHG crystals can be inserted into the cavity to produce powerful output in the green: for many applications, replacing the argon laser as a powerful source of green light. The laser can be Q-switched making it useful for many materials-processing applications.

The most common host crystal is YAG, but other host materials such as Vanadate (YVO_4) or glass may also be used. The wavelength of the resulting laser beam depends on the host material itself, which modifies the energy levels of the neodymium ion embedded in it. Common host materials and resulting lasing wavelengths are listed in Table 4. While not an extensive list (many other materials exist or are under development) it should give the reader an idea of the types of hosts that can be used and the variations in wavelength (which are minimal and all emit in the near-ir). Of all the materials listed, YAG is the most common material, especially for medium to high power units with vanadate being the favored material for low power (<1 W) and compact solid-state lasers.

Most YAG lasers, especially those used industrially or in the laboratory, feature separate optics with all components mounted on a rail for stability. The optics for YAG lasers are usually straightforward consisting of two mirrors of which one or both are slightly spherical. The YAG lasers frequently include a Q-switch allowing the production of fast, intense pulses (many applications depend on these type of pulses). Q-switches are simply attached to the mounting rail between the rod and optics and most have adjustment screws allowing alignment with the intracavity beam. Q-switches are usually the AO type using inexpensive quartz or similar glass (which is quite transparent at this wavelength). Another popular option is a second (or third) harmonic generator, generally a simple nonlinear crystal in a holder that also attaches to the rail.

YAG lasers are a workhorse for many applications involving cutting, drilling, and trimming. In competition with the carbon dioxide laser for many applications, the short pulse length possible with a Q-switched YAG laser makes it ideal for many applications where the CW carbon dioxide laser is not optimal. Most YAG lasers for materials processing have Q-switches already installed at the factory. For drilling applications, especially in metals, the fast pulse ablates materials without creating heat in the substrate being drilled. For semiconductor processing, the fast pulse allows a tiny amount of resistor material to be obliterated without heating the surrounding material, and hence affecting the properties of the material. The YAG lasers also work well for many marking applications. In the entertainment industry, frequency-doubled YAG lasers

have been used for numerous laser-light displays, especially high power applications, such as cloud-writing, where they offer an alternative to argon-ion lasers.

9.2. Semiconductor Lasers. No laser has gained such widespread applications as the semiconductor (diode) laser. Found in applications ranging from laser pointers to DVD players, these tiny, efficient lasers have made possible many of the optical devices we take for granted. Consider that the very first CD players (built in the early 1980s) used a HeNe gas laser. Such an arrangement is hardly portable and does not lend itself well to, say, an in-dash player in a car. As well as being tiny, these devices are also inexpensive and require only a simple power supply to operate.

Like any semiconductor diode, a laser diode is based on a p–n junction manufactured by combining two materials such as gallium–arsenide: one material doped to have excess holes and the other doped doped to have excess electrons. A voltage develops across this junction that prevents electrons in the conduction band of the n-type material from diffusing across the barrier and combining with holes in the p-type material. When a voltage equal to this potential is applied across the device, current flows and electrons recombine with holes, producing photons in the process. Charge carriers (holes and electrons) in the energy bands in the material (conduction and valence bands) can lie anywhere in these bands, so a range of wavelengths are possible, with the longest wavelengths corresponding to the bandgap energy.

The simplest (and oldest) structure for a laser diode is the *homojunction laser diode*, which uses a single junction. These are fabricated of a single junction between two direct-bandgap materials of the same type, one p- and one n-type, called a *homojunction* since both materials are of the same type. Light is emitted by electron–hole pair recombinations in the thin active region formed by the junction of the two materials (the depletion region). Usually, gallium arsenide (GaAs) is used, with each part of the device doped slightly differently: one part with an electron donor and one part with an electron acceptor. Mirrors for the laser cavity are fabricated simply by cleaving the crystal at right angles to the laser axis. Having an index of refraction of 3.7, the reflectivity of each mirror (simply the interface between air and GaAs) may be calculated to be 33%. This represents a large loss in the cavity; however, most semiconductor laser materials have ample gain to allow such a simple configuration. Improved performance may be achieved by fabricating a single dielectric mirror, composed of alternating quarter-wavelength-thick layers of high and low index-of-refraction materials, at the HR end of the laser diode. Improved mirrors are used on almost all modern laser diodes. A diagram of a simple homojunction structure is shown in Figure 28. Homojunction lasers are characterized by large threshold currents with a typical device requiring tens of amperes to lase. Such currents prohibit continuous operation at room temperature, so CW homojunction devices require cryogenic cooling, making them impractical for many applications.

Improvements in the structure of the laser may be made by confining the intracavity beam in a dielectric waveguide structure formed from the semiconductor material itself. Such a structure requires two interfaces of different indexes of refraction, one on top and one below the active region, so two junctions are formed in what is called a *heterostructure laser diode*, or in this case, a *double heterostructure*, since there are two confining interfaces. To obtain different

indexes of refraction, two different materials are required. The gallium–arsenide is generally used as the higher index material and aluminum–gallium–arsenide (AlGaAs) as the lower index material. As depicted in Figure 29, AlGaAs is doped to form p- and n-type materials that essentially have identical indexes of refraction given that dopant concentration is small. Between layers of these materials, GaAs is sandwiched as the active-region material, from which laser light is emitted. Differences between the indexes of refraction occurring at each interface form a reflector that confines light inside the GaAs layer, which drastically improves efficiency, and more important, lowers threshold current for the device by increasing gain. The active region (GaAs) is typically only 0.1 μm in thickness.

Usually, a stripe contact is used on the top of the structure to make an electrical connection to the device. This further limits the area of the active region in the GaAs laser (since current is not spread out over the top surface area of the entire structure), which serves to increase current density and further lower threshold current. In a real laser diode of this type, more than three layers are generally required, and a layer that serves as an electrical interface between metal contacts and each AlGaAs layer is usually employed. Double-heterojunction laser diodes commonly operate at room temperatures with low threshold currents, in the tens-of-milliamperes range.

The double-heterostructure arrangement confines intracavity light in only one direction (top and bottom) of the GaAs layer, but a further improvement in performance can be made by manufacturing the device so that a confining dielectric interface exists on all four sides of the active region in a *buried heterostructure laser*. In such an arrangement, the entire stack of three layers of a typical heterojunction laser (p-type AlGaAs, GaAs active region, and n-type AlGaAs) is confined on each side by an n-type AlGaAs layer. The interface between the GaAs material in the active region, and this lower index-of-refraction material on the sides of the active layer, serves to further confine light in the laser cavity. Such lasers are often called *index guided*, since light inside the cavity is guided in a manner similar to that of an index-graded fiber optic.

One of the newest structures for semiconductor lasers is the vertical cavity surface-emitting laser (the VCSEL). In a laser of this type, light is not emitted from the edge of the device, but rather through the entire top layer of the semiconductor crystal itself. Like a conventional laser diode, VCSELs features a thin active layer (100–200 nm), but whereas the gain length of a conventional diode is 200–500 μm (the length of the structure), the gain length in a VCSEL is the length of the active layer. Resonator optics are fabricated above and below the semiconductor crystal. With a short active layer and low gain, cavity optics must be fabricated from multiple layers of dielectrics—alternating quarter-wavelength-thick layers of high and low index-of-refraction materials-for high reflectivity. Current in the device flows along the optical axis through electrodes on the top and bottom of the device instead of perpendicular to it.

While edge-emitting laser diodes produce an elliptically shaped output beam that has high divergence, requiring an external lens to collimate it into a usable beam, a VCSEL produces a round beam of much higher quality so the output from the VCSEL is preferred for coupling to a fiber. In addition to a better beam shape, VCSELs feature single longitudinal mode operation with a narrow spectral linewidth.

Finally, it is possible to *optically pump* some semiconductor materials. One design uses an 808-nm pump diode to optically pump a 946-nm semiconductor laser that is, in turn, frequency doubled to produce 488-nm light. The arrangement, VECSEL (for vertical external cavity surface-emitting laser) has external optics allowing inclusion of a nonlinear crystal inside the cavity to accomplish frequency doubling. This particular laser, which has more in common with a diode-pumped solid-state laser than a semiconductor laser, is designed as a replacement for the blue argon ion laser.

The inherent spectral width of semiconductor lasers is quite large compared with, eg, a gas laser with a typical multimode diode having a spectral width of 2–4 nm in the form of many closely spaced longitudinal modes. To reduce the spectral width, wavelength-selective optics may be employed in a manner similar to that used for other lasers: namely, the use of a grating in place of the HR. Given the tiny dimensions of a semiconductor laser, though, implementation of a discrete diffraction grating is difficult, so the grating may be constructed as an elongated structure called a *distributed Bragg reflector* (DBR). Such a reflector resembles a corrugated surface (like the center of cardboard) manufactured from dielectric materials. Reflection of waves at the interface between two materials of different indexes of refraction leads to constructive interference at a single well-defined wavelength (determined by the distance between peaks in the corrugation). Such a reflector acts much like a high performance dielectric mirror, with the specific wavelength of maximum reflectivity called the *Bragg wavelength*.

The wavelength stability of a laser diode depends highly on the temperature of the diode. As the temperature of a laser diode increases, the refractive index of the semiconductor material itself changes. Since the resonant wavelength of a cavity depends on the refractive index of the material, the wavelength shifts toward longer wavelengths as the temperature of the device increases.

Solid-state lasers are by far the most commonly used today. Low powered ir lasers are employed for many optical storage applications, including CD and DVD players. Visible laser diodes are also common and are used for applications ranging from laser pointers and levels to scanning applications, often replacing helium–neon lasers. For storage applications, a shorter wavelength is preferred since it can be used to yield a smaller spot size, and hence increase the density of optical storage media. While short-wavelength diode lasers have recently been developed, their use in optical storage devices is relatively new. In addition to short-wavelength diodes, another alternative is the DPSS laser, in which a diode laser is used to pump a solid-state laser usually employing a YAG or vanadate crystal. High powered arrays of semiconductor lasers allow the construction of DPSS lasers yielding second-harmonic powers in the tens-of-watts range, usually in the green region of the spectrum.

9.3. Dye Lasers. Aside from previous examples of vibrational, rotational, electronic, and hybrids of these types of levels seen in the active species of several lasers in this article, an energy level in a species might well involve all three types of levels. This leads to energy levels that are quite broad (called energy *bands* at this point), so the resulting laser line is also quite broad spectrally. Dye lasers usually involve all three types of energy-level mechanisms, which leads to a continuum output over a given range. By tuning the feedback

mechanism (eg, with a diffraction grating and etalon), one can selectively amplify a given wavelength: these lasers are truly tunable.

In a dye laser, the active lasing medium is an organic dye dissolved in a solvent such as alcohol. These lasers may be pumped by either flashlamps (10) (like a solid-state laser) or by another laser. Laser-pumped dye lasers normally employ nitrogen or excimer pump lasers and hence are pulsed, but continuous dye lasers are possible using a CW argon ion laser as a pump source. The major advantage of this laser over other types is continuous tunability over a wide range. A laser employing rhodamine-6G, eg, can be tuned continuously through a range of wavelengths spanning visually from a shade of green–yellow yellow to a shade of red. By changing the dye employed, the range can be selected. Hundreds of dyes are known to lase. The dyes involved are large molecules with molecular weights in the range 400–500 amu. Most laser dyes belong to one of several families of dyes such as rhodamines or coumarins. Lasing begins when incident energy is absorbed by the dye, exciting it from the lowest singlet state to a high energy level within the upper singlet band, as shown in Figure 30. From the high energy level the dye falls to a slightly lower state within the same singlet band, which serves as an upper lasing level. A laser transition can then occur between the upper lasing level and the lower singlet state, which serves as a lower lasing level.

An alternative pathway exists to foil laser action in the triplet states of the dye, also shown on Figure 30. Triplet states originate when excited electrons in the dye molecule spin in the same direction as that of the remaining electrons in the dye molecule (the singlet states result when the excited electron spins in the direction opposite to the lower energy-state valence electrons still in the dye molecule). Because triplet states have lower energies than corresponding singlet states, dye molecules can easily migrate to those states and in doing so depopulate the upper lasing level. To make matters worse, triplet states are metastable and have much longer lifetimes than the singlet levels. When a short pump pulse such as that from a nitrogen laser (at 10 ns) is employed, triplet states do not form and do not present a problem for lasing, but when a flashlamp is used, which generally have pulse widths of over 1 μs, triplet states can form. For this reason, flashlamps must be designed to discharge as quickly as possible. Ordinary photographic strobes, eg, often have pulse lengths of 1 ms and will not work for pumping most dye lasers. In addition to a fast pump pulse to prevent triplet formation, triplet quenching additives can be mixed with many dyes, such as cyclooctatetraene (COT). These additives work by providing a deexcitation pathway from the triplet states, allowing the dye molecule involved to reenter the lasing process.

While flashlamp pumping of a dye is possible, in a configuration resembling that of a solid-state laser except with a quartz dye cell in place of a rod, most dyes are pumped by another laser, such as an excimer or nitrogen laser as in Figure 31. These lasers are fairly simplistic and, compared to flashlamp- pumped types, exhibit unusually high gain, allowing the use of a short dye cell.

In this arrangement, the pump laser beam is focused to a line on a dye cell using a cylindrical lens, in which the focused beam from a nitrogen or excimer pump laser strikes the side of a dye cell exciting dye molecules along the length of the cell (11). In most cases, it is possible to operate the dye laser superradiantly

by focusing the pump laser tightly onto the dye cell, but this is not desirable since a dye laser operated in such a way is not tunable, losing the major advantage of this laser over other lasers, so pump light is not normally focused so tightly that superradiance occurs (and while not desirable for tunability, the ability to operate the laser superradiantly makes alignment of this type of laser particularly simple). Pump light striking the dye cell excites dye molecules in a narrow channel physically just inside the dye cell itself: Penetration of the pump light into the cell is minimal, and essentially all is absorbed within the first few millimeters of dye within the cell.

Optics for a laser-pumped dye laser usually consist of an output coupler and a diffraction grating. Although a simple mirror (HR) could be used in the cavity, a grating is invariably used to allow tuning of the laser across the wide gain bandwidth of the laser, the biggest advantage of a dye laser. In a practical dye laser, a beam expander consisting of two lenses is placed between the dye cell and the diffraction grating to utilize more of the grating area (12). An etalon may also be placed in the optical path to reduce spectral width (13). Finally, a MOPA configuration may be used in which a separate amplifier is placed in front of the oscillator (complete with wavelength selector). The configuration for a complete dye laser with narrow spectral output width is shown in Figure 32.

The pump laser is usually a nitrogen or excimer laser emitting in the uv range, but frequency-doubled YAG lasers (at 532 nm) may also be used. Because the pump laser is pulsed, the dye laser also has a pulsed output, but it is possible to build a CW dye laser pumped by a CW laser source such as an ion laser. In this case, the biggest problem becomes heat management and degradation of the dye itself. Both problems are alleviated by forming the dye into a continually flowing sheet of liquid called a *laminar flow*. Flowing dye is pumped through a nozzle to create a broad, flat stream onto which pump laser light, usually from an argon-ion laser given that many laser dyes absorb strongly in the violet-to-green band where the argon-ion laser has the strongest output, is focused by a lens (or a concave mirror) onto the surface of the flowing dye. In a manner similar to that of the laser-pumped arrangement of Figure 31, a column of excited dye molecules serves as the gain medium, which, is off-axis (usually at the Brewster angle) from the axis of the optical resonator.

The output characteristics of a dye laser are highly dependent on the optics employed. A laser employing broadband optics would feature a naturally broad spectral width, typically spanning almost the entire range of the dye (especially in a laser-pumped dye laser, in which gain is usually large); this can be as large as 100 nm. A laser with such characteristics is not particularly useful, so wavelength selection is invariably employed.

Use of a diffraction grating alone as a wavelength selector (with suitable beam- expanding optics allowing utilization of a large area of the grating surface) renders a spectral width of 0.01 nm, which is suitable for many applications, but to reduce linewidth an intracavity etalon is often included in the optical path. Use of an etalon along with a diffraction grating can render spectral widths as low as 0.0005 nm. Such linewidths are required when a dye laser is used as a tunable laser source for spectroscopy. Etalons may be angle or pressure tuned, with pressure tuning preferred for simplicity. As a modelocked laser, dye lasers have produced the shortest pulses produced from any laser source. The wide gain

bandwidth of a typical dye allows the production of a series of extremely short pulses, in the femtosecond range.

As a source for spectroscopy, the dye laser is ideal given the wide range over which tuning may be accomplished and the narrow spectral width of the output. It is used in situations such as atomic absorption spectroscopy, where the beam is passed through a sample in a cell or a hot gas such as exhaust gas from a flame or the gases burning inside the cylinder of an internal combustion engine. Compact flashlamp-pumped dye lasers are occasionally employed in the field of opthalmology for retinal photocoagulation. Tunability allows the output to be optimized for the peak absorption wavelength of hemoglobin, and the pulses are short enough to coagulate blood without generating an explosive shock wave within tissue as Q-switched laser pulses can.

BIBLIOGRAPHY

"Lasers" in *ECT* 3rd ed., Vol. 14, pp. 42–81, by J. F. Ready, Honeywell Technology Center; in *ECT* 4th ed., Vol. 15, pp. 1–50, by John F. Ready, Honeywell Technology Center; "Lasers" in *ECT* (online), posting date: December 4, 2000, by John F. Ready, Honeywell Technology Center.

CITED PUBLICATIONS

1. E. S. Fry and co-workers, *Phys. Rev. Lett.* **70** (21), 3235 (1993).
2. D. A. Leonard and co-workers, *App. Phys. Lett.* **7** (6), 175 (1965).
3. M. S. Csele, *Fundamentals of Light Sources and Lasers*, John Wiley & Sons, Inc., Hoboken, New Jersey, 2004.
4. J. A. Giordmaine, *Sci. Am.* **38** (April 1964).
5. W. B. Bridges, *App. Phys. Lett.* **4** (7) (1964).
6. W. Bridges, *IEEE J. Selected Top. Q. Elect.* **6** (6), 885 (Nov./Dec. 2000).
7. D. A. Leonard, *Appl. Phys. Lett.* **7** (1), 4 (1965).
8. J. G. Eden, *IEEE J. Selected Top. in Q. Electronics* **6** (6), 1051 (Nov./Dec. 2000).
9. J. J. Ewing, *IEEE J. Selected Top. in Q. Electronics* **6** (6), 1061 (Nov./Dec. 2000).
10. P. P. Sorokin and co-workers, *J. Chem. Phys.* **48** (10) (1968).
11. J. A. Meyer and co-workers, *Appl. Phys. Lett.* **16** (1), 3 (1970).
12. G. Capelle and co-workers, *Appl. Opt.* **9** (12), 2742 (1970).
13. T. Hansch, *Appl. Opt.* **11** (4), 895 (1972).

GENERAL REFERENCES

M. S. Csele, *Fundamentals of Light Sources and Lasers*, John Wiley & Sons, Inc., Hoboken, New Jersey, 2004.
S. O. Kasap, *Optoelectronics and Photonics*, Prentice-Hall, New Jersey, 2001.
W. T. Silfvast, *Laser Fundamentals*, Cambridge University Press, New York, 1996.
K. Kuhn, *Laser Engineering*, Prentice-Hall, New Jersey, 1998.
B. A. Saleh and M.C. Teich, *Fundamentals of Photonics*, John Wiley & Sons, Inc., New York, 1991.

A. E. Siegman, *Lasers*, University Science Books, Mill Valley, Calif., 1986.
A. Yariv, *Quantum Electronics*, John Wiley & Sons, Inc., New York, 1989.
J. H. Eberly and P. W. Milonni, *Lasers*, John Wiley & Sons, Inc., New York, 1988.
T. Pearsall, *Photonics Essentials*, McGraw-Hill, New York, 2003.
J. Hecht, *The Laser Guidebook*, McGraw-Hill, New York, 1992.
F. G. Smith and T. A. King, *Optics and Photonics*, John Wiley & Sons, Inc., New Jersey, 2000.

MARK CSELE
Niagara College

Table 1. **Commercially Available HeNe Lasers**

Wavelength, nm	Relative gain (compared to 632.8-nm output)
543.5 (green)	0.06
594.1 (yellow)	0.07
611.9 (orange)	0.2
632.8 (red)	1

Table 2. **Comparison of Argon and Krypton Laser Output**

Argon ion (Ar^+)		Krypton ion (Kr^+)	
Wavelength, nm	Line power, mW	Wavelength, nm	Line power, mW
454.5	140 mW	406.7	
457.9	420 mW	413.1	150 mW (combined)
465.8	180 mW	415.4	
472.7	240 mW	476.2	50 mW
476.5	720 mW	482.5	30 mW
488.0	1800[a]	520.8	70 mW
496.5	720 mW	530.9	200 mW
501.7	480 mW	568.2	150 mW
514.5	2400[a]	647.1	500 mW
528.7	420 mW	676.4	120 mW

[a]Line power mW.

Table 3. **Excimer Characteristics**

Laser species	Wavelength, nm	Relative power output
ArF	193	0.5
KrF	249	1.0
XeCl	308	0.7
XeF	350	0.6

Table 4. **Common Nd^{3+} Hosts and Wavelengths**

Common name	Chemical formula and name	Wavelength, nm
YAG	$Y_3Al_5O_{12}$ (yttrium aluminum garnet)	1064
vanadate	YVO_4 (yttrium *o*-vanadate)	1064
glass	Various phosphate and silicate glasses	1060/1054
YLF	YLF (yttrium lithium fluoride)	1053
LSB	$LaSc_3(BO_3)_4$	1062

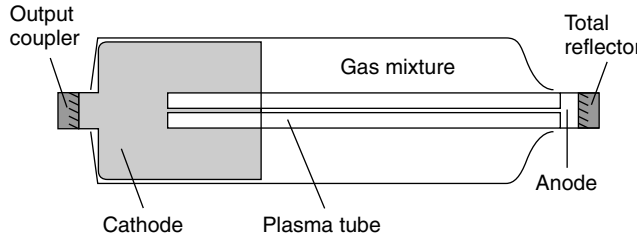

Fig. 1. Elements of a helium–neon laser tube.

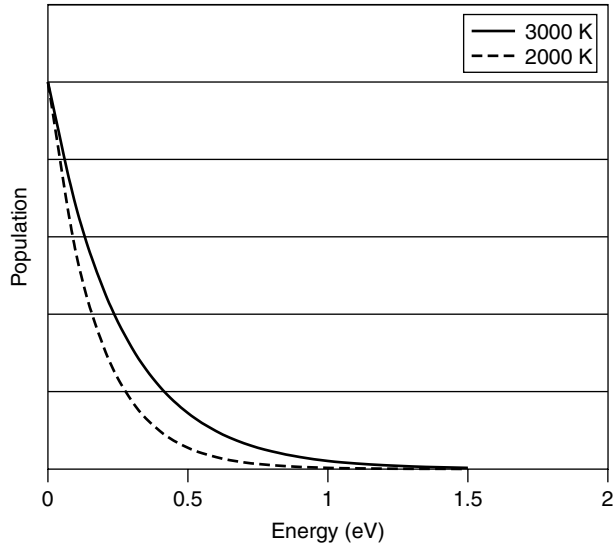

Fig. 2. Boltzmann distribution of atomic energies at various temperatures.

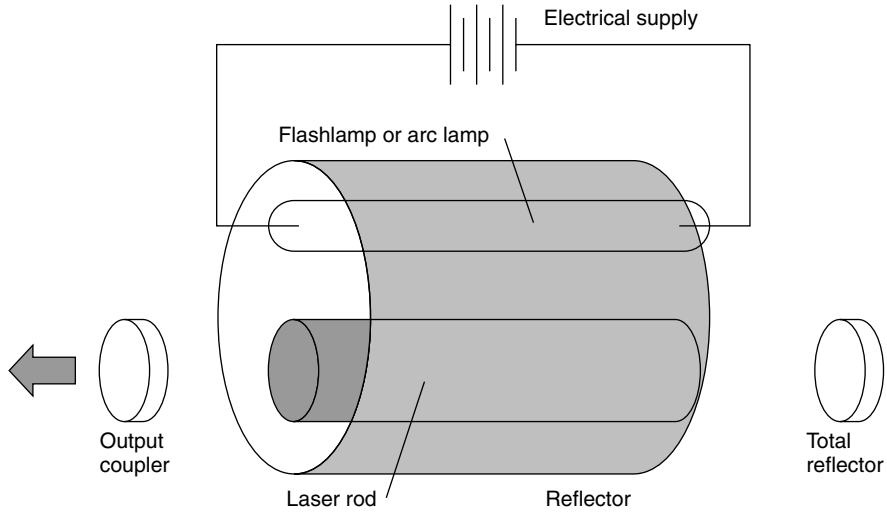

Fig. 3. Optical pumping for lasers.

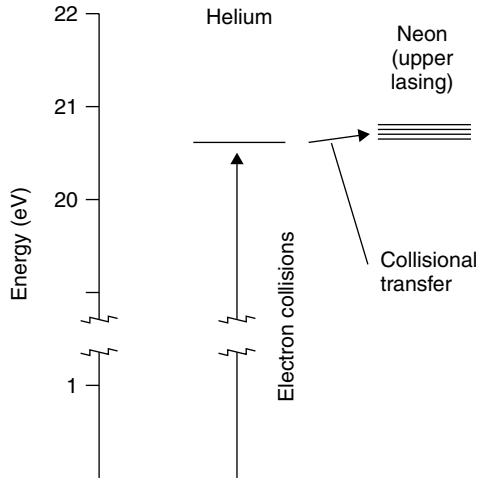

Fig. 4. Pumping mechanism in the helium–neon laser.

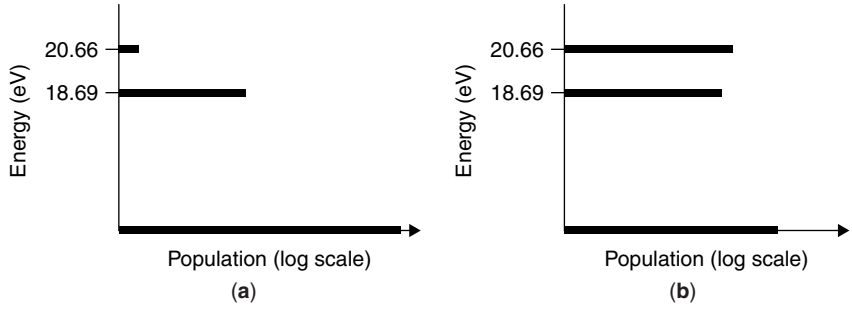

Fig. 5. Population energies of neon atoms (**a**) at equilibrium and (**b**) in a HeNe laser.

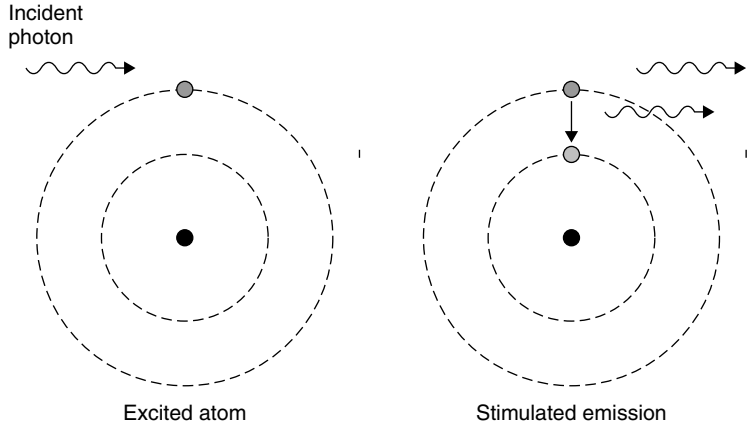

Fig. 6. Stimulated emission from an excited atom.

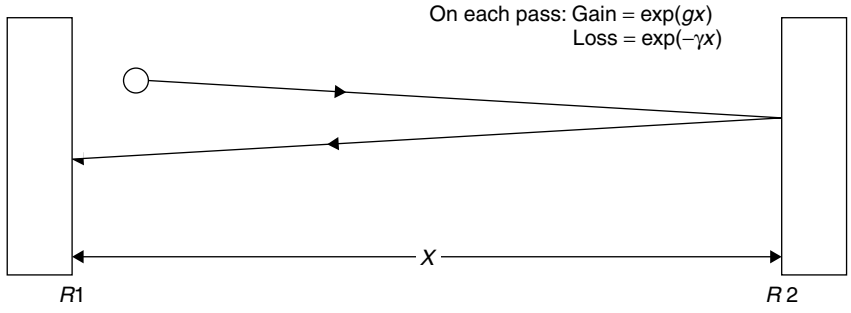

Fig. 7. Gain and loss in a laser.

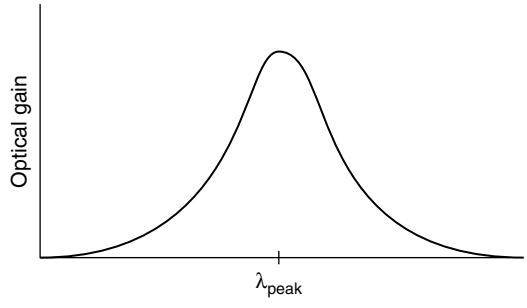

Fig. 8. Gain curve for a practical laser.

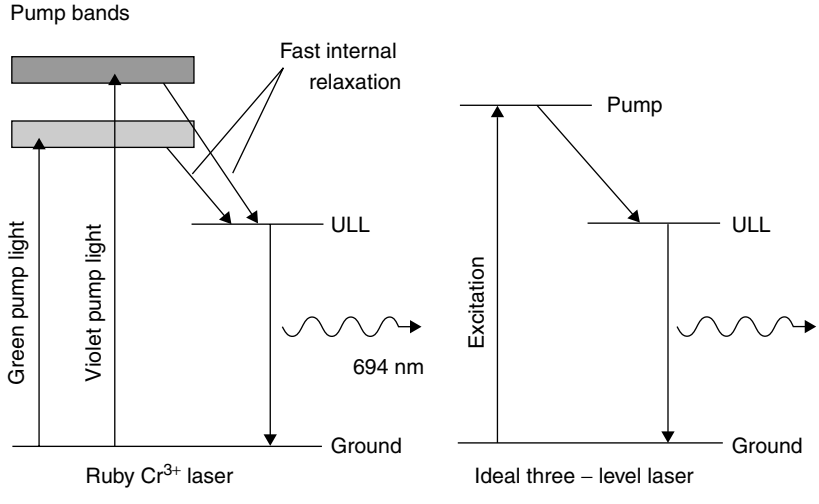

Fig. 9. Dynamics of the ruby laser.

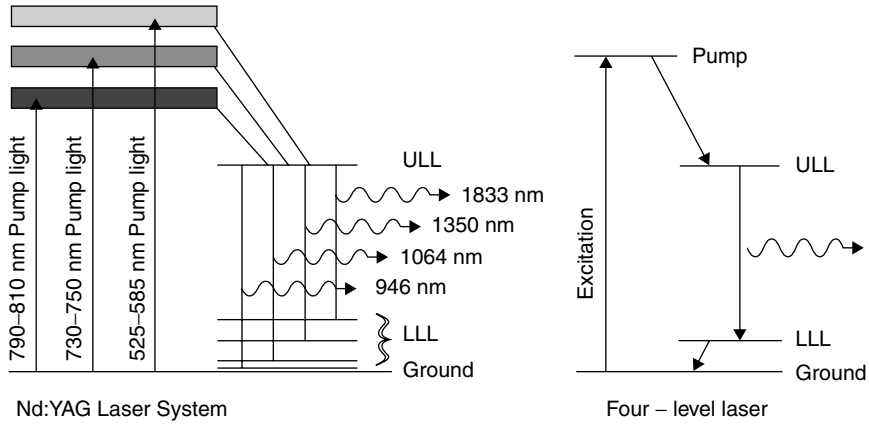

Fig. 10. Energy levels in the YAG laser.

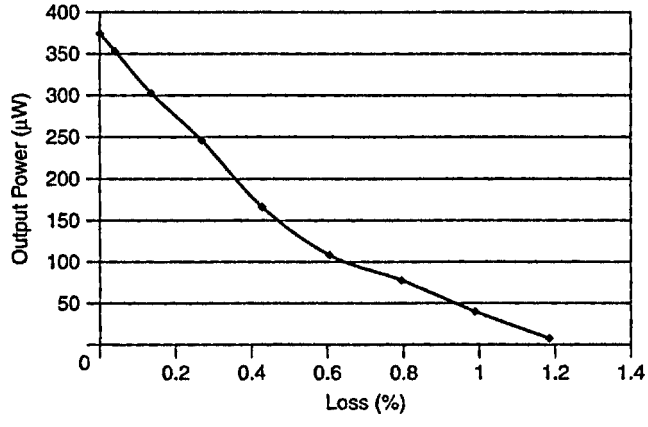

Fig. 11. Effect of cavity loss on output power in a laser.

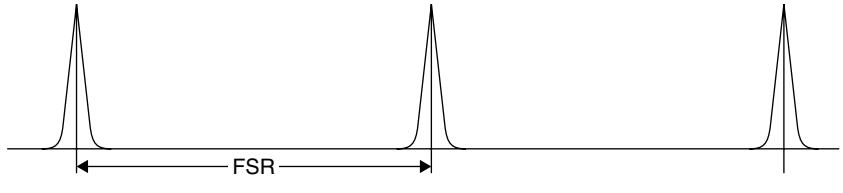

Fig. 12. Response of an interferometer.

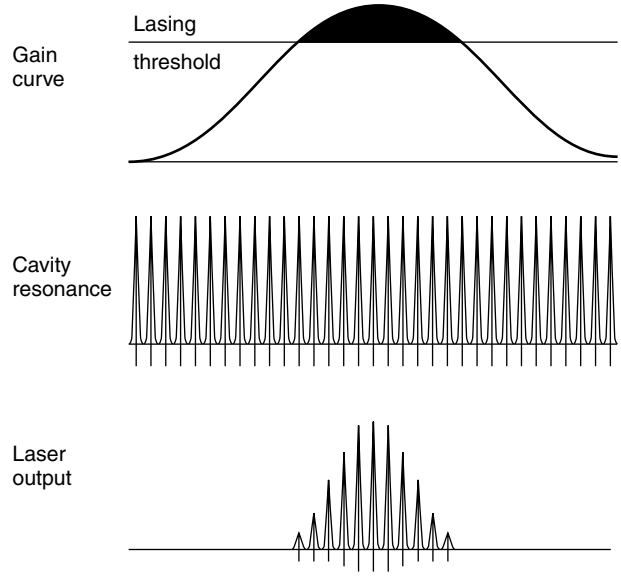

Fig. 13. Origin of longitudinal modes.

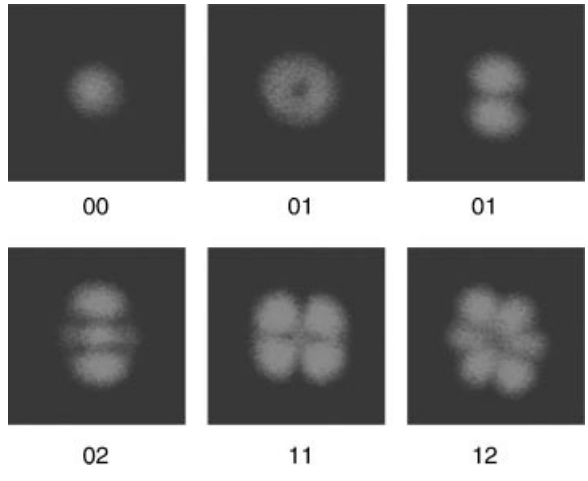

Fig. 14. Transverse electromagnetic modes (TEM).

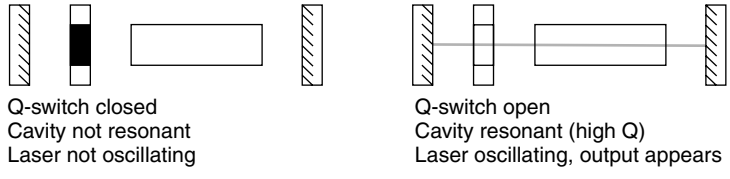

Fig. 15. Q-switching a laser.

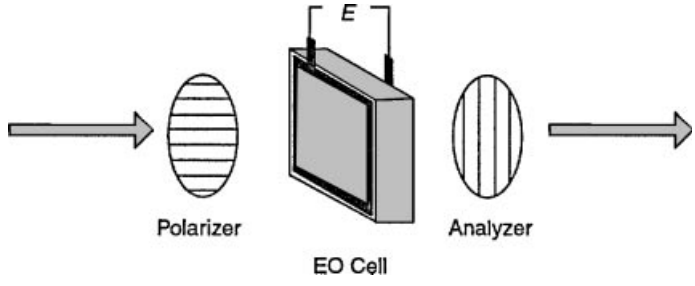

Fig. 16. Electrooptic modulator.

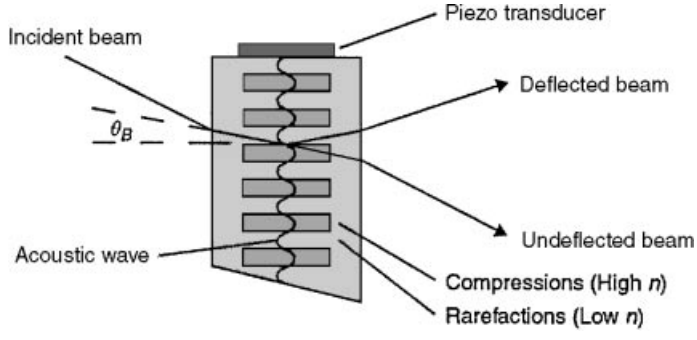

Fig. 17. Acoustic waves in an AO modulator (Bragg diffraction condition).

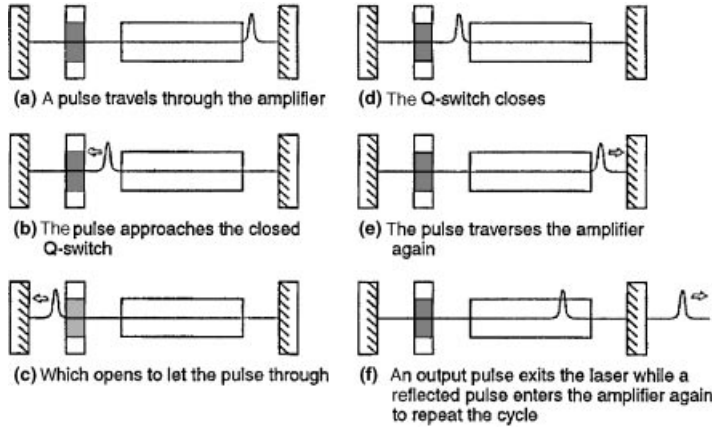

Fig. 18. Modelock pulse development in the time domain.

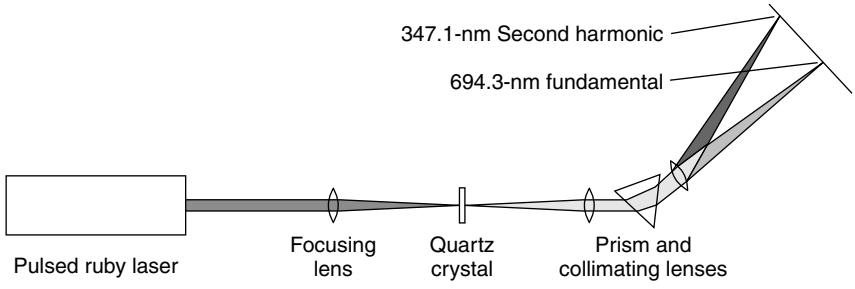

Fig. 19. Generation of second harmonic light.

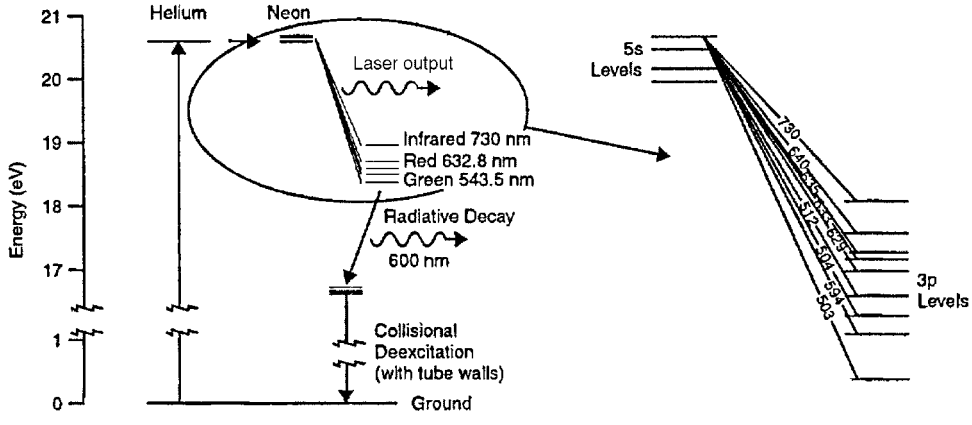

Fig. 20. Helium–neon energy levels.

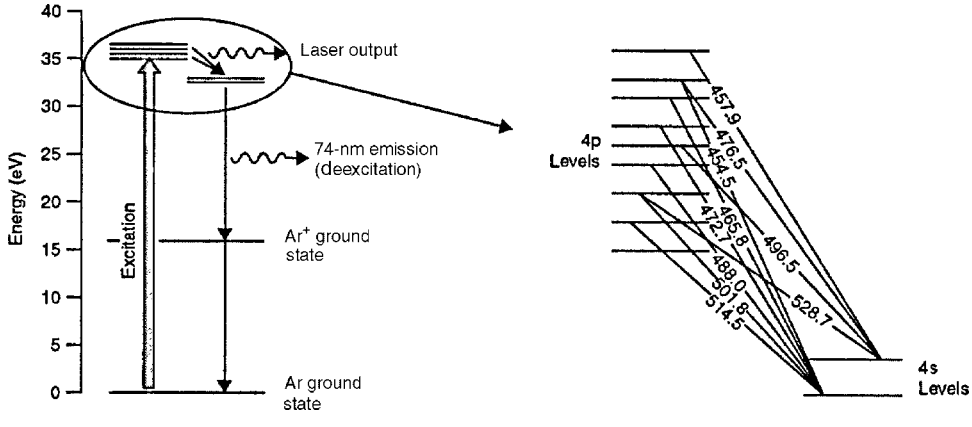

Fig. 21. Argon ion laser energy levels and transitions.

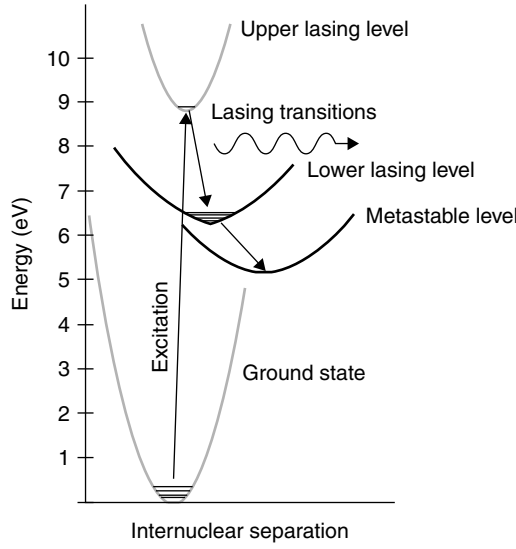

Fig. 22. Molecular nitrogen laser energy levels.

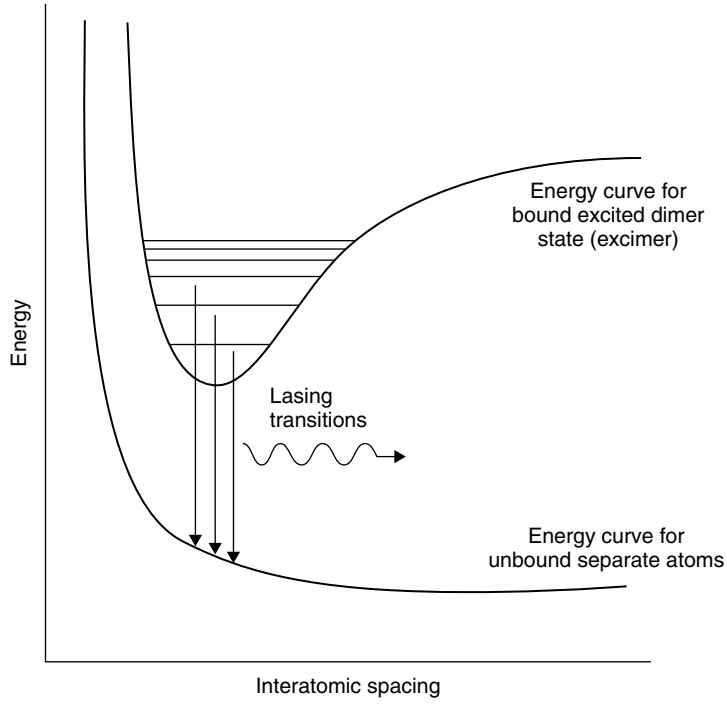

Fig. 23. Excimer energy levels.

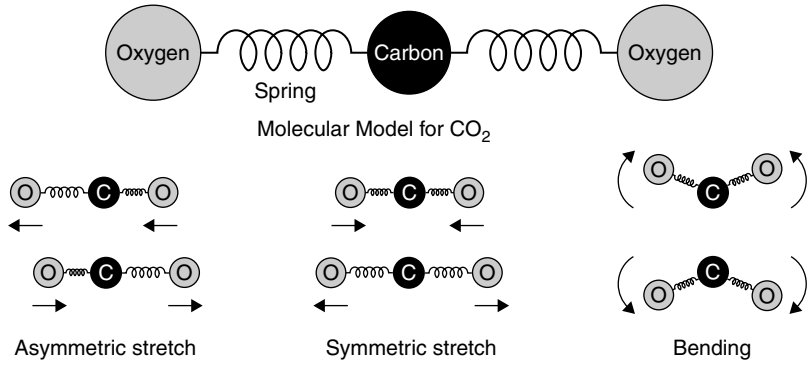

Fig. 24. Model of the CO_2 molecule and vibrational modes.

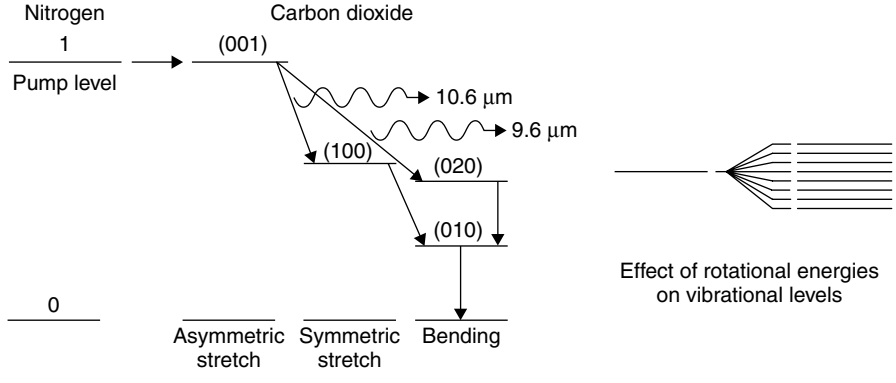

Fig. 25. Energy levels in the CO_2 laser.

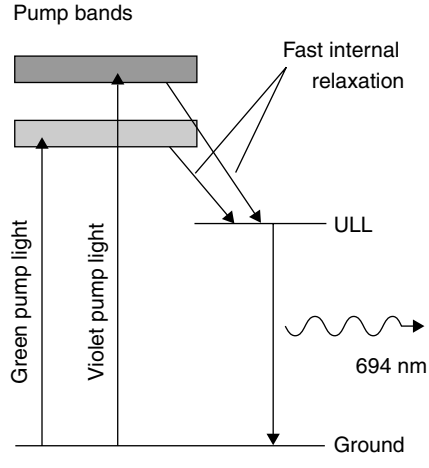

Fig. 26. Energy levels in ruby.

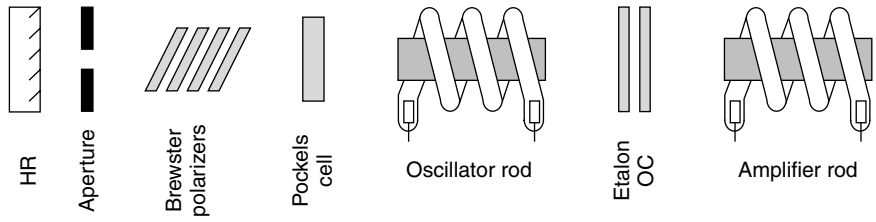

Fig. 27. Optical train of a ruby laser.

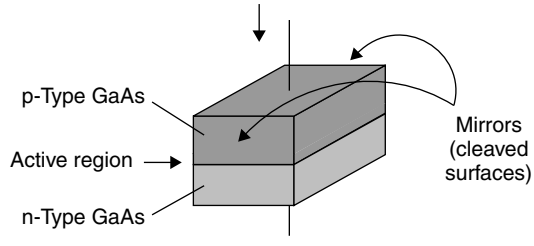

Fig. 28. Homojunction laser structure.

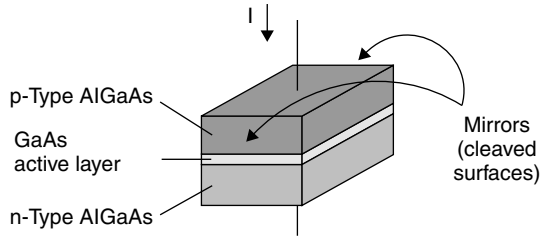

Fig. 29. Heterojunction laser structure.

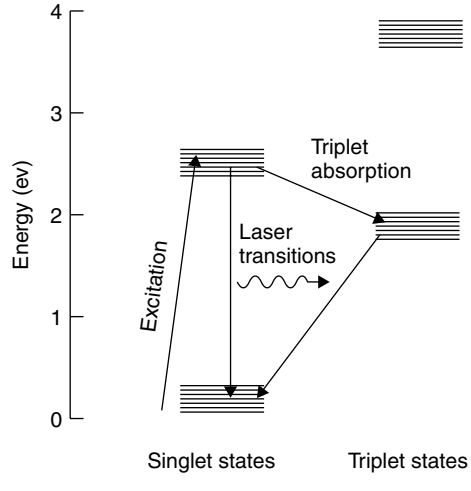

Fig. 30. Laser dye energy levels.

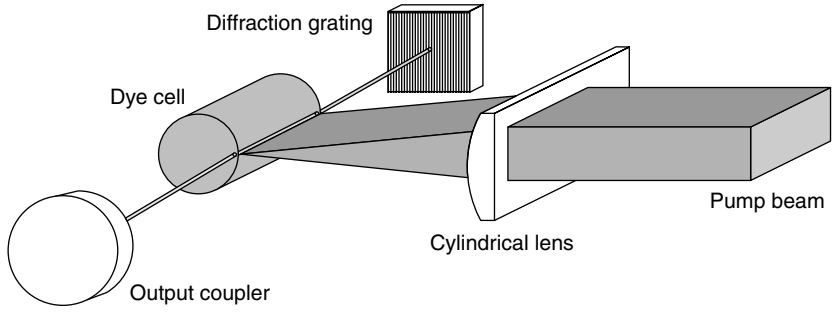

Fig. 31. Laser-pumped dye laser.

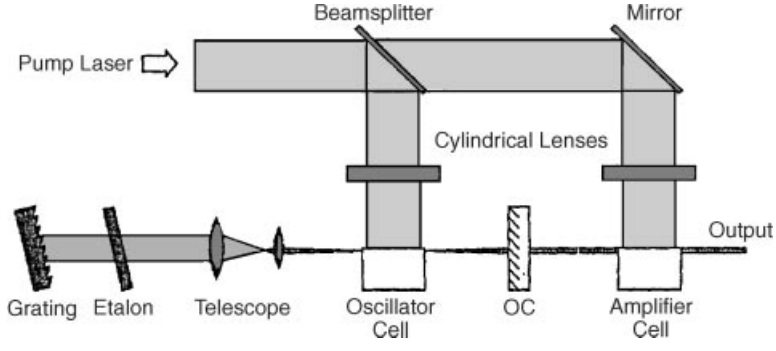

Fig. 32. Optics for a dye laser.