Introduction to Laser Technology

10

Basic Laser Principles	**10.2**
Propagation Characteristics of Laser Beams	**10.6**
Transverse Modes and Mode Control	**10.10**
Single Axial Longitudinal Mode Operation	**10.11**
Frequency and Amplitude Fluctuations	**10.13**
Tunable Operation	**10.15**
Types of Lasers	**10.16**
Laser Applications	**10.28**

Basic Laser Principles

Lasers are devices that produce intense beams of light which are *monochromatic*, *coherent*, and *highly collimated*. The wavelength (color) of laser light is extremely pure (monochromatic) when compared to other sources of light, and all of the photons (energy) that make up the laser beam have a fixed phase relationship (coherence) with respect to one another. Light from a laser typically has very low divergence. It can travel over great distances or can be focused to a very small spot with a brightness which exceeds that of the sun. Because of these properties, lasers are used in a wide variety of applications in all walks of life.

The basic operating principles of the laser were put forth by Charles Townes and Arthur Schalow from the Bell Telephone Laboratories in 1958, and the first actual laser, based on a pink ruby crystal, was demonstrated in 1960 by Theodor Maiman at Hughes Research Laboratories. Since that time, literally thousands of lasers have been invented (including the edible "Jello" laser), but only a much smaller number have found practical applications in scientific, industrial, commercial, and military applications. The helium neon laser (the first continuous-wave laser), the semiconductor diode laser, and air-cooled ion lasers have found broad OEM application. In recent years the use of diode-pumped solid-state (DPSS) lasers in OEM applications has been growing rapidly.

`The term "laser" is an acronym for (L)ight (A)mplification by (S)timulated (E)mission of (R)adiation. To understand the laser, one needs to understand the meaning of these terms. The term "light" is generally accepted to be electromagnetic radiation ranging from 1 nm to 1000 μm in wavelength. The visible spectrum (what we see) ranges from approximately 400 to 700 nm. The wavelength range from 700 nm to 10 μm is considered the near infrared (NIR), and anything beyond that is the far infrared (FIR). Conversely, 200 to 400 nm is called ultraviolet (UV); below 200 nm is the deep ultraviolet (DUV).

To understand stimulated emission, we start with the Bohr atom.

THE BOHR ATOM

In 1915, Neils Bohr proposed a model of the atom that explained a wide variety of phenomena that were puzzling scientists in the late 19th century. This simple model became the basis for the field of quantum mechanics and, although not fully accurate by today's understanding, still is useful for demonstrating laser principles.

In Bohr's model, shown in figure 10.1, electrons orbit the nucleus of an atom. Unlike earlier "planetary" models, the Bohr atom has a limited number of fixed orbits that are available to the electrons. Under the right circumstances an electron can go from its ground state (lowest-energy orbit) to a higher (excited) state, or it can decay from a higher state to a lower state, but it cannot remain between these states. The allowed energy states are called "quantum" states and are referred to by the principal "quantum numbers" 1, 2, 3, etc. The quantum states are represented by an energy-level diagram.

For an electron to jump to a higher quantum state, the atom must receive energy from the outside world. This can happen through a variety of mechanisms such as inelastic or semielastic collisions with other atoms and absorption of energy in the form of electromagnetic radiation (e.g., light). Likewise, when an electron drops from a higher state to a lower state, the atom must give off energy, either as kinetic activity (nonradiative transitions) or as electromagnetic radiation (radiative transitions). For the remainder of this discussion we will consider only radiative transitions.

PHOTONS AND ENERGY

In the 1600s and 1700s, early in the modern study of light, there was a great controversy about light's nature. Some thought that light was made up of particles, while others thought that it was made up of waves. Both concepts explained some of the behavior of light, but not all. It was finally determined that light is made up of particles called "photons" which exhibit both particle-like and wave-like properties. Each photon has an intrinsic energy determined by the equation

$$E = h\nu \tag{10.1}$$

where ν is the frequency of the light and h is Planck's constant. Since, for a wave, the frequency and wavelength are related by the equation

$$\lambda\nu = c \tag{10.2}$$

where λ is the wavelength of the light and c is the speed of light in a vacuum, equation 10.1 can be rewritten as

$$E = \frac{hc}{\lambda}. \tag{10.3}$$

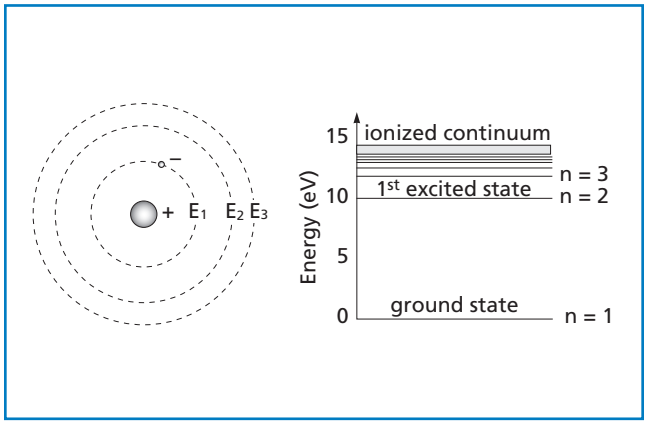

Figure 10.1 **The Bohr atom and a simple energy-level diagram**

It is evident from this equation that the longer the wavelength of the light, the lower the energy of the photon; consequently, ultraviolet light is much more "energetic" than infrared light.

Returning to the Bohr atom: for an atom to absorb light (i.e., for the light energy to cause an electron to move from a lower energy state E_n to a higher energy state E_m), the energy of a single photon must equal, almost exactly, the energy difference between the two states. Too much energy or too little energy and the photon will not be absorbed. Consequently, the wavelength of that photon must be

$$\lambda = \frac{hc}{\Delta E} \qquad (10.4)$$

where

$$\Delta E = E_m - E_n.$$

Likewise, when an electron decays to a lower energy level in a radiative transition, the photon of light given off by the atom must also have an energy equal to the energy difference between the two states.

SPONTANEOUS AND STIMULATED EMISSION

In general, when an electron is in an excited energy state, it must eventually decay to a lower level, giving off a photon of radiation. This event is called "spontaneous emission," and the photon is emitted in a random direction and a random phase. The average time it takes for the electron to decay is called the time constant for spontaneous emission, and is represented by τ.

On the other hand, if an electron is in energy state E_2, and its decay path is to E_1, but, before it has a chance to spontaneously decay, a photon happens to pass by whose energy is approximately $E_2 - E_1$, there is a probability that the passing photon will cause the electron to decay in such a manner that a photon is emitted at exactly the same wavelength, in exactly the same direction, and with exactly the same phase as the passing photon. This process is called "stimulated emission." Absorption, spontaneous emission, and stimulated emission are illustrated in figure 10.2.

Now consider the group of atoms shown in figure 10.3: all begin in exactly the same excited state, and most are effectively within the stimulation range of a passing photon. We also will assume that τ is very long, and that the probability for stimulated emission is 100 percent. The incoming (stimulating) photon interacts with the first atom, causing stimulated emission of a coherent photon; these two photons then interact with the next two atoms in line, and the result is four coherent photons, on down the line. At the end of the process, we will have eleven coherent photons, all with identical phases and all traveling in the same direction. In other words, the initial photon has been "amplified" by a factor of eleven. Note that the energy to put these atoms in excited states is provided externally by some energy source which is usually referred to as the "pump" source.

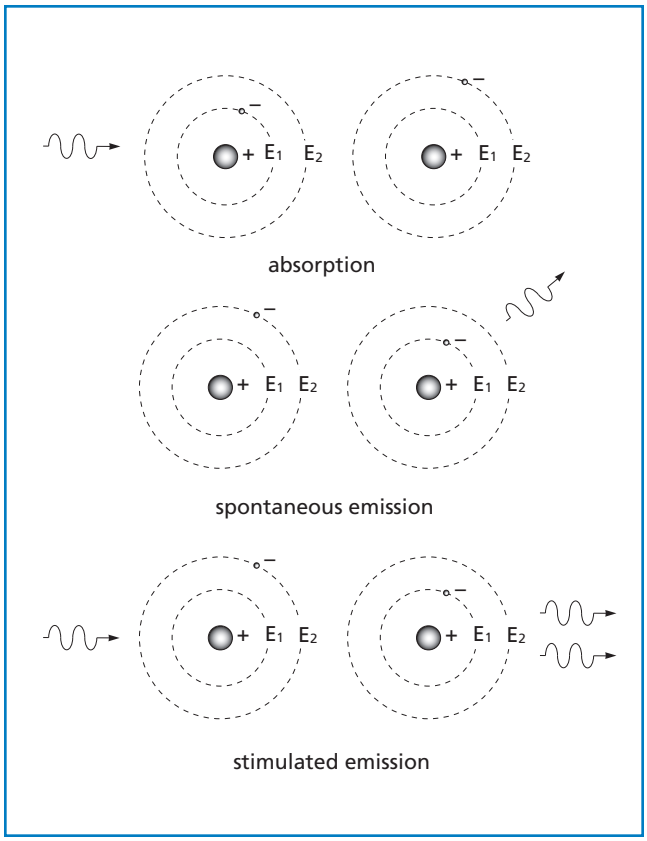

Figure 10.2 **Spontaneous and stimulated emission**

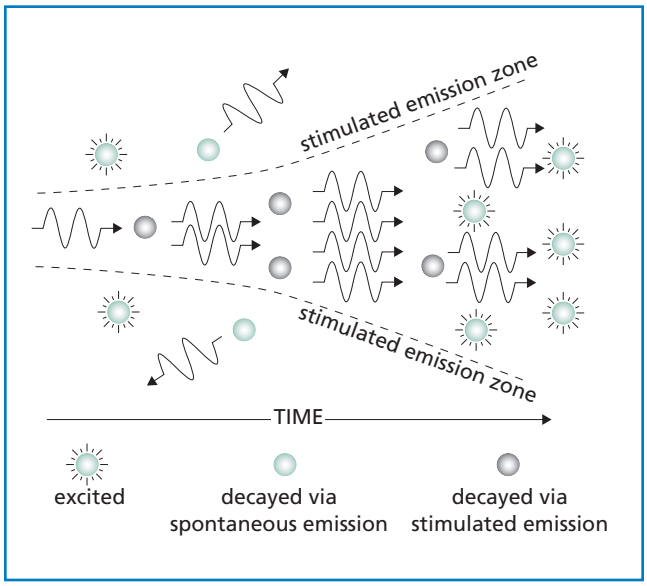

Figure 10.3 **Amplification by stimulated emission**

Of course, in any real population of atoms, the probability for stimulated emission is quite small. Furthermore, not all of the atoms are usually in an excited state; in fact, the opposite is true. Boltzmann's principle, a fundamental law of thermodynamics, states that, when a collection of atoms is at thermal equilibrium, the relative population of any two energy levels is given by

$$\frac{N_2}{N_1} = \exp\left(-\frac{E_2 - E_1}{kT}\right) \tag{10.5}$$

where N_2 and N_1 are the populations of the upper and lower energy states, respectively, T is the equilibrium temperature, and k is Boltzmann's constant. Substituting $h\nu$ for $E_2 - E_1$ yields

$$\Delta N \equiv N_1 - N_2 = \left(1 - e^{-h\nu/kT}\right) N_1. \tag{10.6}$$

For a normal population of atoms, there will always be more atoms in the lower energy levels than in the upper ones. Since the probability for an individual atom to absorb a photon is the same as the probability for an excited atom to emit a photon via stimulated emission, the collection of real atoms will be a net absorber, not a net emitter, and amplification will not be possible. Consequently, to make a laser, we have to create a "population inversion."

POPULATION INVERSION

Atomic energy states are much more complex than indicated by the description above. There are many more energy levels, and each one has its own time constants for decay. The four-level energy diagram shown in figure 10.4 is representative of some real lasers.

The electron is pumped (excited) into an upper level E_4 by some mechanism (for example, a collision with another atom or absorption of high-energy radiation). It then decays to E_3, then to E_2, and finally to the ground state E_1. Let us assume that the time it takes to decay from E_2 to E_1 is much longer than the time it takes to decay from E_2 to E_1. In a large population of such atoms, at equilibrium and with a continuous pumping process, a population inversion will occur between the E_3 and E_2 energy states, and a photon entering the population will be amplified coherently.

THE RESONATOR

Although with a population inversion we have the ability to amplify a signal via stimulated emission, the overall single-pass gain is quite small, and most of the excited atoms in the population emit spontaneously and do not contribute to the overall output. To turn this system into a laser, we need a positive feedback mechanism that will cause the majority of the atoms in the population to contribute to the coherent output. This is the resonator, a system of mirrors that reflects undesirable (off-axis) photons out of the system and reflects the desirable (on-axis) photons back into the excited population where they can continue to be amplified.

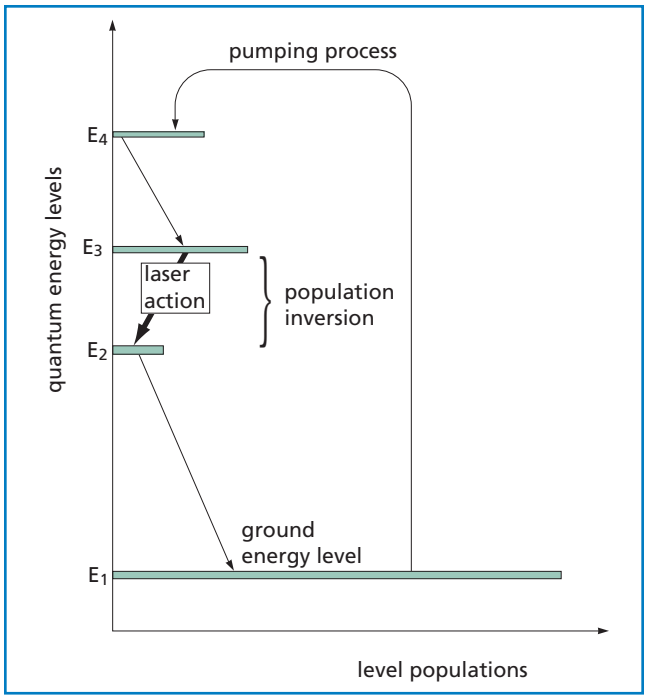

Figure 10.4 **A four-level laser pumping system**

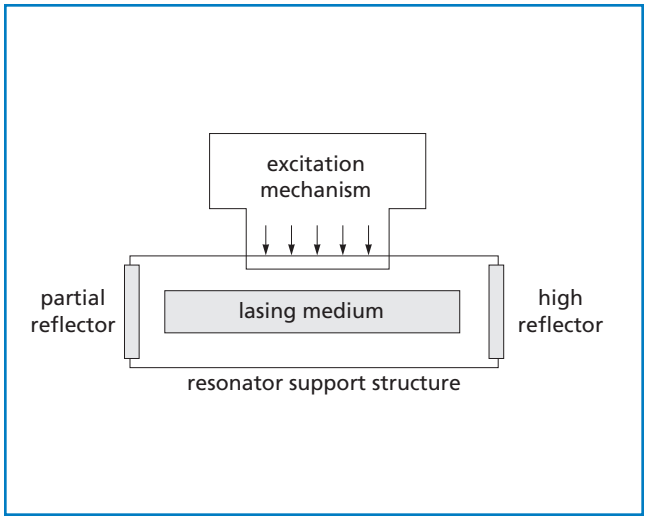

Figure 10.5 **Schematic diagram of a basic laser**

Now consider the laser system shown in figure 10.5. The lasing medium is pumped continuously to create a population inversion at the lasing wavelength. As the excited atoms start to decay, they emit photons spontaneously in all directions. Some of the photons travel along the axis of the lasing medium, but most of the photons are directed out the sides. The photons traveling along the axis have an opportunity to stimulate atoms

they encounter to emit photons, but the ones radiating out the sides do not. Furthermore, the photons traveling parallel to the axis will be reflected back into the lasing medium and given the opportunity to stimulate more excited atoms. As the on-axis photons are reflected back and forth interacting with more and more atoms, spontaneous emission decreases, stimulated emission along the axis predominates, and we have a laser.

Finally, to get the light out of the system, one of the mirrors has a partially transmitting coating that couples out a small percentage of the circulating photons. The amount of coupling depends on the characteristics of the laser system and varies from a fraction of a percent for helium neon lasers to 50 percent or more for high-power lasers.

Practical Optical Coatings

In the design of a real-world laser, the optical resonator is often the most critical component, and, particularly for low-gain lasers, the most critical components of the resonator are the mirrors themselves. The difference between a perfect mirror coating (the optimum transmission and reflection with no scatter or absorption losses) and a real-world coating, capable of being produced in volume, can mean a 50-percent (or greater) drop in output power from the theoretical maximum. Consider the 543-nm green helium neon laser line. It was first observed in the laboratory in 1970, but, owing to its extremely low gain, the mirror fabrication and coating technology of the day was incapable of producing a sufficiently loss-free mirror that was also durable. Not until the late 1990s had the mirror coating technology improved sufficiently that these lasers could be offered commercially in large volumes.

The critical factors for a mirror, other than transmission and reflection, are scatter, absorption, stress, surface figure, and damage resistance. Coatings with low damage thresholds can degrade over time and cause output power to drop significantly. Coatings with too much mechanical stress not only can cause significant power loss, but can also induce stress birefringence, which can result in altered polarization and phase relationships. The optical designer must take great care when selecting the materials for the coating layers and the substrate to ensure that the mechanical, optical, and environmental characteristics are suitable for the application.

The equipment used for both substrate polishing and optical coating is a critical factor in the end result. Coating scatter is a major contributor to power loss. Scatter arises primarily from imperfections and inclusions in the coating, but also from minute imperfections in the substrate. Over the last few years, the availability of "super-polished" mirror substrates has led to significant gains in laser performance. Likewise, ion-beam sputtering and next-generation ion-assisted ion deposition has increased the packing density of laser coatings, thereby reducing absorption, increasing damage thresholds, and enabling the use of new and exotic coating materials.

Propagation Characteristics of Laser Beams

BEAM WAIST AND DIVERGENCE

Diffraction causes light waves to spread transversely as they propagate, and it is therefore impossible to have a perfectly collimated beam. The spreading of a laser beam is in accord with the predictions of diffraction theory. Under ordinary circumstances, the beam spreading can be so small it can go unnoticed. The following formulas accurately describe beam spreading, making it easy to see the capabilities and limitations of laser beams. The notation is consistent with much of the laser literature, particularly with Siegman's excellent *Lasers* (University Science Books).

Even if a Gaussian TEM_{00} laser-beam wavefront were made perfectly flat at some plane, with all rays there moving in precisely parallel directions, it would acquire curvature and begin spreading in accordance with

$$R(z) = z\left[1+\left(\frac{\pi w_0^2}{\lambda z}\right)^2\right] \tag{10.7}$$

and

$$w(z) = w_0\left[1+\left(\frac{\lambda z}{\pi w_0^2}\right)^2\right]^{1/2} \tag{10.8}$$

where z is the distance propagated from the plane where the wavefront is flat, λ is the wavelength of light, w_0 is the radius of the $1/e^2$ irradiance contour at the plane where the wavefront is flat, $w(z)$ is the radius of the $1/e^2$ contour after the wave has propagated a distance z, and $R(z)$ is the wavefront radius of curvature after propagating a distance z. $R(z)$ is infinite at $z = 0$, passes through a minimum at some finite z, and rises again toward infinity as z is further increased, asymptotically approaching the value of z itself.

The plane $z = 0$ marks the location of a beam waist, or a place where the wavefront is flat, and w_0 is called the beam waist radius.

The irradiance distribution of the Gaussian TEM_{00} beam, namely,

$$I(r) = I_0 e^{-2r^2/w^2} = \frac{2P}{\pi w^2} e^{-2r^2/w^2}\ , \tag{10.9}$$

where $w = w(z)$ and P is the total power in the beam, is the same at all cross sections of the beam. The invariance of the form of the distribution is a special consequence of the presumed Gaussian distribution at $z = 0$. Simultaneously, as $R(z)$ asymptotically approaches z for large z, $w(z)$ asymptotically approaches the value

$$w(z) = \frac{\lambda z}{\pi w_0} \tag{10.10}$$

where z is presumed to be much larger than $\pi w_0^2/\lambda$ so that the $1/e^2$ irradiance contours asymptotically approach a cone of angular radius

$$\theta = \frac{w(z)}{z} = \frac{\lambda}{\pi w_0}. \tag{10.11}$$

This value is the far-field angular radius (half-angle divergence) of the Gaussian TEM_{00} beam. The vertex of the cone lies at the center of the waist (see figure 10.6).

It is important to note that, for a given value of λ, variations of beam diameter and divergence with distance z are functions of a single parameter, w_0, the beam waist radius.

NEAR-FIELD VS. FAR-FIELD DIVERGENCE

Unlike conventional light beams, Gaussian beams do not diverge linearly, as can be seen in figure 10.6. Near the laser, the divergence angle is extremely small; far from the laser, the divergence angle approaches the asymptotic limit described in equation 10.11 above. The Raleigh range (z_R), defined as the distance over which the beam radius spreads by a factor of $\sqrt{2}$, is given by

$$z_R = \frac{\pi w_0^2}{\lambda} \tag{10.12}$$

The Raleigh range is the dividing line between near-field divergence and mid-range divergence. Far-field divergence (the number quoted in laser specifications) must be measured at a point $>z_R$ (usually $10z_R$ will suffice). This is a very important distinction because calculations for spot size and other parameters in an optical train will be inaccurate if near- or mid-field divergence values are used. For a tightly focused beam, the distance from the waist (the focal point) to the far field can be a few millimeters or less. For beams coming directly from the laser, the far-field distance can be measured in meters.

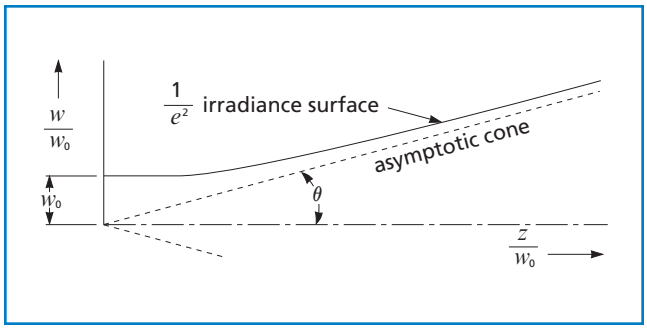

Figure 10.6 **Growth in beam diameter as a function of distance from the beam waist**

LOCATING THE BEAM WAIST

For a Gaussian laser beam, the location (and radius) of the beam waist is determined uniquely by the radius of curvature and optical spacing of the laser cavity mirrors because, at the reflecting surfaces of the cavity mirrors, the radius of curvature of the propagating beam is exactly the same as that of the mirrors. Consequently, for the flat/curved cavity shown in figure 10.7 (a), the beam waist is located at the surface of the flat mirror. For a

symmetric cavity (b), the beam waist is halfway between the mirrors; for non-symmetric cavities (c and d), the beam waist is located by using the equation

$$z_1 = \frac{L(R_2 - L)}{R_1 + R_2 - 2L} \qquad (10.13)$$

and

$$z_1 + z_2 = L$$

where L is the effective mirror spacing, R_1 and R_2 are the radii of curvature of the cavity mirrors, and z_1 and z_2 are the distances from the beam waist of mirrors 1 and 2, respectively. (Note that distances are measured from the beam waist, and that, by convention, mirror curvatures that are concave when viewed from the waist are considered positive, while those that are convex are considered negative.)

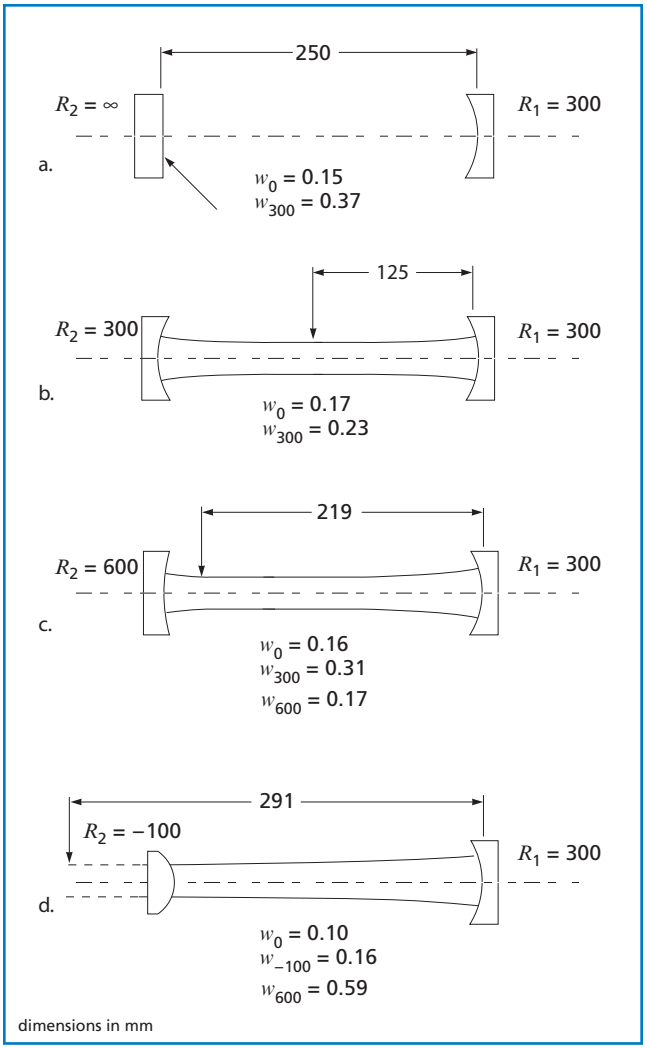

Figure 10.7 **Location of beam waist for common cavity geometries**

In any case but that of a flat output mirror, the beam waist is refracted as it passes through the mirror substrate. If the output coupler's second surface is flat, the effective waist of the refracted beam is moved toward the output coupler and is reduced in diameter. However, by applying a spherical correction to the second surface of the output coupler, the location of the beam waist can be moved to the output coupler itself, increasing the beam waist diameter and reducing far-field divergence. (See Calculating a Correcting Surface.)

It is useful, particularly when designing laser cavities, to understand the effect that mirror spacing has on the beam radius, both at the waist and at the curved mirror. Figure 10.8 plots equations 10.7 and 10.8 as a function of R/z (curved mirror radius divided by the mirror spacing). As the mirror spacing approaches the radius of curvature of the mirror ($R/z = 1$), the beam waist decreases dramatically, and the beam radius at the curved mirror becomes very large. On the other hand, as R/z becomes large, the beam radius at the waist and at the curved mirror are approximately the same.

CALCULATING A CORRECTING SURFACE

A laser beam is refracted as it passes through a curved output mirror. If the mirror has a flat second surface, the waist of the refracted beam moves closer to the mirror, and the divergence is increased. To counteract this, laser manufacturers often put a radius on the output coupler's second surface to collimate the beam by making a waist at the output coupler.

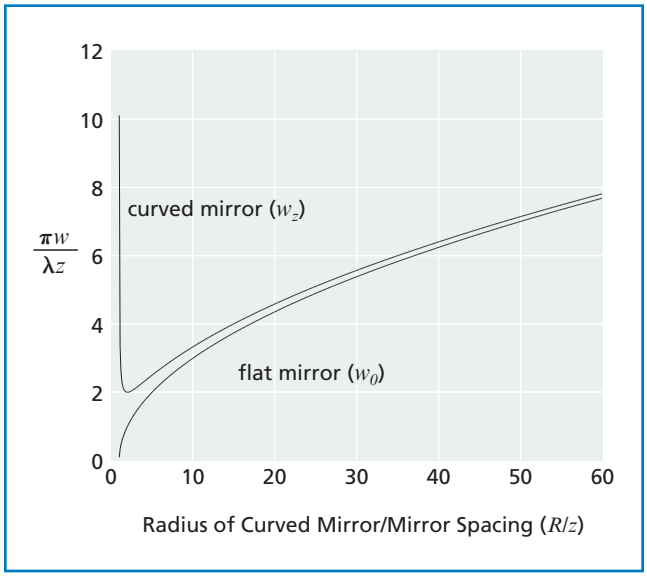

Figure 10.8 **Beam waist and output diameter as a function of mirror radius and separation**

This is illustrated by the case of a typical helium neon laser cavity consisting of a flat high reflector and an output mirror with a radius of curvature of 20 cm separated by 15 cm. If the laser is operating at 633 nm, the beam waist radius, beam radius at the output coupler, and beam half-angle divergence are

$$w_0 = 0.13 \text{ mm}, \; w_{200} = 0.26 \text{ mm}, \text{ and } \theta = 1.5 \text{ mrad},$$

respectively; however, with a flat second surface, the divergence nearly doubles to 2.8 mrad. Geometrical optics would give the focal length of the lens formed by the correcting output coupler as 15 cm; a rigorous calculation using Gaussian beam optics shows it should be 15.1 cm. Using the lens-makers formula

$$\frac{1}{f} = (n-1)\left(\frac{1}{R_1} - \frac{1}{R_2}\right) \qquad (10.14)$$

with the appropriate sign convention and assuming that $n = 1.5$, we get a convex correcting curvature of approximately 5.5 cm. At this point, the beam waist has been transferred to the output coupler, with a radius of 0.26 mm, and the far-field half-angle divergence is reduced to 0.76 mrad, a factor of nearly 4.

Correcting surfaces are used primarily on output couplers whose radius of curvature is a meter or less. For longer radius output couplers, the refraction effects are less dramatic, and a correcting second surface radius is unnecessary.

HIGHER ORDER GAUSSIAN LASER BEAMS

In the real world, the truly 100-percent, single transverse mode, Gaussian laser beam (also called a pure or fundamental mode beam) described by equations 10.7 and 10.8 is very hard to find. Low-power beams from helium neon lasers can be a close approximation, but the higher the power of the laser, and the more complex the excitation mechanism (e.g., transverse discharges, flash-lamp pumping), or the higher the order of the mode, the more the beam deviates from the ideal.

To address the issue of higher order Gaussian beams and mixed mode beams, a beam quality factor, M^2, has come into general use. A mixed mode is one where several modes are oscillating in the resonator at the same time. A common example is the mixture of the lowest order single transverse mode with the doughnut mode, before the intracavity mode limiting aperture is critically set to select just the fundamental mode. Because all beams have some wavefront defects, which implies they contain at least a small admixture of some higher order modes, a mixed mode beam is also called a "real" laser beam.

For a theoretical single transverse mode Gaussian beam, the value of the waist radius–divergence product is (from equation 10.11):

$$w_0\theta = \lambda / \pi. \qquad (10.15)$$

It is important to note that this product is an invariant for transmission of a beam through any normal, high-quality optical system (one that does not add aberrations to the beam wavefront). That is, if a lens focuses the single mode beam to a smaller waist radius, the convergence angle coming into the focus (and the divergence angle emerging from it) will be larger than that of the unfocused beam in the same ratio that the focal spot diameter is smaller: the product is invariant.

For a real laser beam, we have

$$W_0\Theta = M^2\lambda / \pi \qquad (10.16)$$

where W_0 and Θ are the $1/e^2$ intensity waist radius and the far-field half-divergence angle of the real laser beam, respectively. Here we have introduced the convention that upper case symbols are used for the mixed mode and lower case symbols for the fundamental mode beam coming from the same resonator. The mixed-mode beam radius W is M times larger than the fundamental mode radius at all propagation distances. Thus the waist radius is that much larger, contributing the first factor of M in equation 10.16. The second factor of M comes from the half-angle divergence, which is also M times larger. The waist radius–divergence half-angle product for the mixed mode beam also is an invariant, but is M^2 larger. The fundamental mode beam has the smallest divergence allowed by diffraction for a beam of that waist radius. The factor M^2 is called the "times-diffraction-limit" number or (inverse) beam quality; a diffraction-limited beam has an M^2 of unity.

For a typical helium neon laser operating in TEM_{00} mode, $M^2 < 1.05$. Ion lasers typically have an M^2 factor ranging from 1.1 to 1.7. For high-energy multimode lasers, the M^2 factor can be as high as 30 or 40. The M^2 factor describes the propagation characteristics (spreading rate) of the laser beam. It cannot be neglected in the design of an optical train to be used with the beam. Truncation (aperturing) by an optic, in general, increases the M^2 factor of the beam.

The propagation equations (analogous to equations 10.7 and 36.8) for the mixed-mode beam $W(z)$ and $R(z)$ are as follows:

$$W(z) = W_0\left[1 + \left(\frac{zM^2\lambda}{\pi W_0^2}\right)^2\right]^{1/2} = W_0\left[1 + \left(\frac{z}{Z_R}\right)^2\right] \qquad (10.17)$$

and

$$R(z) = z\left[1 + \left(\frac{\pi W_0^2}{zM^2\lambda}\right)^2\right] = z\left[1 + \left(\frac{Z_R}{z}\right)^2\right]. \qquad (10.18)$$

The Rayleigh range remains the same for a mixed mode laser beam:

$$Z_R = \frac{\pi W_0^2}{M^2\lambda} = \frac{\pi w_0^2}{\lambda} = z_R. \qquad (10.19)$$

Now consider the consequences in coupling a high M^2 beam into a fiber. Fiber coupling is a task controlled by the product of the focal diameter ($2W_f$) and the focal convergence angle (θ_f). In the tight focusing limit, the focal diameter is proportional to the focal length f of the lens, and is inversely proportional to the diameter of the beam at the lens (i.e., $2W_f \propto f/D_{\text{lens}}$).

The lens-to-focus distance is f, and, since $f \times \theta_f$ is the beam diameter at distance f in the far field of the focus, $D_{\text{lens}} \propto f\theta_f$. Combining these proportionalities yields

$$W_f \theta_f = \text{constant}$$

for the fiber-coupling problem as stated above. The diameter-divergence product for the mixed-mode beam is M^2 larger than the fundamental mode beam in accordance with equations 10.15 and 10.16.

There is a threefold penalty associated with coupling a beam with a high M^2 into a fiber: 1) the focal length of the focusing lens must be a factor of $1/M^2$ shorter than that used with a fundamental-mode beam to obtain the same focal diameter at the fiber; 2) the numerical aperture (NA) of the focused beam will be higher than that of the fundamental beam (again by a factor of $1/M^2$) and may exceed the NA of the fiber; and 3) the depth of focus will be smaller by $1/M^2$ requiring a higher degree of precision and stability in the optical alignment.

APPLICATION NOTE

Stable vs Unstable Resonator Cavities

A stable resonator cavity is defined as one that self-focuses energy within the cavity back upon itself to create the typical Gaussian modes found in most traditional lasers. The criterion for a stable cavity is that

$$0 \le g_1 g_2 \le 1$$

where

$$g_1 = 1 - \frac{L}{R_1} \text{ and } g_2 = 1 - \frac{L}{R_2}$$

where R_1 and R_2 are the radii of the cavity mirrors and L is the mirror separation.

The mode volumes of stable resonator cavities are relatively small. This is fine when the excitation regions of a laser are also relatively small, as is the case with a HeNe or DPSS laser. However, for large-format high-energy industrial lasers, particularly those with high single-pass gain, stable resonators can limit the output. In these cases, unstable resonators, like the one shown in the illustration below, can generate higher output with better mode quality. In this case, the output coupling is determined by the ratio of the diameters of the output and high-reflecting mirrors, not the coating reflectivity. In the near field, the output looks like a doughnut, because the center of the beam is occluded by the output mirror. At a focus, however, the beam has most of the propagation characteristics of a fundamental-mode stable laser.

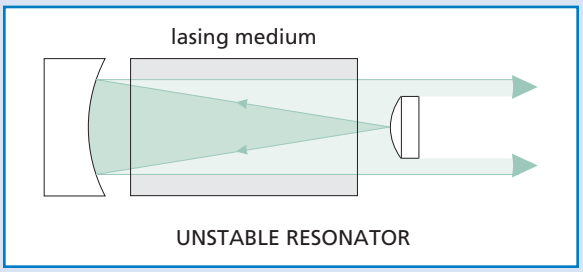

Unstable Resonator Design

Transverse Modes and Mode Control

The fundamental TEM_{00} mode is only one of many transverse modes that satisfies the condition that it be replicated each round-trip in the cavity. Figure 10.9 shows examples of the primary lower-order Hermite-Gaussian (rectangular) modes.

Note that the subscripts m and n in the mode designation TEM_{mn} are correlated to the number of nodes in the x and y directions. The propagation equation can also be written in cylindrical form in terms of radius (ρ) and angle (ϕ). The eigenmodes ($E_{\rho\phi}$) for this equation are a series of axially symmetric modes, which, for stable resonators, are closely approximated by Laguerre-Gaussian functions, denoted by $TEM_{\rho\phi}$. For the lowest-order mode, TEM_{00}, the Hermite-Gaussian and Laguerre-Gaussian functions are identical, but for higher-order modes, they differ significantly, as shown in figure 10.10.

The mode, TEM_{01}*, also known as the "bagel" or "doughnut" mode, is considered to be a superposition of the Hermite-Gaussian TEM_{10} and TEM_{01} modes, locked in phase and space quadrature. (See W.W. Rigrod, "Isolation of Axi-Symmetric Optical-Resonator Modes," *Applied Physics Letters*, Vol. 2 (1 Feb. '63), pages 51–53.)

In real-world lasers, the Hermite-Gaussian modes predominate since strain, slight misalignment, or contamination on the optics tends to drive the system toward rectangular coordinates. Nonetheless, the Laguerre-Gaussian TEM_{10} "target" or "bulls-eye" mode is clearly observed in well-aligned gas-ion and helium neon lasers with the appropriate limiting apertures.

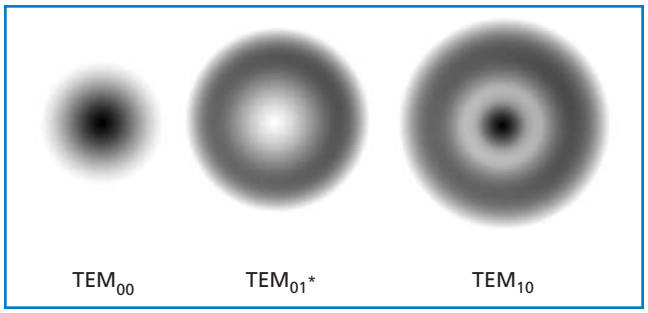

Figure 10.10 **Low-order axisymmetric resonator modes**

MODE CONTROL

The transverse modes for a given stable resonator have different beam diameters and divergences. The lower the order of the mode is, the smaller the beam diameter, the narrower the far-field divergence, and the lower the M^2 value. For example, the TEM_{01}* doughnut mode is approximately 1.5 times the diameter of the fundamental TEM_{00} mode, and the Laguerre TEM_{10} target mode is twice the diameter of the TEM_{00} mode. The theoretical M^2 values for the TEM_{00}, TEM_{01}*, and TEM_{10} modes are 1.0, 2.3, and 3.6, respectively (R. J. Freiberg et al., "Properties of Low Order Transverse Modes in Argon Ion Lasers"). Because of its smooth intensity profile, low divergence, and ability to be focused to a diffraction-limited spot, it is usually desirable to operate in the lowest-order mode possible, TEM_{00}. Lasers, however, tend to operate at the highest-order mode possible, either in addition to, or instead of, TEM_{00} because the larger beam diameter may allow them to extract more energy from the lasing medium.

The primary method for reducing the order of the lasing mode is to add sufficient loss to the higher-order modes so that they cannot oscillate without significantly increasing the losses at the desired lower-order mode. In most lasers this is accomplished by placing a fixed or variable aperture inside the laser cavity. Because of the significant differences in beam diameter, the aperture can cause significant diffraction losses for the higher-order modes without impacting the lower-order modes. As an example, consider the case of a typical argon-ion laser with a long-radius cavity and a variable mode-selecting aperture.

When the aperture is fully open, the laser oscillates in the axially symmetric TEM_{10} target mode. As the aperture is slowly reduced, the output changes smoothly to the TEM_{01}* doughnut mode, and finally to the TEM_{00} fundamental mode.

In many lasers, the limiting aperture is provided by the geometry of the laser itself. For example, by designing the cavity of a helium neon laser so that the diameter of the fundamental mode at the end of the laser bore is approximately 60 percent of the bore diameter, the laser will naturally operate in the TEM_{00} mode.

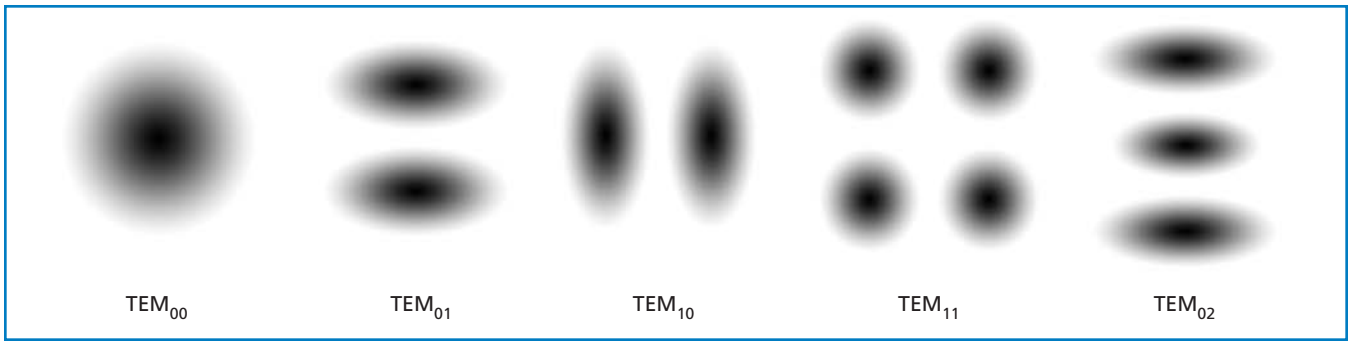

Figure 10.9 **Low-order Hermite-Gaussian resonator modes**

Single Axial Longitudinal Mode Operation

THEORY OF LONGITUDINAL MODES

In a laser cavity, the requirement that the field exactly reproduce itself in relative amplitude and phase each round-trip means that the only allowable laser wavelengths or frequencies are given by

$$\lambda = \frac{P}{N} \text{ or } \nu = \frac{Nc}{P} \qquad (10.20)$$

where λ is the laser wavelength, ν is the laser frequency, c is the speed of light in a vacuum, N is an integer whose value is determined by the lasing wavelength, and P is the effective perimeter optical path length of the beam as it makes one round-trip, taking into account the effects of the index of refraction. For a conventional two-mirror cavity in which the mirrors are separated by optical length L, these formulas revert to the familiar

$$\lambda = \frac{2L}{N} \text{ or } \nu = \frac{Nc}{2L}. \qquad (10.21)$$

These allowable frequencies are referred to as longitudinal modes. The frequency spacing between adjacent longitudinal modes is given by

$$\Delta\nu = \frac{c}{P}\ . \qquad (10.22)$$

As can be seen from equation 10.22, the shorter the laser cavity is, the greater the mode spacing will be. By differentiating the expression for ν with respect to P we arrive at

$$\delta\nu = -\frac{Nc}{P^2}\delta P \text{ or } \delta n = -\frac{Nc}{2L^2}\delta L. \qquad (10.23)$$

Consequently, for a helium neon laser operating at 632.8 nm, with a cavity length of 25 cm, the mode spacing is approximately 600 MHz, and a 100-nm change in cavity length will cause a given longitudinal mode to shift by approximately 190 MHz.

The number of longitudinal laser modes that are present in a laser depends primarily on two factors: the length of the laser cavity and the width of the gain envelope of the lasing medium. For example, the gain of the red helium neon laser is centered at 632.8 nm and has a full width at half maximum (FWHM) of approximately 1.4 GHz, meaning that, with a 25-cm laser cavity, only two or three longitudinal modes can be present simultaneously, and a change in cavity length of less than one micron will cause a given mode to "sweep" completely through the gain. Doubling the cavity length doubles the number of oscillating longitudinal modes that can fit under the gain curve doubles.

The gain of a gas-ion laser (e.g., argon or krypton) is approximately five times broader than that of a helium neon laser, and the cavity spacing is typically much greater, allowing many more modes to oscillate simultaneously.

A mode oscillating at a frequency near the peak of the gain will extract more energy from the gain medium than one oscillating at the fringes. This has a significant impact on the performance of a laser system because, as vibration and temperature changes cause small changes in the cavity length, modes sweep back and forth through the gain. A laser operating with only two or three longitudinal modes can experience power fluctuations of 10% or more, whereas a laser with ten or more longitudinal modes will see mode-sweeping fluctuations of 2 percent or less.

SELECTING A SINGLE LONGITUDINAL MODE

A laser that operates with a single longitudinal mode is called a single-frequency laser. There are two ways to force a conventional two-mirror laser to operate with a single longitudinal mode. The first is to design the laser with a short enough cavity that only a single mode can be sustained. For example, in the helium neon laser described above, a 10-cm cavity would allow only one mode to oscillate. This is not a practical approach for most gas lasers because, with the cavity short enough to suppress additional modes, there may be insufficient energy in the lasing medium to sustain any lasing action at all, and if there is lasing, the output will be very low.

The second method is to introduce a frequency-control element, typically a low-finesse Fabry-Perot etalon, into the laser cavity. The free spectral range of the etalon should be several times the width of the gain curve, and the reflectivity of the surfaces should be sufficient to provide 10 percent or greater loss at frequencies half a longitudinal mode spacing away from the etalon peak. The etalon is mounted at a slight angle to the optical axis of the laser to prevent parasitic oscillations between the etalon surfaces and the laser cavity.

Once the mode is selected, the challenge is to optimize and maintain its output power. Since the laser mode moves if the cavity length changes slightly, and the etalon pass band shifts if the etalon spacing varies slightly, it important that both be stabilized. Various mechanisms are used. Etalons can be passively stabilized by using zero-expansion spacers and thermally stabilized designs, or they can be thermally stabilized by placing the etalon in a precisely controlled oven. Likewise, the overall laser cavity can be passively stabilized, or, alternatively, the laser cavity can be actively stabilized by providing a servomechanism to control cavity length, as discussed in Frequency Stabilization.

The Ring Laser: The discussions above are limited to two-mirror standing-wave cavities. Some lasers operate naturally in a single longitudinal mode. For example, a ring laser cavity, (used in many dye and Ti:Sapphire lasers as well as in gyroscopic lasers) that has been constrained to oscillate in only one direction produces a traveling wave without the fixed nodes of the standing-wave laser. The traveling wave sweeps through the laser gain, utilizing all of the available energy and preventing the buildup of adjacent modes. Other lasers are "homogeneously broadened" allowing virtually instantaneous transfer of energy from one portion of the gain curve to another.

FREQUENCY STABILIZATION

The frequency output of a single-longitudinal-mode laser is stabilized by precisely controlling the laser cavity length. This can be accomplished passively by building an athermalized resonator structure and carefully controlling the laser environment to eliminate expansion, contraction, and vibration, or actively by using a mechanism to determine the frequency (either relatively or absolutely) and quickly adjusting the laser cavity length to maintain the frequency within the desired parameters.

A typical stabilization scheme is shown in figure 10.11. A portion of the laser output beam is directed into a low-finesse Fabry-Perot etalon and tuned to the side of the transmission band. The throughput is compared to a reference beam, as shown in the figure. If the laser frequency increases, the ratio of attenuated power to reference power increases.

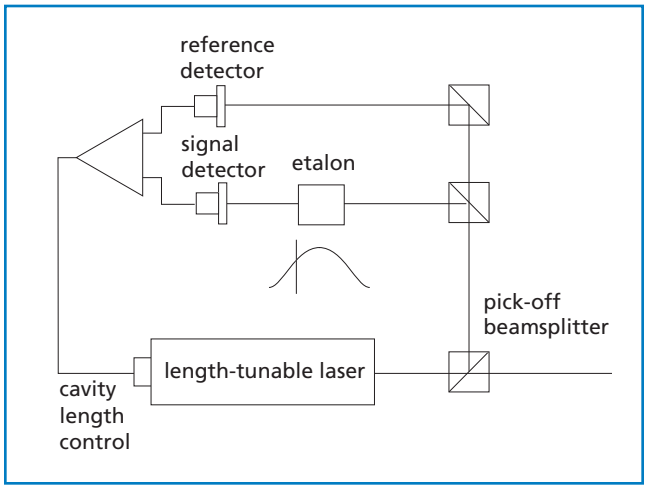

Figure 10.11 **Laser frequency stabilization scheme**

If the laser frequency decreases, the ratio decreases. In other words, the etalon is used to create a frequency discriminant that converts changes in frequency to changes in power. By "locking" the discriminant ratio at a specific value (e.g., 50 percent) and providing negative feedback to the device used to control cavity length, output frequency can be controlled. If the frequency increases from the preset value, the length of the laser cavity is increased to drive the frequency back to the set point. If the frequency decreases, the cavity length is decreased. The response time of the control electronics is determined by the characteristics of the laser system being stabilized.

Other techniques can be used to provide a discriminant. One common method used to provide an ultrastable, long-term reference is to replace the etalon with an absorption cell and stabilize the system to the saturated center of an appropriate transition. Another method, shown in figure 10.12, is used with commercial helium neon lasers. It takes advantage of the fact that, for an internal mirror tube, the adjacent modes are orthogonally polarized. The cavity length is designed so that two modes can oscillate under the gain curve. The two modes are separated outside the laser by a polarization-sensitive beamsplitter. Stabilizing the relative amplitude of the two beams stabilizes the frequency of both beams.

The cavity length changes needed to stabilize the laser cavity are very small. In principle, the maximum adjustment needed is that required to sweep the frequency through one free spectral range of the laser cavity (the cavity mode spacing). For the helium neon laser cavity described earlier, the required change is only 320 nm, well within the capability of piezoelectric actuators.

Commercially available systems can stabilize frequency output to 1 MHz or less. Laboratory systems that stabilize the frequency to a few kilohertz have been developed.

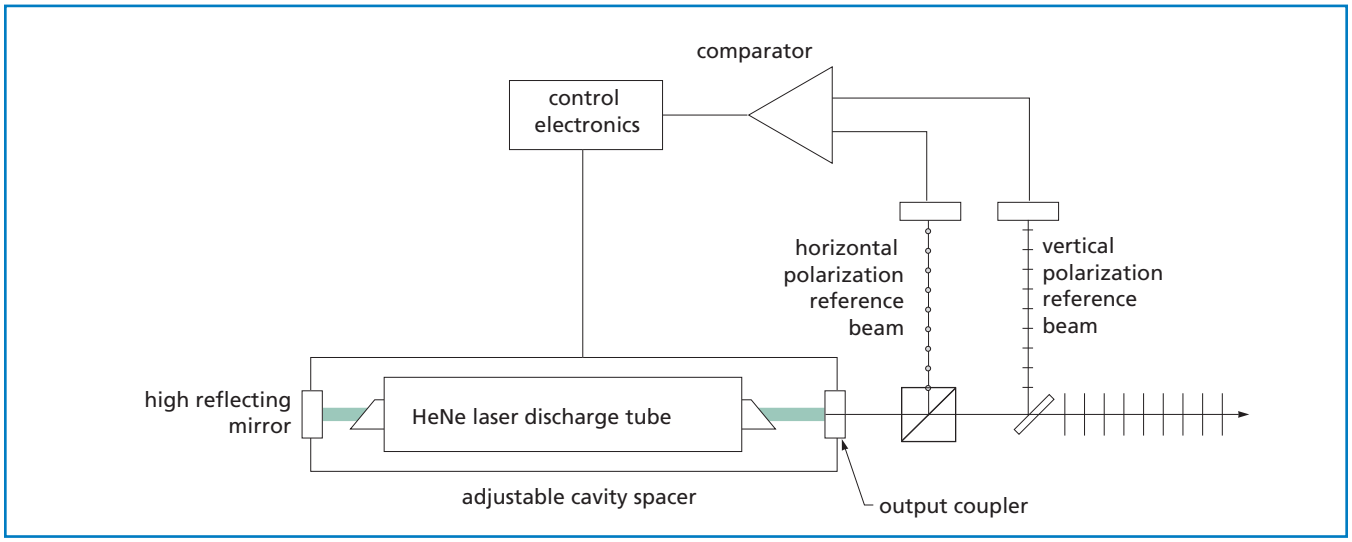

Figure 10.12 **Frequency stabilization for a helium neon laser**

Frequency and Amplitude Fluctuations

The output of a freely oscillating laser will fluctuate in both amplitude and frequency. Fluctuations of less than 0.1 Hz are commonly referred to as "drift"; faster fluctuations are termed "noise" or, when talking about sudden frequency shifts, "jitter."

The major sources of noise in a laser are fluctuations in the pumping source and changes in length or alignment caused by vibration, stress, and changes in temperature. For example, unfiltered line ripple can cause output fluctuations of 5 to 10 percent or more.

Likewise, a 10-μrad change in alignment can cause a 10-percent variation in output power, and, depending upon the laser, a 1-μm change in length can cause amplitude fluctuations of up to 50 percent (or more) and frequency fluctuations of several gigahertz.

High-frequency noise (>1 MHz) is caused primarily by "mode beating." Transverse Laguerre-Gaussian modes of adjacent order are separated by a calculable fraction of the longitudinal mode spacing, typically ~17 MHz in a 1-m resonator with long radius mirrors. If multiple transverse modes oscillate simultaneously, heterodyne interference effects, or "beats," will be observed at the difference frequencies. Likewise, mode beating can occur between longitudinal modes at frequencies of

$$\Delta\nu_{\text{longitudinal}} = \frac{c}{2L} = \frac{c}{2P} \quad . \qquad (10.24)$$

Mode beating can cause peak-to-peak power fluctuations of several percent. The only way to eliminate this noise component is to limit the laser output to a single transverse and single longitudinal mode.

Finally, when all other sources of noise have been eliminated, we are left with quantum noise, the noise generated by the spontaneous emission of photons from the upper laser level in the lasing medium. In most applications, this is inconsequential.

METHODS FOR SUPPRESSING AMPLITUDE NOISE AND DRIFT

Two primary methods are used to stabilize amplitude fluctuations in commercial lasers: automatic current control (ACC), also known as current regulation, and automatic power control (APC), also known as light regulation. In ACC, the current driving the pumping process passes through a stable sensing resistor, as shown in figure 10.13, and the voltage across this resistor is monitored. If the current through the resistor increases, the voltage drop across the resistor increases proportionately. Sensing circuitry compares this voltage to a reference and generates an error signal that causes the power supply to reduce the output current appropriately. If the current decreases, the inverse process occurs. ACC is an effective way to reduce noise generated by the power supply, including line ripple and fluctuations.

With APC, instead of monitoring the voltage across a sensing resistor, a small portion of the output power in the beam is diverted to a photodetector, as shown in figure 10.14, and the voltage generated by the

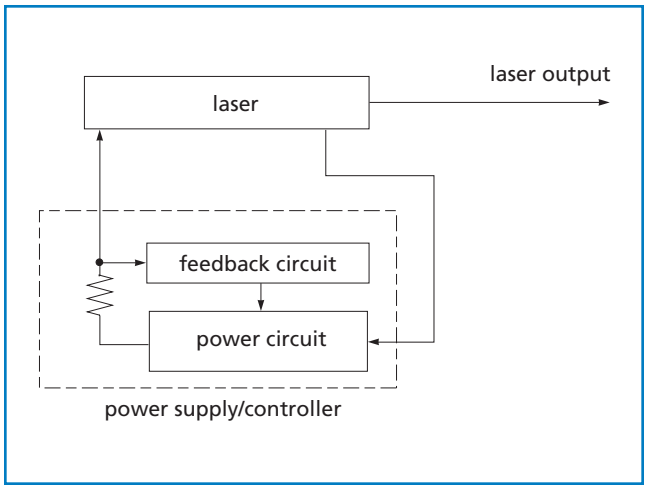

Figure 10.13 **Automatic current control schematic**

detector circuitry is compared to a reference. As output power fluctuates, the sensing circuitry generates an error signal that is used to make the appropriate corrections to maintain constant output.

Automatic current control effectively reduces amplitude fluctuations caused by the driving electronics, but it has no effect on amplitude fluctuations caused by vibration or misalignment. Automatic power control can effectively reduce power fluctuations from all sources. Neither of these control mechanisms has a large impact on frequency stability.

Not all continuous-wave lasers are amenable to APC as described above. For the technique to be effective, there must be a monotonic relationship between output power and a controllable parameter (typically current or voltage). For example, throughout the typical operating range of a gas-ion laser, an increase in current will increase the output power and vice versa. This is not the case for some lasers. The output of a helium neon laser is very insensitive to discharge current, and an increase in current may increase or decrease laser output. In a helium cadmium laser, where electrophoresis determines the density and uniformity of cadmium ions throughout the discharge, a slight change in discharge current in either direction can effectively kill lasing action.

If traditional means of APC are not suitable, the same result can be obtained by placing an acousto-optic modulator inside the laser cavity and using the error signal to control the amount of circulating power ejected from the cavity.

One consideration that is often overlooked in an APC system is the geometry of the light pickoff mechanism itself. One's first instinct is to insert the pickoff optic into the main beam at a 45-degree angle, so that the reference beam exits at a 90-degree angle. However, as shown in figure 10.15, for uncoated glass, there is almost a 10-percent difference in reflectivity for s and p polarization.

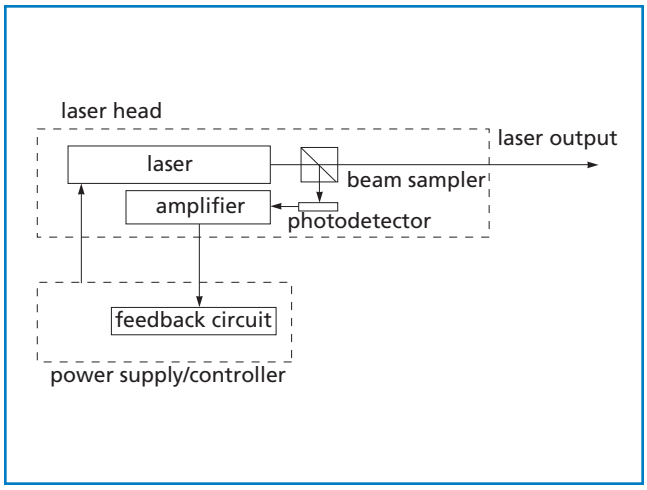

Figure 10.14 **Automatic power control schematic**

In a randomly polarized laser, the ratio of the *s* and *p* components is not necessarily stable, and using a 90-degree reference beam can actually increase amplitude fluctuations. This is of much less concern in a laser with a high degree of linear polarization (e.g., 500:1 or better), but even then there is a slight presence of the orthogonal polarization. Good practice dictates that the pickoff element be inserted at an angle of 25 degrees or less.

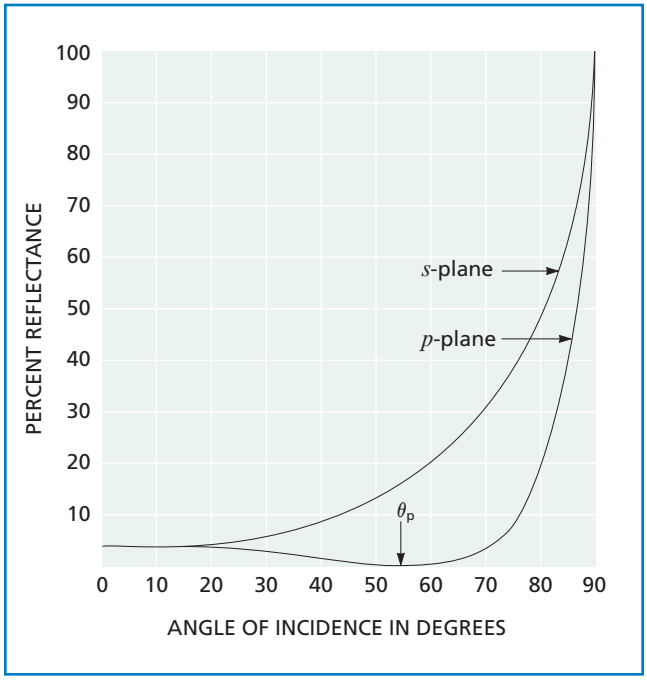

Figure 10.15 **Reflectivity of a glass surface vs. incidence angle for *s* and *p* polarization**

APPLICATION NOTE

Measuring Frequency Stability

The accepted method of measuring long-term frequency stability is to heterodyne the laser to be tested with another laser of equal or greater stability. By observing the variation of the resulting beat frequencies, the combined drift of the two lasers can be measured. The results will be no better than the sum of the two instabilities and will, therefore, provide a conservative measure of frequency drift.

In the charts below, a frequency-stabilized HeNe was hetrodyned with the output from a Zeeman-stabilized laser. The charts show the performance over one minute and over an eight-hour typical workday. The laser can be cycled over a 20°C temperature range without mode hopping.

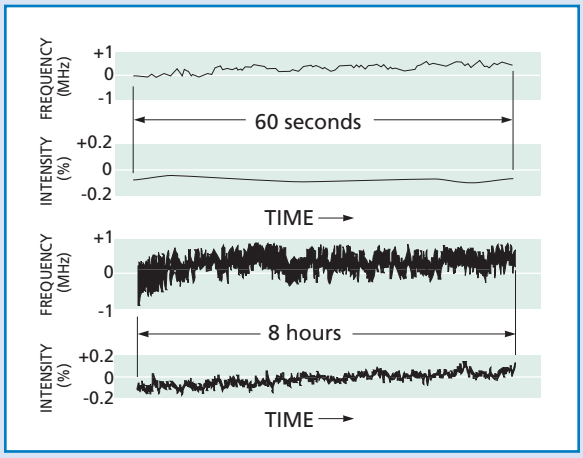

Short- and long-term frequency stability of a frequency stabilized helium neon laser

Tunable Operation

Many lasers can operate at more than one wavelength. Argon and krypton lasers can operate at discrete wavelengths ranging from the ultraviolet to the near infrared. Dye lasers can be continuously tuned over a spectrum of wavelengths determined by the fluorescence bandwidths of the specific dyes (typically about 150 nm). Alexandrite and titanium sapphire lasers can be tuned continuously over specific spectral regions.

To create a tunable laser, the cavity coatings must be sufficiently broadband to accommodate the entire tuning range, and a variable-wavelength tuning element must be introduced into the cavity, either between the cavity optics or replacing the high-reflecting optic, to introduce loss at undesired wavelengths.

Three tuning mechanisms are in general use: Littrow prisms, diffraction gratings, and birefringent filters. Littrow prisms (see figure 10.16) and their close relative, the full-dispersing prism, are used extensively with gas lasers that operate at discrete wavelengths. In its simplest form, the Littrow prism is a 30-60-90-degree prism with the surface opposite the 60-degree angle coated with a broadband high-reflecting coating. The prism is oriented so that the desired wavelength is reflected back along the optical axis, and the other wavelengths are dispersed off axis. By rotating the prism the retroreflected wavelength can be changed. In laser applications, the prism replaces the high-reflecting mirror, and the prism's angles are altered (typically to 34, 56, and 90 degrees) to minimize intracavity losses by having the beam enter the prism exactly at Brewster's angle. For higher-power lasers which require greater dispersion to separate closely spaced lines, the Littrow prism can be replaced by a full-dispersing prism coupled with a high reflecting mirror.

Gratings are used for laser systems that require a higher degree of dispersion than that of a full-dispersing prism.

Birefringent filters have come into general use for continuously tunable dye and Ti:Sapphire lasers, since they introduce significantly lower loss than do gratings. The filter is made from a thin, crystalline-quartz plate with its fast axis oriented in the plane of the plate. The filter, placed at Brewster's angle in the laser beam, acts like a weak etalon with a free spectral range wider than the gain curve of the lasing medium. Rotating the filter around the normal to its face shifts the transmission bands, tuning the laser. Since there are no coatings and the filter is at Brewster's angle (thereby polarizing the laser), there are no inherent cavity reflection losses at the peak of the transmission band. A single filter does not have as significant a line-narrowing effect as does a grating, but this can be overcome by stacking multiple filter plates together, with each successive plate having a smaller free spectral range.

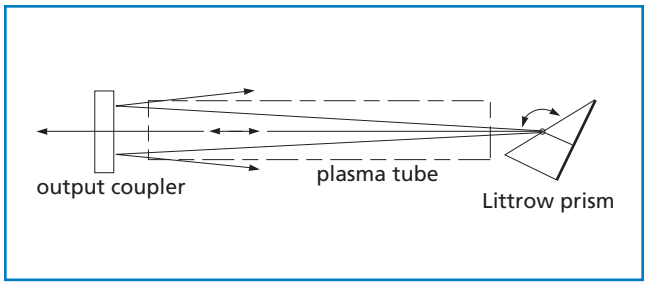

Figure 10.16 **Littrow prism used to select a single wavelength**

Prism-tunable-ion laser

Types of Lasers

Since the discovery of the laser, literally thousands of types of lasers have been discovered. As Arthur Shawlow is purported to have said, "Hit it hard enough and anything will lase." However, only a relative few of these lasers have found broadly based, practical applications.

Lasers can be broadly classified into four categories: gas discharge lasers, semiconductor diode lasers, optically pumped lasers, and "other," a category which includes chemical lasers, gas-dynamics lasers, x-ray lasers, combustion lasers, and others developed primarily for military applications. These lasers are not discussed further here.

GAS-DISCHARGE LASERS

In principle, gas-discharge lasers are inherently simple—fill a container with gas, put some mirrors around it, and strike a discharge. In practice, they are much more complex because the gas mix, discharge parameters, and container configuration must be specifically and carefully designed to create the proper conditions for a population inversion. Furthermore, careful consideration must be given to how the discharge will react with its container and with the laser optics. Finally, since the temperature of the gas can affect the discharge conditions, questions of cooling must be addressed.

Figure 10.17 below shows a cutaway of a helium neon laser, one of the simplest gas-discharge lasers. An electrical discharge is struck between the anode and cathode. The laser bore confines the discharge, creating the current densities needed to create the inversion. In this example, the laser mirrors are mounted to the ends of the tube and are effectively part of the gas container. In other cases, the mirrors are external to the container, and light enters and exits the chamber through Brewster's windows or extremely low-loss antireflection-coated normal windows. Because most gas-discharge lasers are operated at extremely low pressures, a getter is needed to remove the impurities generated by outgassing in the walls of the container or by erosion of the electrodes and bore caused by the discharge. The Brewster's window is used to linearly polarize the output of the laser.

The most common types of gas-discharge lasers are helium neon lasers, helium cadmium lasers (a metal-vapor laser), noble-gas ion lasers (argon, krypton), carbon-dioxide lasers, and the excimer-laser family. Each of these will be discussed briefly below.

Helium Neon Lasers

The helium neon (HeNe) laser, shown in figure 10.17, the second laser to be discovered, was the first to be used in volume applications. Today, millions of these lasers are in the field, and only semiconductor diode lasers are sold in greater quantity.

The HeNe laser operates in a high-voltage (kV), low-current (mA) glow discharge. Its most familiar output wavelength is 633 nm (red), but HeNe lasers are also available with output at 543 nm (green), 594 nm (yellow), 612 nm (orange), and 1523 nm (near infrared). Output power is low, ranging from a few tenths to tens of milliwatts, depending on the wavelength and size of the laser tube.

Helium is the major constituent (85 percent) of the gas mixture, but it is the neon component that is the actual lasing medium. The glow discharge pumps the helium atoms to an excited state that closely matches the upper energy levels of the neon atoms. This energy is then transferred to

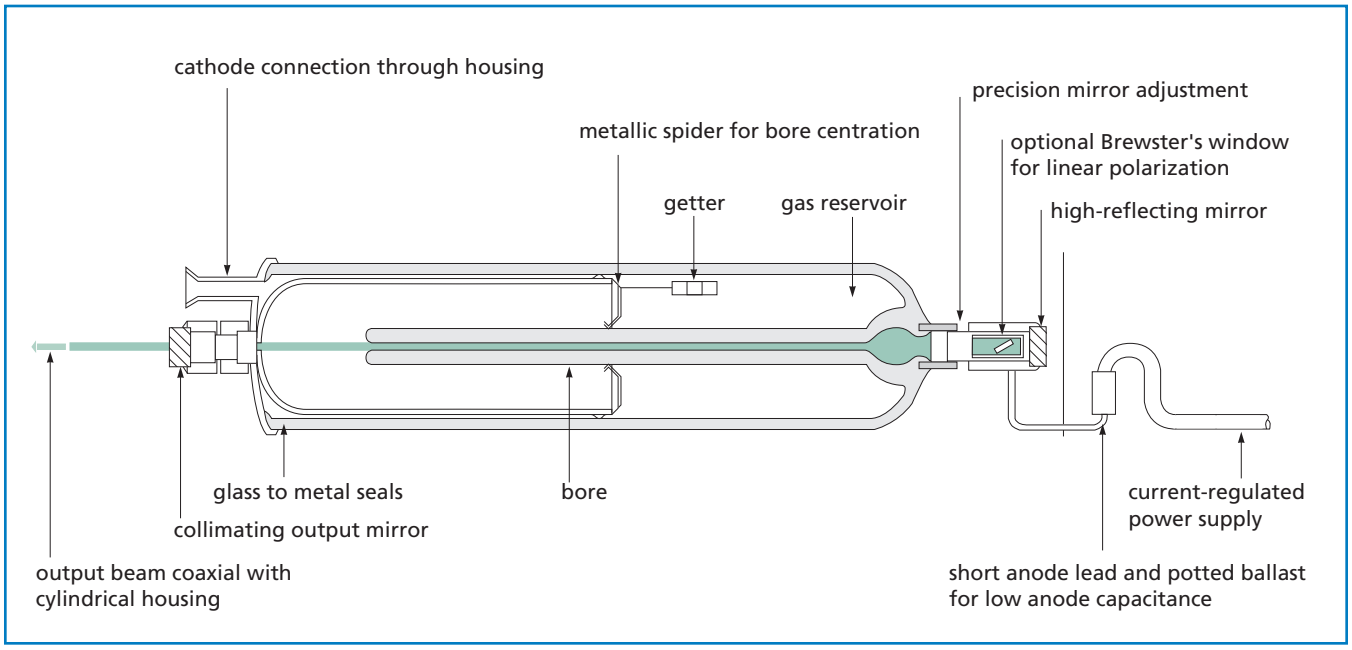

Figure 10.17 **Typical HeNe laser construction**

the neon atoms via collisions of the second kind (i.e., exciting the neon to a higher energy level as opposed to transferring the energy as kinetic motion). One characteristic of the glow discharge is its negative impedance (i.e., increasing the voltage decreases the current); consequently, to function with a standard current-regulated power supply, a ballast resistor must be used in series with the laser to make the overall impedance positive.

The popularity (and longevity) of the HeNe laser is based on five factors: they are (relative to other lasers) small and compact; they have the best inherent beam quality of any laser, producing a virtually pure single transverse mode beam ($M^2 < 1.05$); they are extremely long lived, with many examples of an operating life of 50,000 hours or more; they generate relatively little heat and are convection cooled easily in OEM packages; and they have a relatively low acquisition and operating cost.

Helium Cadmium Lasers

Helium cadmium (HeCd) lasers are, in many respects, similar to the HeNe laser with the exception that cadmium metal, the lasing medium, is solid at room temperature. The HeCd laser is a relatively economical, cw source for violet (442 nm) and ultraviolet (325 nm) output. Because of its excellent wavelength match to photopolymer and film sensitivity ranges, it is used extensively for three-dimensional stereolithography and holographic applications.

As mentioned above, cadmium, a metal, is solid at room temperature. For lasing to occur, the metal must be evaporated from a reservoir, as shown in figure 10.18, and then the vapor must be distributed uniformly down the laser bore. This is accomplished through a process called electrophoresis. Because cadmium will plate out on a cool surface, extreme care must be taken in the design of the laser to contain the cadmium and to protect the optics and windows from contamination, since even a slight film will introduce sufficient losses to stop lasing. The end of life usually occurs when cadmium is depleted from its reservoir.

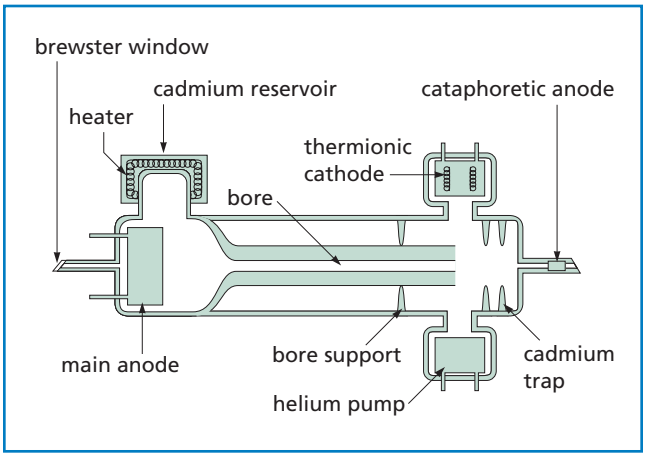

Figure 10.18 **Construction of a HeCd laser**

Noble-Gas Ion Lasers

The noble-gas ion lasers (argon-ion and krypton-ion), have been the mainstay of applications requiring high cw power in the visible, ultraviolet, and near-infrared spectral regions. High-power water-cooled systems can be found in research laboratories around the world; lower-power air-cooled systems are used in a wide variety of OEM applications. Argon-ion lasers are available with output up to 7 W in the ultraviolet (333–354 nm) and 25 W or more in the visible regions (454–515 nm), with primary output at 488 nm (blue) and 514 nm (green). Krypton-ion lasers have their primary output at 568 nm (yellow), 647 nm (red), and 752 nm (near infrared). Mixed-gas lasers combine both argon and krypton to produce lasers with a wider spectral coverage.

Unlike the HeNe laser, ion lasers operate with a high-intensity low-pressure arc discharge (low voltage, high current). A 20-W visible laser will require 10 kW or more power input, virtually all of which is deposited in the laser head as heat which must be removed from the system by some cooling mechanism. Furthermore, the current densities in the bore, which can be as high as 10^5A/cm^2, place large stresses on the bore materials.

Ion lasers can be broken into two groups: high-power (1–20 + W) water-cooled lasers and low-power air-cooled lasers. Both are shown schematically in figure 10.19.

The main features of both lasers are the same. Both use a coiled, directly-heated dispenser cathode to supply the current; both have a gas return path that counteracts gas pumping (non-uniform gas pressure throughout the length of the tube caused by the charged particles moving toward the electrodes).

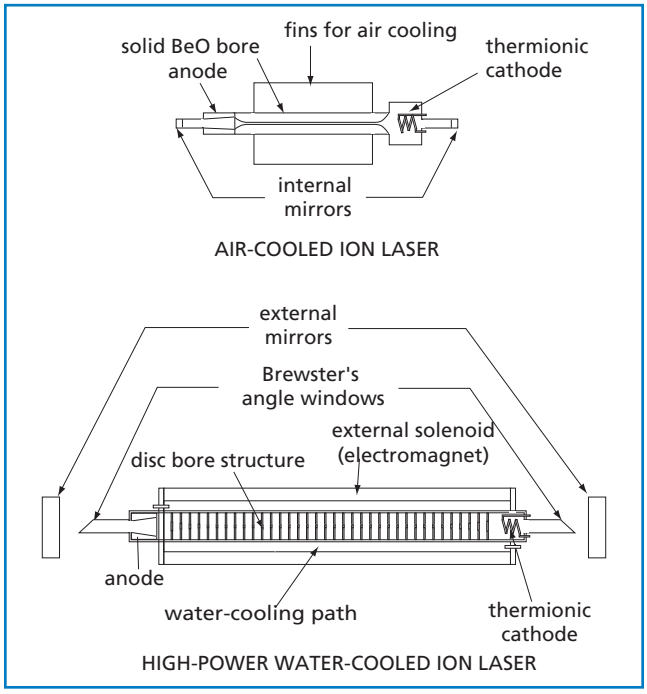

Figure 10.19 **Air-cooled and water-cooled ion lasers**

The bore of an air-cooled system is always made of beryllium oxide (BeO), a ceramic known for its ability to conduct heat. A fin structure is attached to the outside of the ceramic bore, and a blower removes the generated heat, typically less than 1 kW.

Water-cooled systems are available with either BeO bores or a construction wherein tungsten discs are attached to a thin-walled ceramic tube surrounded by a water jacket. The heat from the discs is conducted through the walls of the tube to the surrounding water. The entire bore structure is surrounded by a solenoid electromagnet, which compresses the discharge to increase current density and minimize bore erosion.

The main life-limiting factors in ion lasers are cathode depletion and gas consumption. The intense discharge drives atoms into the walls of the discharge tube where they are lost to the discharge. Over time the tube pressure will decrease, causing the discharge to become unstable. This is particularly a problem with krypton-ion lasers. Water-cooled systems typically have some refill mechanism to keep the pressure constant. Air-cooled systems typically do not, limiting their practical operating life to approximately 5000 operating hours.

Carbon Dioxide Lasers

Because of their ability to produce very high power with relative efficiency, carbon dioxide (CO_2) lasers are used primarily for materials-processing applications. The standard output of these lasers is at 10.6 μm, and output power can range from less than 1 W to more than 10 kW.

Unlike atomic lasers, CO_2 lasers work with molecular transitions (vibrational and rotational states) which lie at low enough energy levels that they can be populated thermally, and an increase in the gas temperature, caused by the discharge, will cause a decrease in the inversion level, reducing output power. To counter this effect, high-power cw CO_2 lasers use flowing gas technology to remove hot gas from the discharge region and replace it with cooled (or cooler) gas. With pulsed CO_2 lasers that use transverse excitation, the problem is even more severe, because, until the heated gas between the electrodes is cooled, a new discharge pulse cannot form properly.

A variety of types of CO_2 lasers are available. High-power pulsed and cw lasers typically use a transverse gas flow with fans which move the gas through a laminar-flow discharge region, into a cooling region, and back again (see figure 10.20). Low-power lasers most often use waveguide structures, coupled with radio-frequency excitation, to produce small, compact systems.

Excimer Lasers

The term excimer or "excited dimer" refers to a molecular complex of two atoms which is stable (bound) only in an electronically excited state. These lasers, which are available only as pulsed lasers, produce intense output in the ultraviolet and deep ultraviolet. The lasers in this family are XeFl (351 nm), XeCl (308 nm), KrF (248 nm), KrCl (222 nm), ArF (193 nm), and F_2 (157 nm). They are used extensively in photolithography, micromachining, and medical (refractive eye surgery) applications.

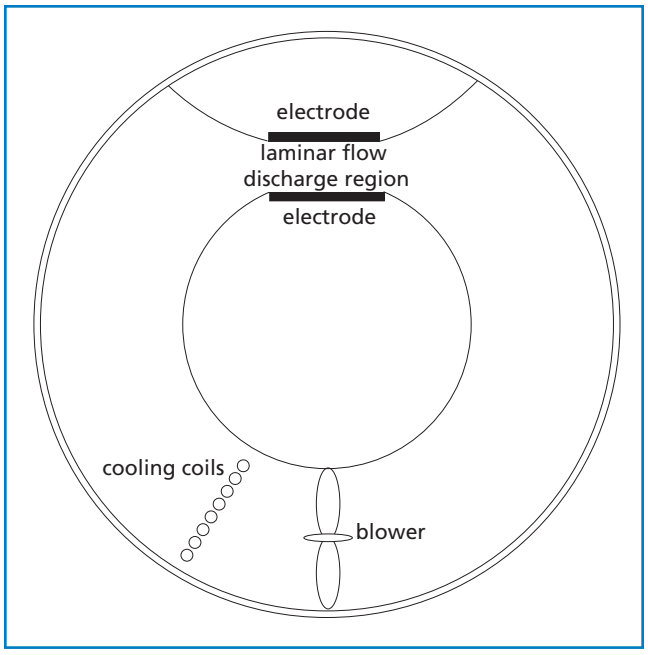

Figure 10.20 **Schematics of transverse flow CO_2 laser system**

At first glance, the construction of an excimer laser is very similar to that of a transverse-flow, pulsed CO_2 laser. However, the major difference is that the gases in the system are extremely corrosive and great care must be taken in the selection and passivation of materials to minimize their corrosive effects. A system built for CO_2 would fail in minutes, if not seconds.

The principal advantage of an excimer laser is its very short wavelength. The excimer output beam can be focused to a spot diameter that is approximately 40 times smaller than the CO_2 laser beam with the same beam quality. Furthermore, whereas the long CO_2 wavelength removes material thermally via evaporation (boiling off material), the excimer lasers with wavelengths near 200 nm remove material via ablation (breaking molecules apart), without any thermal damage to the surrounding material.

SEMICONDUCTOR DIODE LASERS

The means of generating optical gain in a diode laser, the recombination of injected holes and electrons (and consequent emission of photons) in a forward-biased semiconductor pn junction, represents the direct conversion of electricity to light. This is a very efficient process, and practical diode laser devices reach a 50-percent electrical-to-optical power conversion rate, at least an order of magnitude larger than most other lasers. Over the past 20 years, the trend has been one of a gradual replacement of other laser types by diode laser based–solutions, as the considerable challenges to engineering with diode lasers have been met. At the same

time the compactness and the low power consumption of diode lasers have enabled important new applications such as storing information in compact discs and DVDs, and the practical high-speed, broadband transmission of information over optical fibers, a central component of the Internet.

Construction of a double-heterostructure diode laser

In addition to a means to create optical gain, a laser requires a feedback mechanism, a pair of mirrors to repeatedly circulate the light through the gain medium to build up the resulting beam by stimulated emission. The stripe structures needed to make a laser diode chip are formed on a single crystal wafer using the standard photolithographic patterning techniques of the semiconductor industry. The substrate crystal axes are first oriented relative to the patterning such that, after fabrication, a natural cleavage plane is normal to the stripe direction, and cleaving both ends of the chip provides a pair of plane, aligned crystal surfaces that act as a Fabry-Perot resonator for optical feedback. These mirrors use either the Fresnel reflectivity of the facet (often sufficient because of the high gain of diode lasers), or they can be dielectric coated to other reflectivities. This might be desired, for instance, to protect against damage from the high irradiance at the facets. This geometry gives the familiar edge-emitting diode laser (see figure 10.21).

The semiconductor crystal must be defect free to avoid scattering of carriers and of light. To grow crystal layers without defects, commercial semiconductor lasers use III-V compounds, elements taken from those columns of the periodic table. These form varying alloys with the addition of dopants that can be lattice-matched to each other and to the initial crystal substrate. The band gap of the semiconductor chosen determines the lasing wavelength region. There are three main families: GaN-based lasers with UV-blue outputs, GaAs-based lasers with red-near infrared outputs, and InP-based lasers with infrared outputs. These base crystals are precisely doped with Ga, Al, In, As, and P to precisely control the band gap and index of refraction of the layers in the diode structure.

These compounds are direct band-gap semiconductors with efficient recombination of injected holes and electrons because no phonons (lattice vibrations) are required to conserve momentum in the recombination interaction. The injection layers surrounding the junction, the cladding layers, can be indirect band-gap semiconductors (where phonons are involved).

To make a planar waveguide that concentrates the light in the junction region (confinement between the top and bottom horizontal planes of the active region in figure 10.21), the cladding layers are made of an alloy of lower refractive index (larger band gap) than the active junction region. This is then termed a double-heterostructure (DH) laser. The output power of the laser is horizontally polarized because the reflectivity of the planar waveguide is higher for the polarization direction parallel to the junction plane. Because the junction region is thin for efficient recombination (typically 0.1 μm), some light spreads into the cladding layers which are therefore made relatively thick (typically 1 μm) for adequate light confinement.

Gain guiding and index guiding in diode lasers

To confine the light laterally (between planes perpendicular to the junction plane), two main methods (with many variants) are used. The first and simplest puts a narrow conductive stripe on the p-side of the device to limit the injected current to a line, giving a gain-guided laser. There is some spreading of current under the stripe, and the light is restricted only by absorption in the unpumped regions of the junction. The transverse mode of the laser light is therefore not tightly controlled. Many high-power diode lasers, used for instance in side-pumping another solid-state laser (where mode control is less critical), are gain guided.

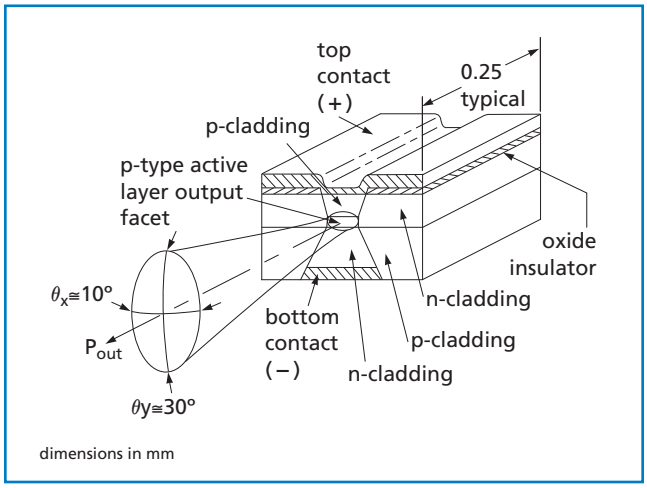

Figure 10.21 **Schematic of a double heterostructure index-guided diode laser**

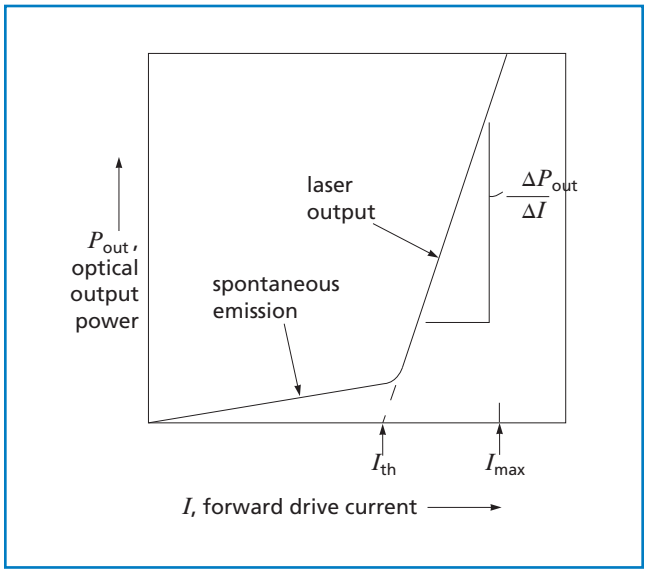

Figure 10.22 **Definition of threshold current, I_{th}, and slope efficiency from the curve of light output, P_{out} vs drive current I**

More efficient lateral laser mode control is achieved by fabricating, with multiple photolithographic, epitaxial, and etching steps, regions of low index of refraction on either side of the lasing stripe (the two lateral n-cladding regions in the upper half of figure 10.21). This confines the light by waveguiding between planes perpendicular to the junction plane as well giving an index-guided laser. These lasers produce a stable single transverse mode of lowest order, as required in data storage applications to read compact discs, and telecommunications applications where coupling into a fiber optic is important.

Threshold current and slope efficiency definitions

Output power from a diode laser increases linearly with the drive current excess above the threshold current (see figure 10.22). This steeply rising light output curve is extrapolated backward to the zero light output intercept to define the threshold current; the weak incoherent light emission for currents below threshold is due to the spontaneous recombination of carriers such as occurs in LEDs.

When divided by the drive voltage V, the slope of the output vs current curve yields the differential (above threshold) electrical-to-optical power conversion efficiency (also termed the slope or quantum efficiency) which ranges from 50 to 80 percent for various devices.

$$\text{Slope efficiency} = \frac{\Delta P_{\text{out}}}{V \Delta I} \ . \qquad (10.25)$$

Fabrication methods and quantum wells

Three types of epitaxial crystal growth are employed in fabricating the layers of semiconductor alloys for diode laser chips: liquid phase epitaxy (LPE), metal-organic chemical vapor deposition (MOCVD), and molecular beam epitaxy (MBE).

Most early diode lasers were made by the LPE process, and it is still in use for many commercial diode lasers and LEDs. In this process, a heated, saturated solution is placed in contact with the substrate, and it is cooled, leaving an epitaxial film grown on the substrate. High-quality crystal layers are readily produced by this technique, but it is hard to control alloy composition. Furthermore, making thin layers is difficult. Because Quantum well (QW) structures, discussed below, require very thin layers. The LPE process is not appropriate for these devices; they are fabricated using the MOCVD or MBE process.

In the MOCVD process, gases transport the reactants to the heated substrate, where they decompose and the epitaxial layer slowly grows. In the high-vacuum MBE process, the reactants are evaporated onto the substrate, giving a very slow, controlled epitaxial growth. The equipment for MBE is more expensive, and the process is slower making this process most suitable for critical and complex devices of low production volume.

The emergence of the MOCVD and MBE processes made possible improved diode lasers employing quantum well structures as their active regions. A quantum well is a layer of semiconductor of low electron (or hole) potential energy between two other layers of higher potential energy. The well layer is made thin enough, typically less than 0.01 μm, to be comparable in size to the Bohr radius of the electron (or hole) in the material. This brings in quantum effects—the confined carrier acts, in the direction perpendicular to the layer plane, as a one-dimensional particle in a potential well. In practical terms, the density of carriers is greatly increased in this QW structure, and the laser threshold current decreases by an order of magnitude. The laser's active region is effectively an engineered, man-made material whose properties can be designed.

There is a disadvantage to QW lasers: the active region is too thin to make a reasonable waveguide. This problem is solved by inserting intermediate layers of graded index between the QW and both cladding layers. This is termed the graded-index separate-confinement heterostructure (GRINSCH) since the carriers are confined to the QW while the laser mode is confined by the surrounding layers. The electrical and optical confinements are separate. For higher output power, several QWs separated by buffer layers can be stacked on top of one another,—a multiple quantum well (MQW) structure. A structure with only one quantum layer is designated a single quantum well (SQL) to distinguish it from a MQW.

The lasing wavelength in QW lasers is determined by both the bulk band gap and the first quantized energy levels; it can be tuned by varying the QW thickness. Further adjustment of the wavelength is possible with strained quantum QWs. If an epitaxial layer is kept below a critical thickness, an alloy with a lattice mismatch to the substrate will distort its lattice (in the direction normal to the substrate) to match the substrate lattice instead of causing misfit dislocations. The strain in the lattice of the resulting QW changes its band gap, an effect taken advantage of to put the lasing wavelength into a desired region.

Wavelength stabilization with distributed, surface-emitting output geometries

The wavelength of a AlGaAs diode laser tunes with substrate temperature at a rate of about 0.07 nm/°C, a rapid enough rate that many applications require the baseplate of the device to be mounted on a temperature controlled thermoelectric cooler to maintain wavelength stability. Wavelength, threshold current, and efficiency are all sensitive to changes in temperature. If the laser baseplate temperature is allowed to drift, in addition to this long-term shift in wavelength, the output oscillation will jump between drifting longitudinal cavity modes and thus exhibit small, rapid, discontinuous changes in wavelength and/or output power which often are undesirable.

To address this issue, gratings are fabricated into the laser, either at the ends of the gain stripe to create a distributed Bragg reflector (DBR) structure, or along the whole length of the gain region to create a distributed feedback (DFB) structure. The grating has a period on the order of 200 nm and is fabricated using interferometric techniques. (The beam from an argon or HeCd laser is split into two; the beams are are then overlapped to create fringes, which in turn are used to expose the photoresist in the photolythography process.) The gratings work by providing a small reflected feedback at each index step. The single frequency whose multiple fed-back reflections add up in phase determines the lasing wavelength and stabilizes it against changes in drive current and baseplate temperature.

Because the laser operates in a single frequency, noise is also reduced. DBR and DFB lasers are used extensively as telecommunication light sources.

The DFB laser is an edge emitter. In the second-order gratings fabricated in both DFB and DBR lasers, the first-order diffraction is perpendicular to the surface of the grating. By providing an output window on one of the gratings in a DBR laser, the output can be brought out through the surface of the chip, i.e., a surface-emitting laser.

Recently, another surface-emitting structure, the vertical-cavity surface-emitting laser (VCSEL), has come into use in telecommunication links. In this structure (see figure 10.23) multilayer mirrors are fabricated on the top and bottom of the QW gain region to give feedback. Consequently, the laser output is perpendicular to the active QW plane.

The epitaxial growth process of this structure is more difficult than that for edge emitters This is because provision must be made to channel current flow around the mirrors to reduce device resistance (for clarity, the bypass channels are omitted in figure 10.23) and because precise control of the mirror layer thicknesses is needed to locate standing wave peaks at the QW active layer(s). Countering these drawbacks, by having no facets to cleave, these lasers have a similar topology to LEDs. They can be tested at the wafer level and packaged using similar low-cost manufacturing methods. In addition, VCSELs have large-area circular beams (defined by the circular limiting aperture of the mirrors) and low threshold currents so they couple well into optical fibers and fit well in low-power (~1 mW) communication system applications.

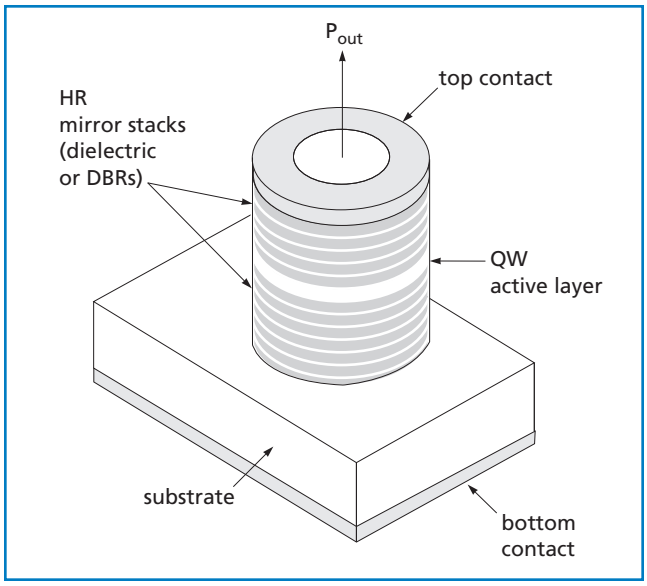

Figure 10.23 **Schematic of the VCSEL structure, with light emitted perpendicular to the active layer**

Diode laser beam conditioning

Because the emitting aperture is small on a typical diode laser, beam divergences are large. For example, the emitting area for the index-guided laser shown in figure 10.21 might be as small as 3 × 1 μm, resulting in diverges of 10 × 30 degrees. The optics needed to collimate this beam or to focus it into a fiber must work at a high numerical aperture, resulting in potential lens aberrations, and requiring critical focusing because of short depth of field. Focal lengths must be kept short as well or the optics rapidly become large. The beam inself is elliptical and may be astigmatic. It is often desireable to first circularize the beam spot with an anamorphic prism pair or cylinder lens before coupling the laser output into an optical train. Higher-power lasers with high-order modes cause additional problems when coupling their beams into a fiber or optical system. A wide variety of specialized components are available to address these issues, from molded miniature aspherical lenses to hyperbolic profile fiber cylinder lenses, but all require critical focusing adjustment in their mounting into the laser diode housing. For these reasons many diode lasers are offered with beam-correcting optics built in by the manufacturer who has the appropriate tooling for the task. Typically these lasers are available as collimated units, or as fiber-coupled ("pigtailed") devices.

High-power diode lasers

Single transverse mode diode lasers are limited to 200 mW or less of output power by their small emitting aperture. The facet area is so small (about 3 × 1 μm) that this power still represents a high irradiance (~7 MW/cm^2). The output is limited to this level to stay safely under the irradiance that would cause damage to the facet.

Enlarging the emitting area with an increase of the lateral width of the active stripe is the most common method of increasing the laser output power, but this also relaxes the single transverse mode constraint. Multiple transverse lateral modes, filaments, and lateral mode instabilities arise as the stripe width increases. For example, in a GaAs laser running at 808 nm, the output power rises linearly from 500 mW to 4 W as the lateral width of the emitting aperture increases from 50 to 500 μm. Howerer, the M^2 value of the beam in this plane increases from 22 to 210. The M^2 increase makes it difficult to couple these devices to fibers, but they find considerable application in pumping solid-state laser chips designed to accept a high-numerical-aperture focus.

The pump diode lasers for even higher-power DPSS lasers are made as linear arrays of 20 or more stripe emitters integrated side by side on a 1-cm-long semiconductor bar. The bar is mounted in a water-cooled housing to handle the heat load from the high drive current. These diode laser arrays provide from 20 to 40 W of continuous output power at wavelengths matching the absorption bands of different laser crystals (e.g., 808 nm for pumping Nd:YAG lasers). The individual stripe emissions are not coherently related, but bars can be used to side pump a laser rod, just as the arc lamps they replace formerly did. Another common delivery geometry is a bundle of multimode fibers, fanned into a line of fibers on one end with each fiber butted against an individual stripe on the bar, with the other end of the bundle gathered into a circular grouping. This converts the bar output into a round spot focusable onto the end of the crystal to be pumped.

Finally, for even more output, a few to a dozen bars are mounted like a deck of cards one on top of another in a water-cooled package, connected in series electrically, and sold as a stacked array. These can deliver in excess of 500 W output power from one device.

Packaging, power supplies, and reliability

For low-power lasers, the industry uses standard semiconductor device package designs, hermetically sealed with an output window. Lasers with higher power dissipation come with a copper baseplate for attachment to a finned heat sink or thermoelectric cooler (TEC). Many are offered coupled into a fiber at the manufacturing plant in a pigtailed package (with an output fiber attached) because of the criticality in mounting the coupling optics as mentioned above.

Careful heat sinking is very important because all the major device parameters—wavelength, threshold current, slope efficiency, and lifetime—depend on device temperature (the cooler, the better). Temperature-servoed TECs are preferred for stable operation with the temperature sensor for feedback mounted close to the diode laser.

Diode lasers are susceptible to permanent damage from static electricity discharges or indeed any voltage transient. Their low operating voltage (~2 V) and ability to respond at high speed means that a static discharge transient can be a drive current spike above the maximum safe level and result in catastrophic facet damage. All the usual antistatic electricity precautions should be taken in working with diode lasers: cotton gloves, conductive gowns, grounded wrist straps, work tables, soldering irons, and so on. Correspondingly, the drive current power supply should be filtered against surges and include "slow starting" circuitry to avoid transients.

Diode lasers degrade with high power and long operating hours as crystal defects migrate and grow, causing dark lines or spots in the output mode pattern and increases in threshold current or decreases in slope efficiency. The best way to prolong life is to keep the laser baseplate running cool. Remember that accelerated life tests are run by operating at high baseplate temperature. Expectations for the median life of a device are set from measurements of large populations—individual devices can still suddenly fail. Nevertheless, the industry expectations today for standard diode lasers run within their ratings is ~10^5 hours of operation for low-power diode lasers and perhaps an order of magnitude less for the high-power versions.

Summary of applications

The applications mentioned in the discussion above, and a few others, are summarized in the following table and ordered by wavelength. The newer GaN lasers provide low power (10–100 mW) blue and UV wavelengths finding applications as excitation sources for biomedical fluorescence studies (DNA sequencing, confocal microscopy). The dominant application for diode lasers is as readouts for optical data storage, followed by growing numbers in use in telecommunications. For high-power (>1 W) diode lasers, the main application is as optical pumps for other solid state lasers.

OPTICALLY PUMPED LASERS

Optically pumped lasers use photons of light to directly pump the lasing medium to the upper energy levels. The very first laser, based on a synthetic ruby crystal, was optically pumped. Optically pumped lasers can be separated into two broad categories: lamp-pumped and laser pumped. In a lamp-pumped laser, the lasing medium, usually a solid-state crystal, is placed near a high-intensity lamp and the two are surrounded by an elliptical reflector that focuses the light from the lamp into the crystal, as shown in figure 10.24. In laser-pumped systems, the light from another laser is focused into a crystal (or a stream of dye), as shown in figure 10.25.

In general, ignoring the efficiency of the pump laser itself, laser pumping is a much more efficient mechanism than lamp pumping because the wavelength of the pump laser can be closely matched to specific absorption bands of the lasing medium, whereas most of the light from a broad-spectrum lamp is not usefully absorbed in the gain medium and merely results in heat that must be removed from the system. Furthermore, the size of the laser pump beam can be tightly controlled, serving as a gain aperture for improving the output mode characteristics of the pumped laser medium. On the other hand, laser pumping is often not suitable for high-energy applications where large laser crystals are required.

Diode-pumped solid-state (DPSS) lasers, a class of laser-pumped lasers, will be discussed in detail below.

DIODE-PUMPED SOLID STATE LASERS

The DPSS laser revolution

The optical difficulties encountered with diode lasers—difficulty in coupling to the high divergence light, poor mode quality in the slow axis of wide-stripe lasers, low output power from single-transverse-mode lasers— led to a new philosophy (figure 10.26) about how best to use these efficient, long-lived, compact light sources. This concept, championed in the 1980s by a group at Stanford University headed by Prof. Bob Byer, has been termed the diode-pumped solid state (DPSS) laser revolution.

The logic is simple. The primary light source (the diode laser) pumps another laser (an infrared crystal laser) to convert to a good mode, the beam of which is wavelength converted (by nonlinear optics techniques) to a visible output. The diode laser source replaces the discharge lamp for optically pumping the gain crystal in a traditional, high-efficiency, infrared laser. The infrared beam is generated in that independent resonator with a good mode, and consequently it can be efficiently converted with an intracavity nonlinear crystal to a visible beam with a good mode. Though power is lost at each step, the result is still a single-mode visible beam generated with a total electrical-to-optical conversion efficiency of several percent. These DPSS lasers are replacing the older visible gas lasers whose conversion efficiencies rarely reach 0.1 percent.

End- and side-pumping geometries

The first DPSS lasers were made by focusing the diode light from a single laser diode emitter through the high-reflector coating (at 1.06 μm) on the

Diode Laser Applications

Wavelength λ (nm)	Lattice Material*	Application
375	GaN	Biomedical fluorescence
405	GaN	Biomedical fluorescence, DVD mastering, direct-to-plate
440	GaN	Biomedical fluorescence, HeCd laser replacement
473	GaN	Biomedical fluorescence
488	GaN	Biomedical fluorescence
635–640	GaAs	Pointers, alignment, HeNe laser replacement
650–680	GaAs	Biomedical fluorescence, barcode scanners, pointers, alignment, surgical
780	GaAs	Audio CD readouts
785	GaAs	Raman spectroscopy
808	GaAs	Optical pumps for Nd:YAG lasers, thermal printing
940	InP	Optical pumps for Yb:YAG lasers
980	InP	Optical pumps for Er fiber telecom amplifiers
1310	InP	Input source for telecom short-wavelength channels, OCT
1455	InP	Optical pump for Raman gain in standard telecom fiber
1550	InP	Input source for telecom long-wavelength channels

*Ga, Al, In, P dopants are added to form the required layered strutures

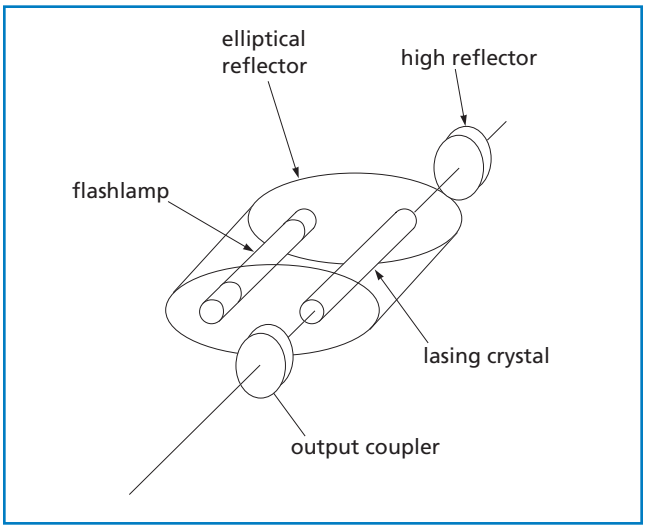

Figure 10.24 **Schematic of a lamp-pumped laser**

end of the Nd:YAG rod. This "end-pumping" geometry provide good overlap between the pumped volume and the lasing volume, but it limited the pump power to that available from single-mode diode emitters.

In order to increase laser output and reduce cost (diode lasers suitable for end pumping are twice as expensive as diode laser arrays), diode arrays were mounted along the length of the laser rod. However, because of poor overlap of the pump beam with the 1.06 μm beam, the efficiency of this "side-pumping" technique was only half that of end-pumping geometries. No pump diode cost savings resulted.

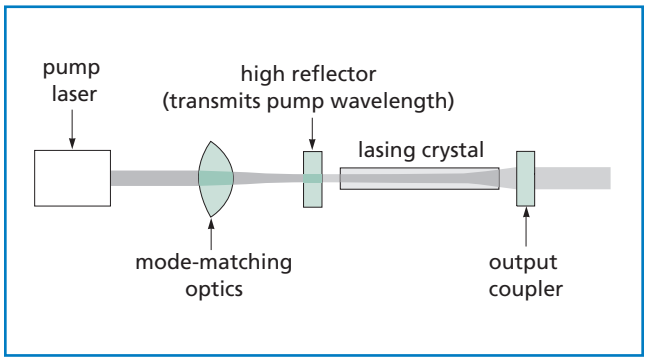

Figure 10.25 **Schematic of a laser-pumped laser**

Then in the late 1980s two advances were made. First, a variety of new laser rod materials, better tailored to take advantage of diode laser pumps, were introduced. Nd:YVO_4 crystals have five times the gain cross section of Nd:YAG, and the Nd can be doped into this crystal at much higher concentrations. This decreases the absorption depths in the crystal from cm to mm, easing the collimation or focusing quality required of the pump beam. This crystal had been known, but could be grown only to small dimensions, which is acceptable for diode-pumped crystals. Another crystal introduced was Yb:YAG, pumped at 980 nm and lasing at 1.03 μm leaving very little residual heat in the crystal per optical pumping cycle and allowing small chips of this material to be pumped at high levels.

Second, means were devised to make micro-cylindrical lenses (focal lengths less than a mm) with the correct surfaces (one type is a hyperbolic profile) for collimating or reducing the fast-axis divergence of the diode laser output. With good tooling and beam characterization these are correctly

positioned in the diode beam and bonded in place to the diode housing. This allows more conventional lenses, of smaller numerical aperture, to be used in subsequent pump light manipulations.

End-pumping with bars

With these two new degrees of freedom, laser designers realized they could create optical trains that would give them end-pumping system efficiencies (achieve good overlap between pump and lasing modes) with diode arrays as pump sources to obtain a lower diode cost per watt in their systems. This produced an explosion of unique DPSS laser designs generically described as "end-pumping with bars."

Figure 10.27 shows the example previously mentioned, delivering the array light through a fiber bundle, with the fibers at one end spread out to butt align with the linear stripes of an array, and the other end of the bundle gathered to an approximately circular spot.

Although the circular spot is large, its focal image, formed with high numerical aperture (NA) optics, is small enough to satisfactorily overlap the IR cavity laser mode. The small depth of focus, from the high NA optics, is inconsequential here because of the short absorption depth in the $Nd:YVO_4$ laser crystal. The laser head can be disconnected from the diode modules at the fiber coupler without loss of alignment.

In another example, an even higher-NA optic (comprising a cylinder lens and a molded aspheric lens) was used to directly focus the 1-cm width of a micro-lensed array bar onto the end of a $Nd:YVO_4$ gain crystal. This produced an oblong pump spot, but good overlap with the IR cavity mode was achieved by altering the infrared cavity (inserting two intracavity beam expansion prisms in that arm) to produce a 5:1 elliptical cavity mode in the gain crystal. Another design used a nonimaging pyramidal "lens duct" to bring in the pump light from a diode laser stack to the end of a gain crystal. Yet another brought light from several arrays into a lasing rod centered in a diffuse-reflecting cavity by means of several planar (glass-slide) waveguides, each piercing a different sector of the reflector sidewall. These are but a few of the design approaches that have been successfully taken.

488 nm Solid-State Lasers

A variety of solid-state approaches are used to achieve 488 nm output, a popular wavelength used for excitation of fluorophores in biotech applications. One approach, optically pumped semiconductor lasers (OPSL), utilizes a diode to pump a solid state gain chip atop a Bragg reflector which then passes through an external non linear crystal to double the 976 nm fundamental radiation down to 488 nm. The OPSL approach typically requires a large number of components and more complex fabrication, coating and assembly compared to other approaches that follow. True DPSS approaches utilize a pump diode, two gain mediums and a frequency mixing of 914 and 1064 nm radiation along with a single frequency generator to produce 491 nm, a wavelength close, but not quite optimum for the narrow bandwidth filters and fluorophores used in these applications. Directly doubled diodes (DDD) offer the best of both worlds by generating both the optimum 488 nm wavelength and a minimal number of components to achieve it (see figure 10.28). DDD's use a pump diode at 976 nm and unlike other doubled diode laser approaches which use external cavity components, an integrated nonlinear crystal/cavity. This approach results in a laser that is highly efficient, robust and volume manufacturable.

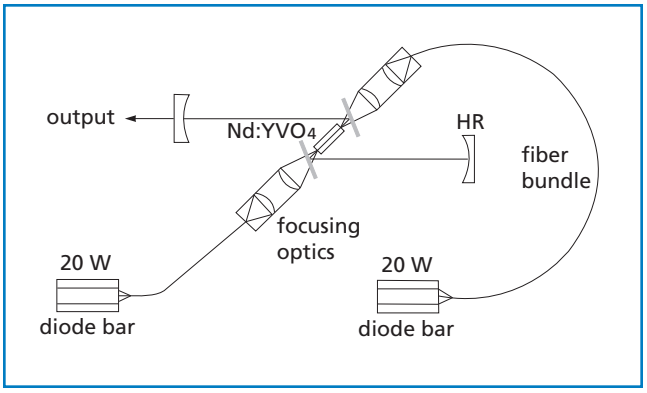

Figure 10.27 **Schematic of an "end-pumping with bars" geometry using fiber bundle delivery, one of many variants on the DPSS laser theme**

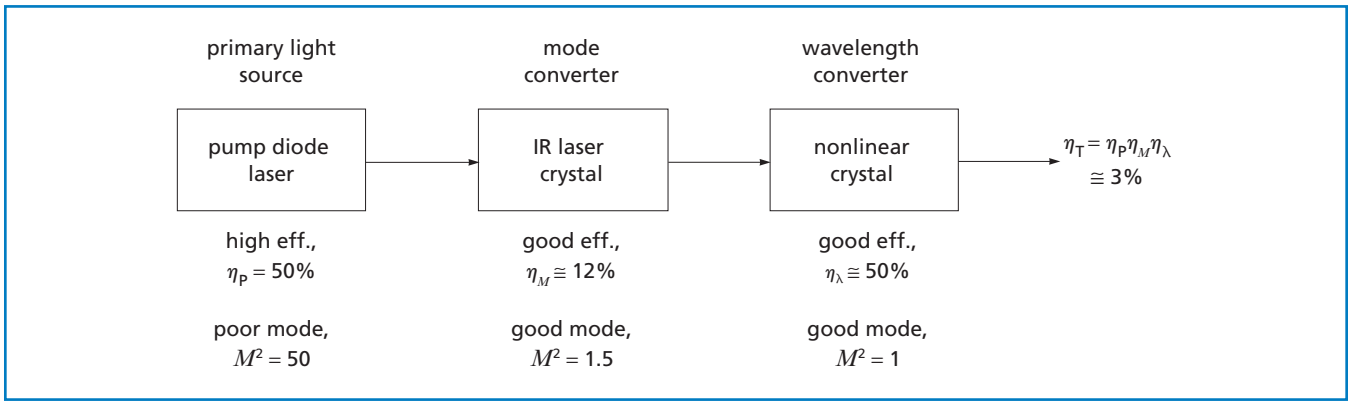

Figure 10.26 **The logic for DPSS lasers**

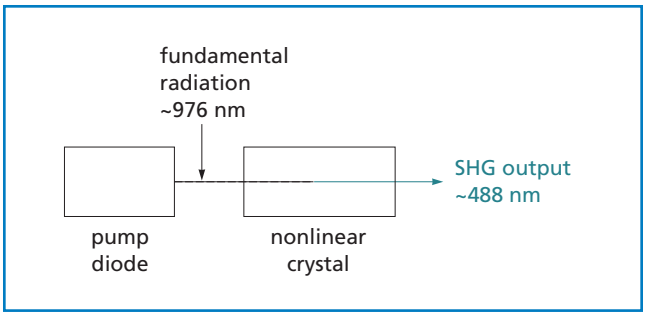

Figure 10.28 **Directly doubled diodes**

Microchip lasers

Another procedure that can be used to make available potentially inexpensive, mass produced, low power, visible output DPSS lasers is mimicking semiconductor chip processing methods. In the late 1980s, MIT Lincoln Labs took this approach and created the "microchip" laser. A thin Nd:YVO_4 plate is polished flat and then diced into ~2-mm-square chips. Each of these chips is then optically contacted to similar, flat, diced KTP doubling crystal plates to make a cube. Prior to dicing, the surfaces that will become the outer cube surfaces are coated for high reflectivity at 1.06 μm. When single-mode diode laser pump light is focused through one mirrored end of the cube, the heat produced makes a thermally-induced waveguide that creates a stable cavity for IR lasing. Since the KTP crystal is within this cavity, the IR lasing is converted to a 532-nm (green) output beam with 10's of milliwats of output. The diode temperature must be controlled to maintain a stable pump wavelength and thermal waveguide. In addition, the cube temperature should be stabilized. Because of the short cavity, the IR laser operates at a single longitudinal mode, and the cavity length must be thermally tuned to keep the mode at the peak of the gain curve. Laser operation in a single frequency suppresses "green noise," discussed next.

The "green noise" problem

As the early DPSS laser designs giving visible-output beams were being introduced, it became apparent that there was a problem unique to this architecture. The visible output power, 532 nm in the green spectrum, could break into high-frequency chaotic oscillations of nearly 100-percent peak-to-peak amplitude. This was named the "green problem" by Tom Baer (then at Spectra-Physics), who in 1986 showed the effect to be due to the dominance of sum-frequency-mixing (SFM) terms coupling different longitudinal modes over second-harmonic-generation (SHG) terms, in the nonlinear conversion step from IR to visible output. Several conditions (all met by the new laser designs) lead to this effect: (1) the IR laser cavity is short (~10 cm or less) with only a few longitudinal modes oscillating, (2) nonlinear conversion efficiencies are high (20% or more), and (3) nonlinear phase-matching bandwidths span several longitudinal mode spacings (true of the commonly used KTP doubling crystal). Then the sum frequency mixing output losses couple the longitudinal modes in relaxation oscillations where the turn on of one mode turns off another.

Two early solutions to this problem emerged. The first is to make the IR cavity long enough to give hundreds of oscillating modes, so that the noise terms average to insignificance as in a long gas laser. The second is to make the IR oscillation run on a single longitudinal mode so that there are no SFM terms. This can be done by using intracavity frequency control elements such as an etalon, or by using a ring cavity (with a Faraday-effect biasing element to maintain the direction of light travel around the ring). Ring cavities eliminate the standing-wave interference effect of linear cavities, termed "spatial hole burning," and the laser runs single frequency when this is done. As more experience was gained with DPSS laser design, other clever solutions to the "green problem" were found, tailored to each particular device and often held as trade secrets. It can be surmised that these involve precise control of wavelength, spatial hole burning, beam polarization, and cavity-element optical path differences to reduce the strength of longitudinal mode SFM terms.

Uniqueness of DPSS laser designs and laser reliability

Unlike the gas lasers they replace, no universal approach is applied in the details of different DPSS laser designs and laser models. There is a large variety of solutions to the major problems, many solutions are unique, and many are held as proprietary. Major design differences are found in:

- the means for optically coupling the pumping light into the gain medium,
- the management of the thermal lens produced by absorption of the pump light in the cavity,
- the control of green noise,
- the strain-free mounting, heat sinking, and placement of the small lasing and nonlinear crystals in the laser cavity, and
- the hermetic sealing of the laser cavity to protect the often delicate crystals and critical alignments of components

Note that because the intracavity space must be hermetically sealed there usually is no field repair, maintenance, or adjustment of a DPSS laser head. If it fails, it is returned to the manufacturer.

It is evident that DPSS lasers are a lot less generic than the gas lasers they replace. For a problem with a particular laser model, there may be no standard solution available in the technical literature. With so many variables, there often are surprises when new designs are first manufactured and introduced. Under these circumstances, the user is advised to pick a supplier with a record of years of consistent manufacture, who has over time dealt with his own unique set of component and assembly problems. If this advice is followed, then the expectation with current products is that a new DPSS laser will operate reliably for 10,000 hours or more.

561 nm DPSS Laser

The newest addition to the CVI Melles Griot laser product line is a DPSS laser with yellow output at 561 nm, an ideal excitation wavelength for biomedical fluorescence.

This 16.5-cm-long laser head delivers up to 75 mW of output power, and consumes less than 10 W of wall plug power at 25 mW and less than 18 W at 75 mW. The laser is pumped by a single stripe diode laser. Frequency-selective elements in the cavity limit IR oscillation to the 1.123 μm Nd line (one of the weaker lines in the 1.064 μm manifold) and constrain this oscillation to a single longitudinal mode. The output is low noise (0.5 percent rms) with excellent mode quality ($M^2 < 1.2$). Polarization is vertical with respect to the mounting surface with an extinction ratio of >100:1.

An example of a DPSS laser product line—the CVI Melles Griot visible output lasers

Figure 10.29 depicts the mix of laser crystals, laser operating wavelengths, and doubling crystals generating the four visible output wavelengths of the present CVI Melles Griot product line of continuous wave DPSS lasers.

Other notable DPSS lasers

A brief discussion of three other significant DPSS laser developments conclude this section.

The *Er-doped fiber amplifier (EDFA)* is not a laser, but it is an optically pumped amplifier for the 1550-nm long-wavelength long-haul fiberoptic channels that make the worldwide web possible. Pumping an Er-doped silica fiber with 980-nm diode laser light inverts the populations of energy levels in the Er ions to provide gain for optical telecommunication signals run through the same fiber. This optical amplifier is much simpler than the discrete electronic repeaters it replaced. A Lucent Technologies executive expressed the importance of this when he said: "What broke [wavelength division multiplexing telecommunications] free was the invention of the [EDFA] optical amplifier."

The *Q-switched industrial DPSS laser* is a 1-W-average-power, ultraviolet (355 nm), high-repetition-rate (30 kHz) system. Output is obtained by doubling the 1.064-μm Q-switched fundamental to green at 532 nm, and then mixing the green beam with the residual transmitted infrared to 355 nm. This process is straightforward in a high-peak-power pulsed beam—just a matter of inserting the appropriate doubling and tripling crystals. What is remarkable is that DPSS laser designs have matured sufficiently to make this possible in a hands-off, long lived, system rugged enough to survive and be useful in an industrial environment.

The *double-clad fiber laser* is shown in figure 10.30. Fiber lasers work by optically pumping (with a diode laser) a doped fiber and adding mirrors for feedback at either end of the fiber. In the dual-clad fiber, the Yb-doped single-mode fiber core is surrounded by a large diameter cladding (with a corrugated star-shaped cross section in the figure) that is itself clad by a low-index polymer coating. Diode laser light at 940 nm is readily launched into and guided in the large diameter outer cladding, and the corrugated cross-section of this fiber suppresses the helical ray modes of propagation that would have poor overlap with the inner core. Over the length of the fiber, the pump light is absorbed by the single-mode core, and high- power lasing near 1.03 μm in a low- order mode is produced. The quantum efficiency of the Yb lasing cycle (ratio of pump wavelength to lasing wavelength) is 91 percent, which leaves little heat deposited in the fiber. Over 1 kW of output at 80-percent slope efficiency has been produced in such a fiber laser. These will become important laser sources for industrial applications.

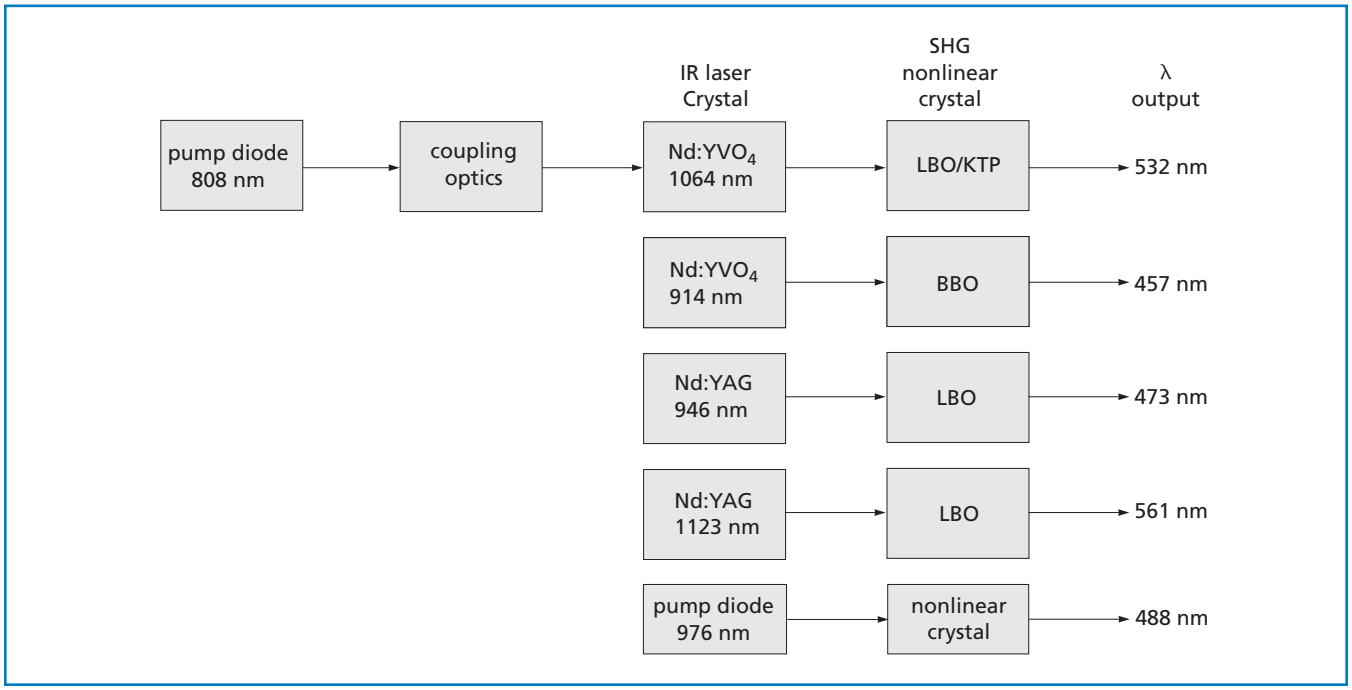

Figure 10.29 **CVI Melles Griot solid-state laser optical trains for producing five different visible output wavelengths**

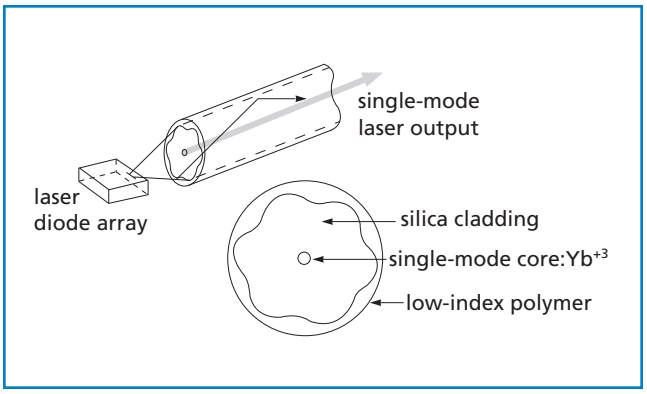

Figure 10.30 **Schematic diagram of the structure of a double-clad fiber, and the method of pumping the inner core by direct illumination into the large diameter of the outer cladding**

561-nm yellow diode pumped solid-state laser

Available Wavelengths

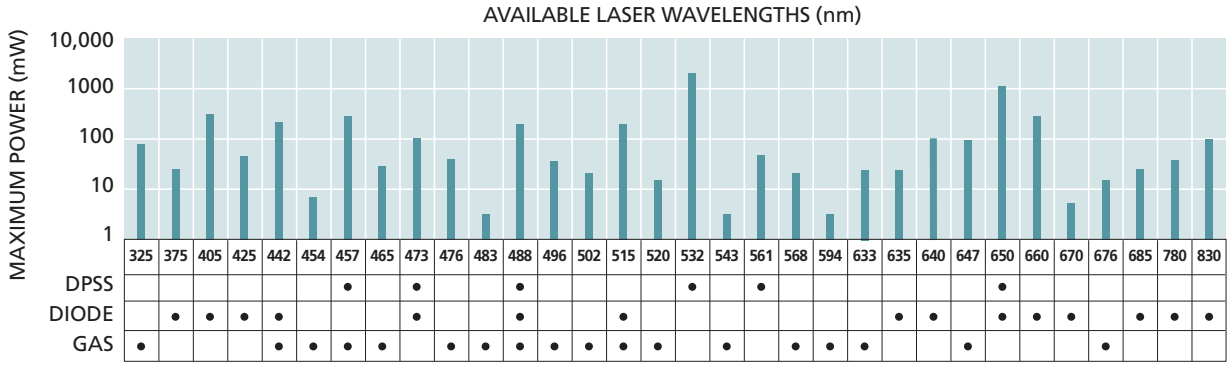

Technology and spectral offering

Laser Applications

Lasers have become so much a part of daily life that many people may not realize how ubiquitous they are. Every home with a CD player has a laser; hardware stores are now selling a wide variety of laser levels; many, if not most, computers, printers, and copiers are using laser technology. Laser applications are so numerous that it would be fruitless to try to list them all; however; one can give some illustrative examples of how lasers are used today.

INDUSTRIAL APPLICATIONS

High-power lasers have long been used for cutting and welding materials. Today the frames of automobiles are assembled using laser welding robots, complex cardboard boxes are made with laser-cut dies, and lasers are routinely used to engrave numbers and codes on a wide variety of products. Some less well-known applications include three-dimensional stereolithography and photolithography.

Three-Dimensional Stereolithography

Often a designer, having created a complex part on a CAD machine, needs to make a prototype component to check out the dimensions and fit. In many cases, it is not necessary for the prototype to be made of the specified (final) material for this checking step, but having a part to check quickly is important. This is where rapid prototyping, i.e., three-dimensional stereolithography, comes in. The stereolithography machine consists of a bath of liquid photopolymer, an ultraviolet laser, beam-handling optics, and computer control (see figure 10.31). When the laser beam is absorbed in the photopolymer, the polymer solidifies at the focal point of the beam. The component design is fed directly from the CAD program to the stereolithography computer. The laser is scanned through the polymer, creating, layer by layer, a solid, three-dimensional model of the part.

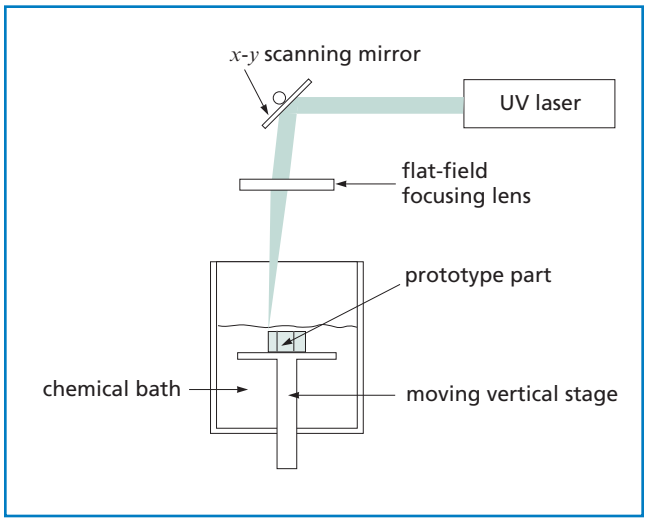

Figure 10.31 **A laser stereolithography system for rapid prototyping of three-dimensional parts**

Photolithography

Lasers are used throughout the manufacture of semiconductor devices, but nowhere are they more important than in exposing photoresist through the masks used for creating the circuits themselves. Originally, ultraviolet mercury lamps were used as the light sources to expose the photoresist, but as features became smaller and more complex devices were put on a single wafer, the mercury lamp's wavelengths were too long to create the features. In the 1990's, manufacturers started to switch to ultraviolet lasers operating at approximately 300 nm to expose the photoresist. Manufacturers are now using wavelengths as short as 193 nm to get the resolution needed for today's semiconductor integrated circuit applications.

Marking and Scribing

Lasers are used extensively in production to apply indelible, human and machine-readable marks and codes to a wide variety of products and packaging. Typical applications include marking semiconductor wafers for identification and lot control, removing the black overlay on numeric display pads, engraving gift items, and scribing solar cells and semiconductor wafers.

The basic marking system consists of a laser, a scanning head, a flat-field focusing lens, and computer control. The computer turns the laser beam on and off (either directly or through a modulator) as it is scanned over the surface to make the mark. Depending upon the application, scanning may occur in a raster pattern (typical for making dot-matrix marks) or in a cursive pattern, with the beam creating letters one at a time. The mark itself results either from ablation of the surface of the material, or by a photochemically induced change in the color of the material. Another marking technique, used with high-energy pulsed CO_2 and excimer lasers, is to shine the light through a mask containing the marking pattern and focusing the resulting image onto the marking surface.

Laser scribing is similar to laser marking, except that the scan pattern is typically rectilinear, and the goal is to create microscoring along the scan lines so that the substrate can be easily broken apart.

A wide variety of materials, including metal, wood, glass, silicon, and rubber, are amenable to laser marking and scribing. Each material has different absorption and thermal characteristic, and some even have directional preferences due to crystalline structure. Consequently, the type of laser used depends, to some extent, on the material to be marked (e.g., glass transmits the 1.06 μm output from a YAG laser but absorbs the 10.6 μm output from a CO_2 laser). Other considerations are the size of the pattern, the speed of the scan, cosmetic quality, and cost.

Currently, most volume marking applications are performed with lamp-pumped YAG-based pulsed or Q-switched lasers. Pulsed and cw CO_2 lasers make up the bulk of the remainder. However, DPSS and fiber lasers are encroaching on this field owing to their higher reliability and lower

operating cost. Because of their very short wavelengths (100–300 nm), excimer lasers are used in applications requiring extremely high resolution, or whose materials would thermally damage at longer wavelengths.

Noncontact measurement

There are many types of laser-based noncontact measurement techniques in use today including scatter measurement, polarimetry and ellipsometry, and interferometric measurement.

Scatter Measurement: In the semiconductor industry, patterns of material are deposited on a wafer substrate using photolithographic processes. Defects on the wafer can result in poor reliability, disconnects in circuitry, or complete circuit failure. Consequently manufacturers need to map the wafer to determine the defects' location and size so that they can either be eliminated or avoided. To do this, they scan the wafer with a laser and measure backscatter with a very sensitive photodetector array.

Lasers used in this application have to have excellent pointing stability, constant wavelength and power stability to calculate the correct size of the defects through complex algorithms, and low noise so the little scatter the defect makes can be distinguished from the background laser light. Blue 488-nm argon ion lasers have been the laser of choice for many years. However; as lithography has shifted to shorter and shorter ultraviolet wavelengths, however, we are beginning to see the metrologic techniques for wafer defect measurement also moving to shorter wavelengths. Ultraviolet diode and solid-state lasers are likely to replace the ion laser in the next generation of instruments.

Polarimetry and Ellipsometry: The optical phase thickness of a thin film can be carefully measured using polarimetry or ellipsometry (see figure 10.32). A beam of known polarization and phase state enters the thin film layer at an angle. The thin film has a known index of refraction. The measured phase change in the reflected beam is then correlated to an optical phase thickness for that layer using the known index of refraction. This technique can also be used with a thicker transparent media, such as glass, where changes in the polarization and phase state of a beam scanned across the substrate indicate variations in index of refraction due to inclusions or stress-induced birefringence. The most common lasers used in these applications are violet, red and near infrared single-emitter laser diodes and mid-visible diode-pumped solid-state lasers owing to their cw output, low noise, and compact sizes.

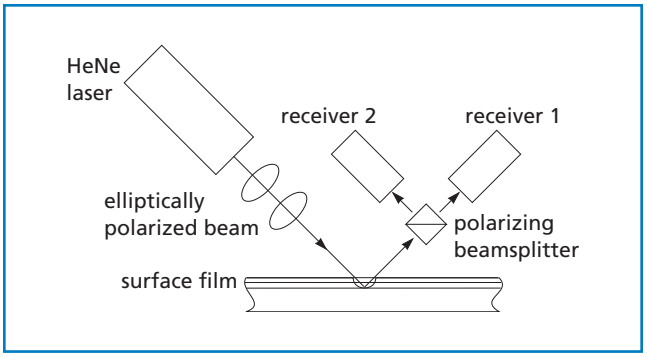

Figure 10.32 **Surface film thickness measurement**

Interferometric Measurement: Interferometric measurement can be used for high-resolution position measurement as well as for measuring waveform deformation of optical beams as they pass through a component or system (see figure 10.33).

The technique uses the wave periodicity of the light beam as a very fine ruler. The position of an object in the path of the beam is computed from the phase of the light reflected from it. Interference between the object beam and a reference beam provides measureable intensity variations which yield this phase information. Distance and velocity measurement can be performed for moving objects as long as the fringe-recording mechanism is paced with it.

Typical applications of this technique include positioning of masks for the lithography process, mirror distance correlation within an FTIR spectrometer, optical feedback in many high-resolution positioning systems, and determining the alignment and flatness of hard disk drive heads.

For applications requiring measurement over a long path length, lasers with a single longitudinal mode and long coherence length are often required. In these cases, frequency-stabilized helium neon lasers or a solid-state lasers with frequency selective elements are used.

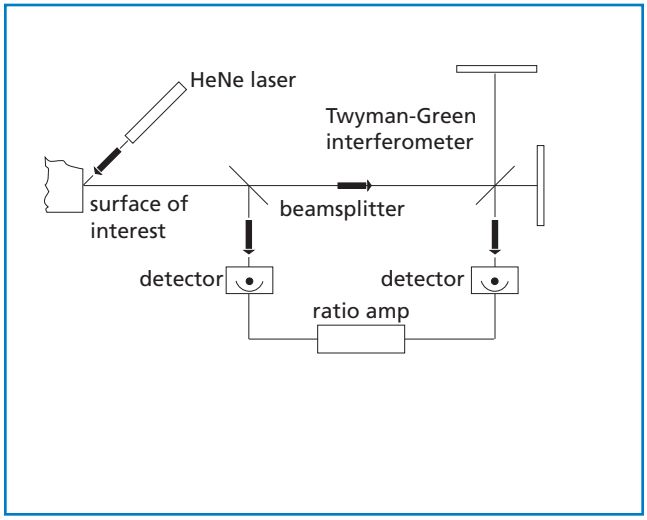

Figure 10.33 **Interferometric measurement**

SCIENTIFIC APPLICATIONS

Lasers are used extensively in the scientific laboratory for a wide variety of spectroscopic and analytic tasks. Two interesting examples are confocal scanning microscopy and time-resolved spectroscopy.

Time-resolved spectroscopy

Time-resolved spectroscopy is a technique used to observe phenomena that occur on a very short time scale. This technique has been used extensively to understand biological processes such a photosynthesis, which occur in picoseconds (10^{-12} seconds) or less. A fluorescing sample is excited by a laser whose pulse length is much shorter than the time duration of the effect being observed. Then, using conventional fluorescence spectroscopy measurement techniques, the time domain of the fluorescence decay process can be analyzed. Because of the speed of the processes, mode-locked lasers are used as the exciting source, often with pulse compression schemes, to generate pulses of the femtosecond (10^{-15} sec) time scale, very much faster than can be generated by electronic circuitry.

Confocal scanning microscopy

Scanning microscopy is used to build up a three-dimensional image of a biological sample. In standard light microscopy, a relatively large volume of the sample is illuminated, and the resultant light gathered by the objective lens comes not only from the plane in focus itself, but also from below and above the focal plane. This results in an image that contains not only the in-focus light, but also the haze or blur resulting from the light from the out-of-focus planes. The basic principle of confocal microscopy is to eliminate the out-of-focus light, thus producing a very accurate, sharp, and high-resolution image. A schematic of a confocal microscope is shown in figure 10.34. A visible laser is used as the light source to produce a distinct and spatially constrained point source of illumination. This light is then focused on the sample. A pinhole is placed in front of the detector at an optical distance that is exactly the same as the optical distance between the focus point and the illuminating source point (the confocal condition). Consequently, only the light generated at the illuminating point will, upon reflection or scattering from the sample, pass through the pinhole in front of the detector; out-of-focus light will be blocked by the pinhole. The signal from the detector is then digitized and passed to a computer. The complete image is digitally built up by scanning the sample in the x and y directions.

TIR and Fluorescence Correlation Spectroscopy

Fluorescence correlation spectroscopy measures the variation in fluorescence emission at the molecular level as fluorochromes travel through a defined field. The data can then used to determine binding and fusion constants for various molecular interactions. Because the measured volumes are so small, measurements are typically made using single-photon or two-photon confocal microscopy techniques. In many cases, the region of interest for fluorescence correlation spectroscopy is the first 100 to 200 nm of the sample's surface. However, the excitation depth (vertical resolution) for conventional confocal spectroscopy is 1 to 1.5 μm, leading to low signal-to-noise ratios and diminished accuracy.

One means of reducing the excitation volume is to use total internal reflection (TIR) techniques (see figure 10.35). If a laser beam, passing through a high index material (e.g., glass at $n \cong 1.5$) strikes an interface with a lower index sample material (e.g., an aqueous solution at $n \cong 1.3$) at an oblique angle, there is an angle of incidence (the *critical angle*) at which all of the light will be completely reflected at the interface, and none will pass into the lower-index material. The critical angle is given by

$$\theta_c = \arcsin\left(\frac{n_t}{n_i}\right) \qquad (10.26)$$

where n_t is the index of the transmitting (lower index) material and n_i of the incident material.

Because the beam is completely reflected at the interface, there is no energy flux across the interface; there is, however, an electromagnetic field generated in the lower index material, determined by the boundary conditions on the electric and magnetic fields at the interface. This transmitted wave is evanescent, propagating along the surface of the interface, but decaying in intensity exponentially with depth, limiting excitation to a few hundred nanometers—five to ten times better resolution than with confocal techniques alone.

Various techniques have been used to obtain TIR. Most commonly, the laser beam is brought in through a prism, as shown in figure 10.35. Another technique is to bring the beam in through the steeply curved edge of the observing microscope itself, and then filtering out the returning beam with a dichroic mirror.

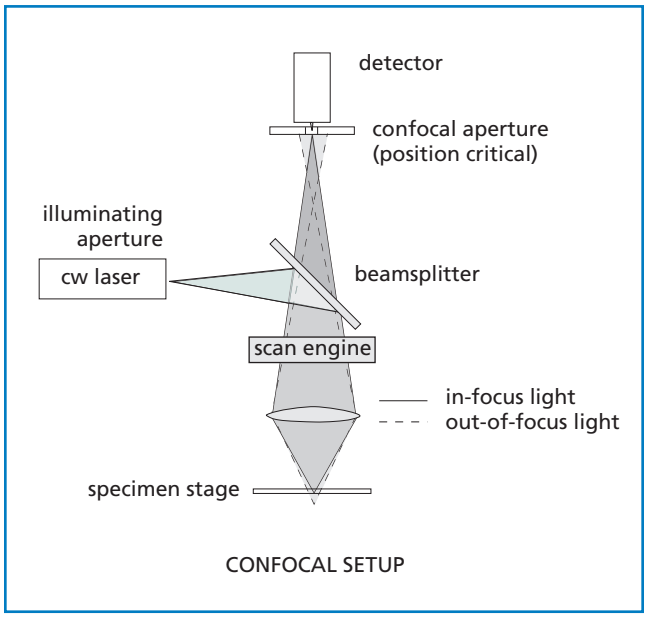

Figure 10.34 **Optical schematic of a confocal microscope**

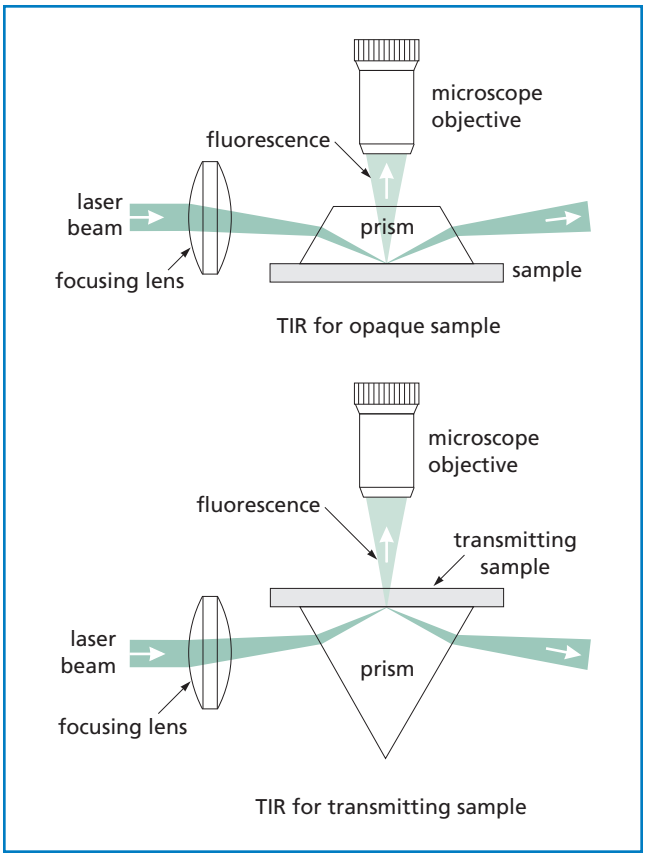

Figure 10.35 **An example of TIR spectroscopy**

Microarray scanning

In DNA research, a microarray is a matrix of individual DNA molecules attached, in ordered sets of known sequence, to a substrate which is approximately the size of a microscope slide. A single array can contain thousands of molecules each tagged with a specific fluorochrome. The array is then put into a microarray reader where each individual site of the matrix is individually probed by a variety of laser wavelengths at, or near, the excitation band of specific protein tags. The resulting fluorescence is measured and the fluorescence, position, and sequence data are stored in a computer database for later analysis.

Microarrays and microarray readers have had a dramatic impact on the speed by which data can be taken. Previously experiments were conducted one or two molecules at a time; preparation and setting up could take hours. With microarray readers, the raw data for analysis of thousands of molecules can be taken in minutes.

The main driver for microarrays is the pharmaceutical industry. If one can identify the differences in the way genes are expressed in a healthy organ and in a diseased organ, and then determine the genes and associated proteins that are part of the disease process, it may be possible to synthesize a drug that would interact with the proteins and cure or reduce the effect of the disease.

The optical system for a typical microarray scanner is shown in figure 10.36. The beam from a laser is focused onto a well (molecule) on the molecular matrix. If the appropriate fluorescent tag is present, the resulting fluorescence is measured by a detector. A filter in front of the detector separates the laser wavelength from the fluorescence signal. The laser beam is then moved to the next well.

Today's microarray scanner systems use two or more cw lasers, each with a different wavelength. Output power typically ranges from 10 to 50 mW, a power level that allows scanning without damaging or changing the material under test. Laser pointing stability is important as the microarray wells are quite small and repeatability is needed to relocate cells. Power stability and low noise are also extremely important due to the small sample size and the resulting weak fluorescence signal.

The most common lasers in use today for excitation are the blue solid-state (473–488 nm), green solid-state (532 nm) and red diode (650–690 nm) lasers. Solid-state and semiconductor laser technology is chosen primarily for its compact size, reliability, and power efficiency. Other wavelengths, including violet (405 nm) and ultraviolet (375 nm) from diode lasers, are currently being tested for application in microarray-reading applications.

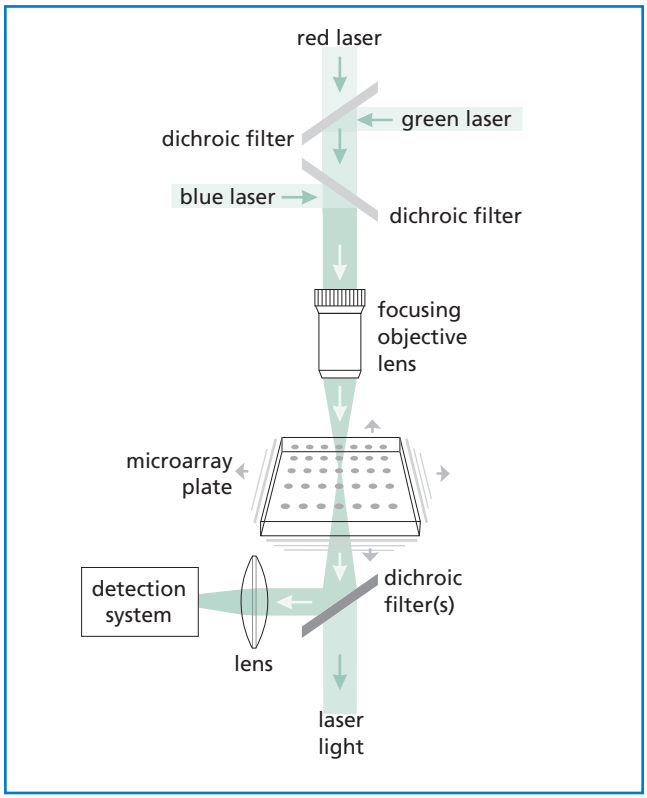

Figure 10.36 **Microarray scanning**

CLINICAL AND MEDICAL APPLICATIONS

One of the earliest applications of lasers in medicine was photocoagulation, using an argon-ion laser to seal off ruptured blood vessels on the retina of the eye. The laser beam passed through the lens and vitreous humor in the eye and focused on the retina, creating scar tissue that effectively sealed the rupture and staunched the bleeding. Today, lasers are used extensively in analytical instrumentation, ophthalmology, cellular sorting, and of course, to correct vision.

Many types of lasers are used in clinical applications including CO_2 , solid state, and diode lasers, as well as an array of gas lasers covering the spectrum from the ultraviolet to the infrared.

Flow cytometry

Flow cytometry is a technique used for measuring single cells. Not only is it a key research tool for cancer and immunoassay disease research, but it is also used in the food industry for monitoring natural beverage drinks for bacterial content or other disease-causing microbes.

In a basic cytometer, the cells flow, one at a time, through a capillary or flow cell where they are exposed to a focused beam of laser light (see figure 10.37). The cell then scatters the light energy onto a detector or array of detectors. The pattern and intensity of the scattered energy helps to determine the cell size, and shape. In many cases the cells are tagged with a variety of fluorochromes designed to selectively adhere to cells or cell components with specific characteristics. When exposed to the laser light, only those with the tag fluoresce. This is used in many systems to assist with separation or sorting of cells or cellular components.

The most popular lasers used in flow cytometry are the 488-nm (blue) argon-ion laser and the 632-nm (red) and 594-nm (yellow) HeNe lasers. However, new violet, blue and red diode lasers and a variety of new DPSS lasers are entering the field.

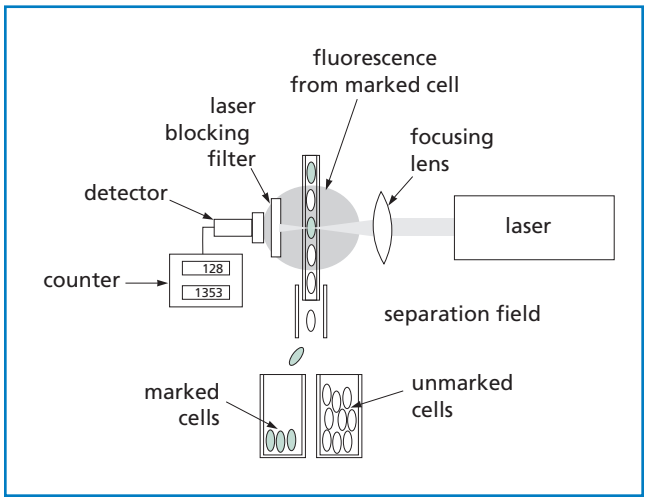

Figure 10.37 **Schematic of a laser cell sorter**

Surgical Applications

Lasers are used in a variety of surgical and dental procedures from cutting tissue, vaporizing tumors, removing tattoos, removing plaque, removing cavities, removing hair and follicles, resurfacing of skin and of course, correcting vision. In many ways, medical applications are like materials processing applications. In some cases material is ablated. In others tissue is cut or welded, and in yet others, photochemical changes are caused in blood vessels to encourage shrinkage and absorption. Understanding tissue absorption characteristics and reaction to wavelength and power are key.

Ultraviolet excimer lasers are used for vision correction because they can ablate material from the lens of the eye without causing thermal damage which could blur vision or make the lens opaque. Ruby lasers are used for tattoo removal because many of the dyes break down when exposed to 694-nm radiation, yet the skin tissue is left undamaged.

Cosmetic treatment of wrinkles, moles, warts, and discolorations (birth marks) is often accomplished with near infrared and infrared lasers. These procedures are often assisted by topical or injected photosensitive chemicals that assist with selective absorption at specific sites.

Lasers are also used to treat macular degeneration, an overgrowth of veins and scar tissue in the retinal region, a condition associated with advancing age. In this procedure, the patient is injected with a selective dye, which enhances the absorption of laser light by the blood in the blood vessels. When the blood vessels absorb laser energy, they wither in size, uncovering the active retina. A multiwatt green DPSS laser is most commonly used for this application because the green wavelength is not absorbed by the lens or aqueous portion of the eye, which allows the laser to affect only the targeted veins.