6. PULSED LASERS

Michael G. Littman and Xiao Wang

Department of Mechanical and Aerospace Engineering
Princeton University
Princeton, New Jersey

6.1 Introduction

Pulsed lasers are versatile tools for the scientist and engineer that have played an important role in the development of modern optical physics. The first laser, invented by T. H. Maiman, was a pulsed laser [1]. In this laser a ruby crystal (chromium-doped sapphire) served as the gain medium, which was excited by light from a pulsed flashlamp that surrounded the crystal. Two parallel faces of the crystal were coated with silver, which served to trap photons emitted along the optic axis so that they would pass through the crystal many times before escaping. The laser operated at room temperature and produced pulses of light with a wavelength of 694 nm (near infrared).

Pulsed lasers have proven to be especially helpful in the study of transient phenomena that occur in time intervals as short as femtoseconds, or as long as milliseconds or beyond. The shortest-duration laser created to date produced pulses of six-femtoseconds duration, a time that corresponds to only three optical cycles [2]. Because of their high peak power, pulsed lasers have also been useful in helping researchers study and apply effects that scale nonlinearly with increasing optical field strength. Pulsed lasers have been operated that have achieved enormous peak powers of as much as 10^{18-19} W/cm^2, which corresponds to a field strength strong enough to strip ground-state electrons from atoms.

In comparison with continuous-wave (CW) lasers, pulsed lasers generally are easier to operate, are more flexible, have higher peak powers, and cover a larger part of the electromagnetic spectrum. Pulsed lasers are also easier to build, because the high gain of the laser medium that is typical of pulsed lasers means that intracavity laser optics can be of relatively low optical quality and lossy. Thus, such inexpensive optics as aluminized mirrors, uncoated lenses, and polarizers are usable in the construction of pulsed lasers. For these reasons, many researchers choose to use pulsed lasers even when CW lasers might appear to be preferable. At present, pulsed lasers are widely used in areas ranging from atomic and molecular spectroscopy to remote sensing, from fiberoptic communications to laser medicine.

EXPERIMENTAL METHODS IN THE PHYSICAL SCIENCES
Vol. 29C

1079-4042/97 $25.00

In this chapter we provide an introduction to a wide variety of pulsed lasers. For an in-depth discussion of lasers in general, the reader is referred to A. Siegman's excellent text [3]. The chapter begins with a section that provides a brief overview of pulsed lasers and introduces some of the basic ideas that have led to the many different forms of pulsed lasers that now exist. This section can be viewed as a guided tour to pulsed lasers, and includes discussion of their operating principles, their attributes that might be useful in a variety of applications, as well as a brief survey of the types of lasers that are commercially available. The next section is a buyer's guide, where a range of issues are discussed that are relevant to the selection of specific lasers for specific applications. The last section gives a step-by-step guide to the building of pulsed tunable dye lasers of varying levels of performance. Pulsed tunable dye lasers are representative of most pulsed lasers, with the added feature that they are especially forgiving of imperfect optical elements and operating conditions. The robustness of pulsed tunable dye lasers derives from the very high gain (e.g., 1000 per pass) of the dye medium. For this reason they are ideal for exploration by the novice.

6.2 Pulsed Lasers

6.2.1 General Concepts

The LASER (**L**ight **A**mplified by the **S**timulated **E**mission of **R**adiation) is a quantum-mechanical device. There are three main elements in the typical laser: (1) a gain medium that has the capacity for light amplification along an optic axis, (2) an exciting energy source that raises the gain medium to a higher internal quantum state, and (3) a cavity that feeds back (i.e., resonates) light along the optic axis so as to efficiently extract energy stored in the gain medium. When the gain medium is excited by the energy source, states of higher energy become more occupied (i.e., populated) than states of lower energy, thereby creating a population inversion. (Population inversions are unusual in nature because quantum systems in thermal equilibrium follow a Boltzmann distribution, in which case lower-energy states are preferentially populated relative to higher-energy states.) When population inversion is maintained in the gain medium, light at or near a resonance frequency corresponding to a transition between the higher- and lower-energy states in the gain medium will be amplified when passing through the gain medium, by the process of stimulated emission. (If the medium were not inverted, light at or near the resonance frequency traveling in the medium would be attenuated by the complementary process of stimulated absorption.) If the cavity gain exceeds the cavity loss, the device will be self-sustaining and form what is generally referred to as a laser oscillator.

Gain media can be excited in numerous ways, including: (1) by electron collisions such as those that excite atoms in a gaseous discharge lamp, (2) by chemical reaction in which product molecules are formed in an excited state, or (3) by direct optical excitation due to absorption from an intense light source such as a flashlamp or laser. When the gain medium is optically excited, it is usually the case that the exciting light will be of higher energy (i.e., shorter wavelength) than the laser transition. So, for example, the flashlamps used in a ruby laser flood the crystal with light in the blue and near-ultraviolet portion of the spectrum, while the laser itself operates in the near-infrared.

There are two possible ways to achieve a pulsed output from a laser: (1) excite the gain medium for a brief but finite duration, or (2) allow for the circulation of light in the gain medium only for a finite time. The first method makes use of a pulsed excitation source. The second makes use of an intracavity optical switch. Depending upon how they are operated, optical switches can allow for either *Q*-switching, discussed later in this chapter, which provides for the rapid dumping of energy stored in the gain medium, or mode locking, which is capable of providing extremely short-duration optical pulses and is a subject treated in another chapter of this volume.

6.2.2 Ruby Laser

The first laser, built by Maiman in early 1960, used a ruby crystal as its gain medium [1]. Figure 1a shows the cavity structure of a ruby laser. Ruby, which is composed of Cr^{3+} ions embedded in sapphire host crystal, is a gain medium representative of the so-called "three-level system." A schematic energy diagram is shown in Fig. 1b. A time history of the laser output and the flashlamp output is shown in Fig. 1c. The electrons in the Cr^{3+} are first pumped by a pulsed flashlamp from ground states to a band of high-energy states. These states then relax rapidly through nonradiative processes into the first excited state of Cr^{3+} ion, which is long-lived. The laser transition, $\lambda = 694$ nm, occurs between the first excited electronic state and the ground state, and in the process the ground state becomes repopulated, thereby destroying the population inversion. Thus, the laser is pulsed, its width determined by the on time of the flashlamp. Since normally all Cr^{3+} ions are in the ground state at room temperature, in order to achieve the condition for population inversion it is necessary to excite more than 50% of the Cr^{3+} ions. Thus, ruby lasers require intense pumping, and as a result they are no longer widely used.

6.2.3 *Q*-Switching

Shortly after the invention of the ruby laser, the method of *Q*-switching was introduced by R. W. Hellwarth to create a high-intensity burst of light [4]. *Q* is the quality factor of a resonator, and it is a measure of its energy storage ability. For a *Q*-switching operation, a population inversion is built up during a pumping period that in the ideal case matches the excited state lifetime of the laser medium.

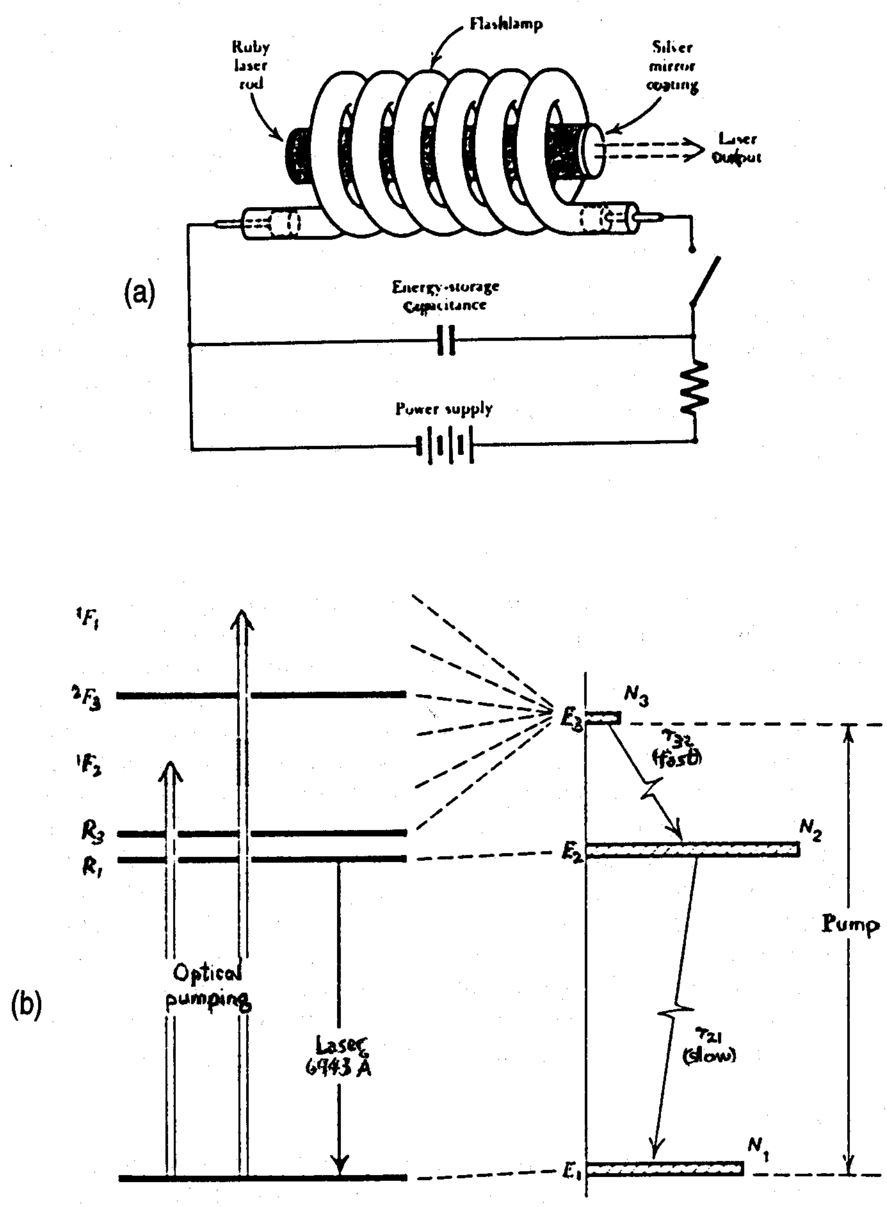

FIG. 1. The ruby laser: (a) physical configuration; (b) energy level diagrams (on right is schematic idealization; on left is real physical system); (c, opposite) time history of laser and flashlamp output. From Siegman [3]. Reprinted with permission from University Science Books.

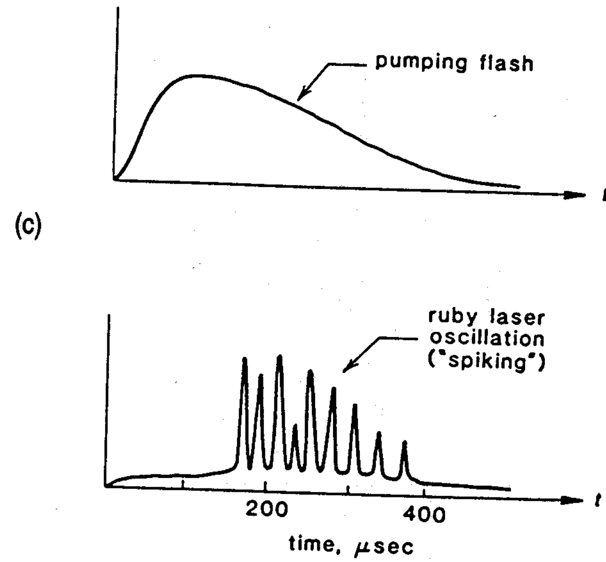

FIG. 1c

(For the ruby laser this time is approximately 3 msec.) During this pumping period the laser cavity Q is kept at a low value by keeping closed an intracavity optical switch in order to prevent circulation of light within the cavity. At the end of this pumping period, when the upper level is heavily populated, the cavity Q is suddenly switched from a low to a high value by opening the optical switch. Laser action then builds up rapidly, beginning with random noise light due to spontaneous emission. (Excited quantum systems decay spontaneously by the emission of light when there are no external driving fields. Spontaneous emission allows laser oscillators to be self-starting. If there was no spontaneous emission, it would be necessary to inject light into a laser cavity to get it oscillating. Spontaneous emission is distinct from stimulated absorption and stimulated emission in that these latter two processes require external driving fields.) Spontaneous light that falls in a narrow range of angles defining the optic axis of the gain medium is reflected by the mirrors and coherently amplified by the stimulated emission process. In a very short time, corresponding to a few passes through the gain medium, most of the energy stored in the medium is extracted. (The stored energy in the gain medium falls to zero typically in several nanoseconds.) For most solid-state laser media, this duration is much shorter than their natural (spontaneous emission) decay time. Q-switching is thus used to create a burst of light that is briefer and more intense than one would obtain without an optical switch.

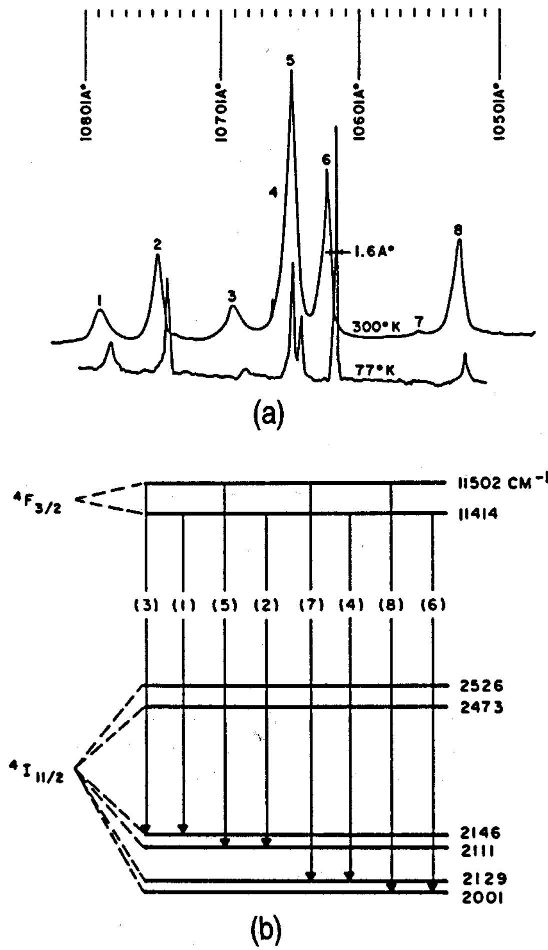

FIG. 2. Nd:YAG laser: (a,b) energy level spectrum and diagram; (c, opposite) energy level diagram idealization. From Geusic *et al.* [5]. Reprinted with permission from the American Institute of Physics.

If a laser medium has a long excited state lifetime, the population in the upper state can be built up gradually to a high level during that time, which can be accomplished using a modest pumping source. Therefore, Q-switching is easier to achieve in laser media with long excited state lifetimes, such as ruby. Other factors, such as the gain of the laser medium and the saturation fluence level of the medium, also influence the energy storage in the medium, and thus the efficiency of Q-switching,

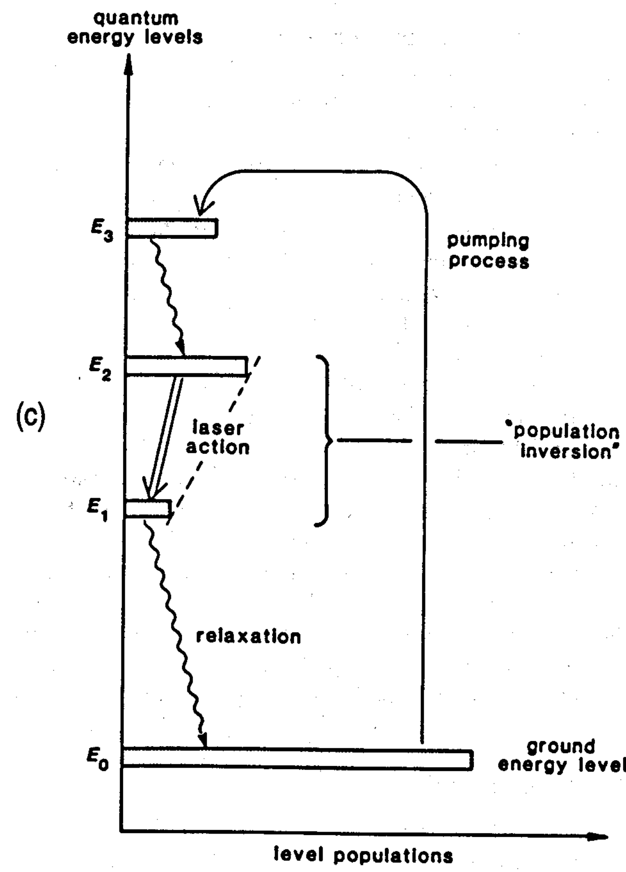

FIG. 2c

6.2.4 Nd:YAG Laser

Another laser in the solid-state laser family, the neodymium-doped yttrium aluminum garnet (Nd:YAG) laser, was first demonstrated by Geusic *et al.* in 1964 [5]. The Nd:YAG laser has been used extensively in a variety of applications, and especially for those requiring high peak powers. It is a four-level laser system (see Fig. 2), where the lower laser level is not the ground state. Therefore, in contrast to the Cr^{3+} ion that is used in the ruby laser, the Nd^{3+} ion is much easier to invert. The YAG host crystal is hard and has high thermal conductivity and good optical quality. The combination of easy pumping of the Nd^{3+} ion and a robust YAG host results in an excellent laser material that has high gain and low threshold for laser operation. Furthermore, the excited state lifetime of Nd:YAG Is 230 μsec, which can be easily matched by the duration of a typical xenon flashlamp; therefore, Q-switched operation can be easily realized in the Nd:YAG laser. Commercially available Q-switched Nd:YAG laser oscillator-amplifier systems can generate pulses that are several nanoseconds long, with 1–2 J of energy per pulse, and operate at repetition rates of up to tens of hertz. (Higher-frequency operation at these power levels leads to cracking of the laser crystal due to thermal stress.)

Although the output of the Nd:YAG laser is in the infrared, at a wavelength of 1.06 μm, by means of frequency doubling, tripling, or quadrupling using nonlinear crystals, laser pulses can be obtained in the visible (532 nm) and ultraviolet (355, 266 nm), with pulse energies of up to hundreds of millijoules achievable from "tabletop" light sources. Harmonic generation can be very efficient, especially at high power, where as much as 80% of light at the fundamental can be frequency shifted to the visible and near-ultraviolet. Crystals are not usable for wavelengths shorter than about 200 nm because of direct optical absorption by the crystal material itself. However, high-order harmonics can be generated by using nonlinear effects in rare gases, which do not absorb readily. These have been used to produce pulsed soft x-rays at wavelengths as short as 7.3 nm [6].

Another laser that uses the Nd^{3+} ion is the Nd:glass laser, where glass is used as the host material. Although glasses have excellent optical quality, they are usually less homogeneous and have lower thermal conductivity compared to the crystalline hosts. Glasses are used mainly because they can be made in large sizes; therefore, more energy can be stored and extracted. Nd:glass lasers are used to generate high-peak-power light pulses that are applicable to thermonuclear fusion.

6.2.5 Cavity Mode and Injection Seeding

In steady state an optical cavity can only resonate light of discrete wavelengths that satisfy the following condition: the cavity length is equal to an integral number of half-wavelengths of the circulating radiation. These wavelengths define the axial or longitudinal cavity modes. Most pulsed lasers operate in quasi-steady state. If the gain bandwidth of the laser medium is narrower than the spacing between the adjacent axial modes, single-mode laser operation can be realized. It is more often, however, that the laser transition linewidth is far larger than the spacing between the adjacent cavity modes. In this case, the transient laser cavity usually will support several modes simultaneously, unless an additional mode selection mechanism is used in the cavity that forces one mode to have higher gain than all the others. In a Q-switched Nd:YAG laser, for example, when the intracavity optical switch is first opened, the laser builds up from spontaneous emission, and many axial modes within the laser gain bandwidth will build up at the same time, so that the output pulse contains many axial modes. Mode-beating in the pulse results in an undesirable irregular temporal and spatial profile, and large near-random fluctuations from pulse to pulse. This unsteadiness is a typical drawback of multimode pulsed lasers.

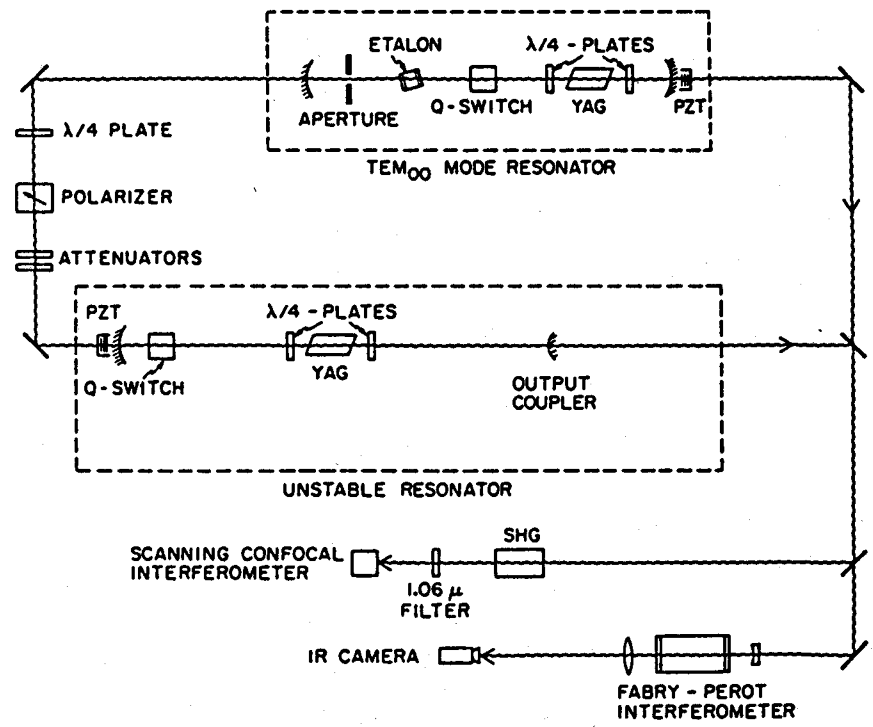

FIG. 3. Schematic of an injection-seeded Nd:YAG laser. From Park *et al.* [7]. Reprinted with permission from the IEEE.

Intracavity Fabry–Perot etalons can be used to reduce the number of modes, but they introduce large insertion loss and therefore reduce the output power. Injection seeding, applied by R. L. Byer [7] for mode selection in Q-switched Nd:YAG lasers, is a better method. Figure 3 is a schematic diagram for injection seeding of a Nd:YAG laser. For injection seeding, an external optical field with a bandwidth narrower than the mode spacing of the high-gain cavity and a field intensity much higher than that of the spontaneous emission noise field is fed into the high-gain laser cavity. If the frequency of the injected field matches an axial mode of the high-gain cavity, the injected photons will "seed" the cavity and be preferentially amplified. This will eventually deplete the gain in the medium, so that the other modes cannot be amplified, and only one mode will oscillate. The injection field is usually obtained from a single-mode CW laser operating at a frequency that falls within the gain bandwidth of the laser medium. Since there is no insertion loss in the cavity, injection seeding does not reduce the output energy of the laser. With injection seeding, Q-switched laser pulses with Fourier-transform-limited spectral linewidths have been obtained. Injection seeding also stabilizes laser operation significantly, reducing the time jitter and power variations from pulse to pulse. Injection-seeded Nd:YAG lasers are commercially available but are significantly more expensive.

6.2.6 Dye Lasers

Organic dye in solution also makes an excellent laser material. The first laser using this medium was a rhodamine 6G dye laser pumped by a frequency-doubled ruby laser, as demonstrated by P. Sorokin and J. Lankard in 1966 [8]. The upper state lifetimes of dyes are typically in the nanosecond range, so that a dye laser pumped by pulses nanoseconds or longer usually produces a pulse width equal to the duration of the pump pulse. A major advantage of dye lasers is their continuous tunability, which is a consequence of the large number of overlapping excited states in the heavy dye molecules. Tuning is achieved by using frequency selecting elements (for example, gratings) in the laser cavity. A grating can be used to feedback light selectively at the desired frequency for recirculation and amplification. By mechanical adjustment of the grating angle, the frequency for oscillation can be adjusted. Tunability is an important property that makes dye lasers popular light sources in laser spectroscopy. Figure 4 is a schematic diagram of a commonly used grating-tuned dye laser [14]. There are many different laser dyes available for operation in a broad range of optical wavelengths. Together, a collection of as few as 10 dyes can cover the spectral range from near-UV to infrared. Harmonics of a Q-switched Nd:YAG laser at $\lambda = 532$ nm and $\lambda = 355$ nm are especially well suited for the pumping of dye lasers to obtain frequency-tunable pulses in the nanosecond range. Short-duration gas discharge lasers such as XeCl excimer lasers at $\lambda = 306$ nm are also very useful in pumping dye lasers.

A disadvantage of dye lasers is that dyes can have short operational time due to bleaching of the dye molecules by the pump light source. Bleaching can require the replacement of dye solution after times as short as 20 minutes for the most fragile dyes. Dyes like rhodamine 6G are very robust, and they typically do not need replenishment for extended periods of time. Fragile dyes also have limited shelf life: sometimes less than one day after they are mixed. Another drawback is that dyes and solvents are sometimes environmentally hazardous.

6.2.7 High-Repetition-Rate Lasers

In some applications, high sampling rate, rather than high peak power, is needed. Therefore, a laser capable of providing a high repetition rate output pulse train is required. This can be realized by continuously pumping the laser medium and repeatedly Q-switching the laser at repetition rates up to several thousand pulses per second. For example, an acoustooptic Q switch can be used in an

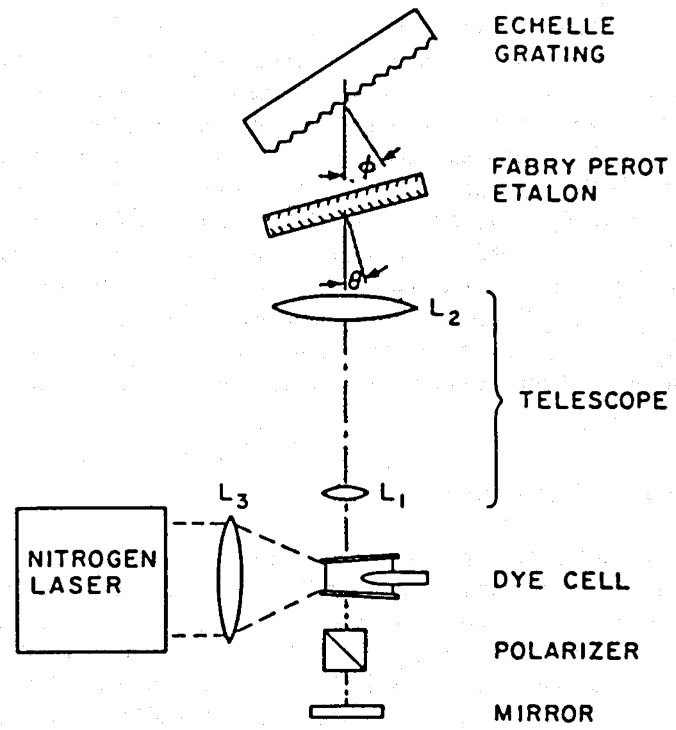

FIG. 4. Schematic of a narrow-band tunable dye laser. From Schawlow [14]. Reprinted with permission from John Wiley & Sons Inc.

Nd:YAG laser to provide a train of pulses at low energy per pulse and up to a 20 kHz repetition rate, with 10–20 watts average output power.

6.2.8 Gas Lasers

Another important class of pulsed lasers uses gas as the laser medium and a high current electron discharge as the pumping source. Examples are the nitrogen (N_2) laser, excimer lasers, and the CO_2 laser. Nitrogen lasers typically generate pulses that are several nanoseconds long, with up to several hundred millijoules of energy. The lasing wavelength is in the near-UV, at 337 nm, and it is suitable for pumping of dye lasers operating in the range of 400 to 800 nm. N_2 lasers were common pumping sources in the 1970s.

Excimer lasers, typically using mixtures of halogen and rare gases as the gain medium, are unique in that the laser media (excimer molecules) are bound only in the excited state. Since the molecule flies apart the moment it emits light, there are no ground state absorbers of light. As a consequence, population inversion is naturally accomplished as soon as the excited state is created. The laser wavelength depends on the kind of laser medium used, ranging from 193 nm (argon fluoride) to 351 nm (xenon fluoride). Typical pulse characteristics are: pulse duration 10–20 nanoseconds; pulse energy several hundred millijoules per pulse; pulse repetition rates up to 200 Hz. Because excimer lasers can conveniently produce high-intensity radiation in the UV region, they are widely used in industrial and medical applications. Like N_2 lasers, excimers are also used to pump dye lasers.

Excited CO_2 radiates at $\lambda = 10.6$ μm and is another important gas laser. CO_2 lasers have very high efficiency, can be operated with pulse lengths ranging from nanoseconds to and milliseconds, and have an average power up to several kilowatts. Aside from applications in scientific research, CO_2 lasers are widely used in industries where high power densities are frequently needed for material processing.

6.2.9 Tunable Solid-State Lasers

During the last decade, progress has been made in the search for tunable solid-state laser materials in an attempt to find better alternatives to tunable dye lasers. Today, the most important laser material in this class is titanium-doped sapphire crystal (Ti:Al_2O_3), introduced by P. Moulton [9]. The main advantage of Ti:sapphire crystal is its huge laser gain bandwidth, from 670 to 1100 nm. It also has a very large gain cross-section that is half that of Nd:YAG at the peak of its tuning range. As a host crystal, sapphire is extremely hard and has high thermal conductivity and excellent optical properties. The absorption band of Ti:sapphire centers at 500 nm; therefore, it can be pumped effectively by argon, copper vapor, and frequency-doubled Nd:YAG lasers, as well as rare gas flashlamps. Laser cavities developed for dyes are perfectly suitable for Ti:sapphire, with the dye cell replaced by Ti:sapphire crystal.

The upper laser level lifetime of Ti:sapphire is only 3.2 μsec. This presents some difficulties for traditional flashlamp pumping for *Q*-switch operation in a Ti:sapphire laser, because it is extremely difficult to make xenon flashlamps that operate with such short pulse duration. Therefore, in the flashlamp-pumped Ti:sapphire laser, the lamp duration is usually about 10 μsec. Although flashlamp pumped and *Q*-switched Ti:sapphire lasers are commercially available, they are not as efficient as a Nd:YAG laser. The usual configuration uses a *Q*-switched Nd:YAG laser to pump the Ti:sapphire laser. Using this pumping method and injection seeding, Fourier-transform-limited pulses that are several nanoseconds long and have several hundred millijoules of energy per pulse have been produced [10]. The Ti:sapphire laser can also be mode-locked when pumped by an intense CW laser such as an argon ion laser, and because of its large bandwidth it can be used to generate extremely short laser pulses in the tens of femtoseconds range.

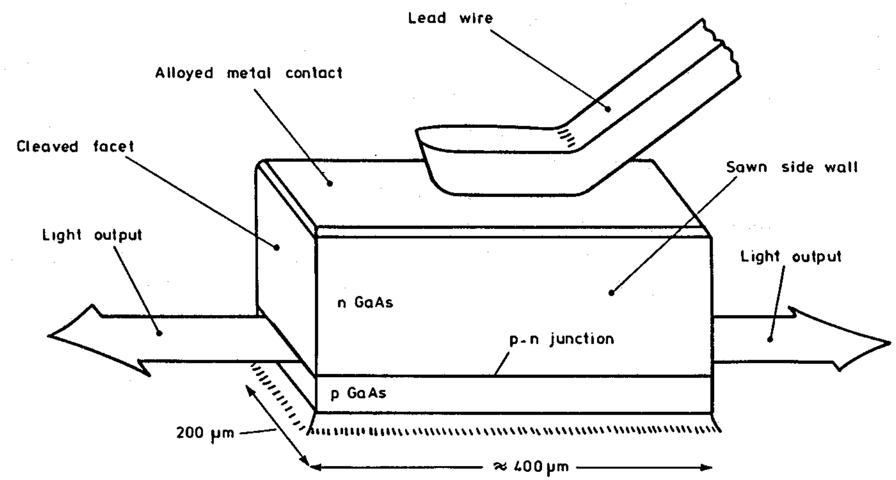

FIG. 5. Schematic of a GaAs semiconductor laser. From Thompson [11]. Reprinted with permission from John Wiley & Sons Limited.

6.2.10 Semiconductor Lasers

The semiconductor laser was invented by Basov *et al.* in 1961 [11], and is rapidly becoming one of the most widely used laser systems. A conducting *p–n* junction emits light; it also exhibits gain and can be converted into a laser if the emitted light is fed back to the junction. Figure 5 is a schematic diagram of a GaAs diode laser. Semiconductor lasers are significant in that they produce light with an efficiency that surpasses almost all other laser systems. Laser diodes are compact, relatively cheap, and have long operating lifetimes. They are commercially available at wavelengths ranging from around 800 nm (GaAlAs) to 10 μm (Pb-salt). GaInAsP laser diodes that emit light at 1 to 1.7 μm are particularly important for optical communication. Diodes that emit coherent blue light have been demonstrated recently [12]. Since the density of optical data storage is inversely proportional to the square of the wavelength of the light used, approximately four times higher data density can be realized if a blue laser diode is used instead of the near-infrared laser diode that is used in the current storage system. Semiconductor lasers are currently a subject of very active ongoing research.

Compared to other laser systems, diode lasers are relatively easy to operate. They only require a power supply with several volts, and typically less than 1 A of current output. (Although currents up to hundreds of amperes may be required for high-power diodes.) Pulsing of the laser diode output is extremely easy: it can be accomplished by directly modulating the driving current. The repetition rate can easily be adjusted from dc to the GHz range, although frequency chirping is usually associated with this kind of direct modulation, because the frequency of the laser also depends on the driving current. Laser diodes are available from CW single-frequency diodes that put out 150 mW of power, to bar stacks that produce up to 1.5 kilowatts peak power with 400-μsec pulses. A diode laser's frequency can be tuned by changing either the temperature or the diode current, or by using an external cavity that is similar to that used in tunable dye lasers.

Diode lasers have become popular pumping sources for solid-state lasers. Since the GaAlAs diode's wavelength can be tuned to match one of the absorption bands of the Nd:YAG crystal, GaAlAs diode can be used to pump an Nd:YAG laser. In the future, the bulky flashlamp-pumped solid-state lasers may well be replaced by compact diode-pumped setups. The laser wavelength of the GaAlAs laser also matches the lasing range of the Ti:sapphire crystal, so that it is possible to use a GaAlAs diode laser to injection seed the Ti:sapphire laser.

6.3 Buyer's Guide

In this section we explore issues frequently encountered when buying a pulsed laser and develop a checklist to aid the decision making process.

Each application places different requirements on the laser system. It is important to know, therefore, what a specific application demands, and to make sure the laser system that is purchased meets all those requirements. Of the many factors, the following are especially important: peak power, energy per pulse, repetition rate, pulse duration, wavelength, tunability and linewidth, transverse mode structure, pulse stability, and longitudinal mode structure. Let us examine them in more detail.

- *Peak power and energy per pulse*. Researchers studying such nonlinear effects as multiphoton ionization or pumping of atoms, harmonic generation in gases, laser-initiated plasma, soft x-ray generation, and relativistic behavior of atoms need lasers with high peak power. Researchers studying linear processes such as photon absorption where the detected signal scales up with energy per pulse, regardless of the peak laser power, do not require pulsed lasers. But pulsed lasers are often used for this purpose because of their simplicity, flexibility, and ease of use.

- *Repetition rate*. In laser-induced chemical reactions, to avoid creating nonlinear effects and unwanted reactions, laser intensity must be kept low so that the signal will be weak. In this case, one may want to sum up a large number of events to get a good signal-to-noise ratio. Therefore, a high-repetition-rate laser is preferred. In many cases the repetition rate of the laser is limited by power supply capacity. For some lasers such as the Nd:YAG, however, the repetition rate is limited by the laser material. At high energy per pulse (e.g., 1000 mJ), an Nd:YAG crystal may crack at a repetition rate in excess of 20 Hz.

- *Pulse duration* can be determined by the properties of the laser medium (e.g., lifetime), the duration of the pump source, or switching elements. To study transient effects, for example, to measure the lifetime of an excited state, one has to make sure the laser pulse duration is shorter than the duration of the phenomenon to be studied.

- *Wavelength*. Solid-state lasers such as Nd:YAG and semiconductor lasers operate in the infrared. To obtain visible or ultraviolet light from high-power solid-state lasers, one can use harmonic generation methods, although conversion efficiency is high only at power densities close to the damage limit of nonlinear crystals. Gas lasers such as N_2 and excimer lasers operate in the ultraviolet. Dye lasers and Ti:sapphire lasers are useful when light is needed in the near-infrared to the near-ultraviolet. The needs for specific wavelengths are quite varied. For example, in laser surgery, light is needed at a variety of wavelengths because biological tissue responds differently at different wavelengths. In molecular spectroscopy, light is needed at specific wavelengths to excite specific molecules.

- *Tunability and linewidth*. In applications such as laser spectroscopy, it is important that the laser be tunable in the wavelength range of interest. In high-precision measurements, the laser linewidth has to be narrow enough to be able to distinguish the features of interest. Dye lasers can be tuned over a large spectral range by adjusting intracavity frequency-selective elements and by changing the dye. The pump laser must be selected with attention to the needs of the gain medium, such as absorption bandwidth, excited state lifetime, and bleaching effects. The Ti:sapphire laser covers a large spectral range in the near-infrared. Its coverage can be further increased by frequency doubling or tripling using nonlinear crystals.

- *Transverse mode*. Many dye lasers and some solid-state lasers produce light in the lowest-order transverse mode, where the spatial intensity distribution in the beam cross-section is a Gaussian function. In some high-power solid-state lasers such as the Nd:YAG, the presence of higher-order transverse modes can lead to hot spots which can damage optics. If the application requires tight focusing of the beam, the fundamental transverse mode output is preferred.

- *Pulse stability*. This describes the variation from pulse to pulse of peak power, pulse energy, or pulse position in time relative to a marker pulse or external trigger pulse. Smaller variation means more reproducible events and therefore is preferred.

- *Longitudinal mode structure*. In lasers with either a single longitudinal mode or many longitudinal modes, the temporal profile of the output beam tends to be stable, and often can be approximated by a Gaussian function in time. In lasers with a few longitudinal modes oscillating simultaneously, beating between modes creates large variations in the temporal profile of the output pulses. This mode-beating can produce some unexpected results. For example, if the pulse-to-pulse variation in energy is 10%, the variation in energy of the second harmonic may be much greater than the 20% that one would expect if these pulses were smooth Gaussian profiles. For similar reasons, multimode lasers also require much higher damage threshold optics (typically by a factor of five) than one would guess by looking at the power specifications of typical lasers.

The next thing to consider is ease of laser operation and reliability. One wants a laser that is well-made, highly reliable, and requires minimum maintenance. Many lasers that are commercially available are operational nightmares. Be careful. Before buying a laser, one should also know the laboratory requirements that the installation will require. Here are some questions one should ask.

- *Is the laser well-designed and easy to operate?* Solid state lasers are usually easy to operate, but there is a large variation between different manufacturers. For example, some lasers are easy to turn on and off, while others may require complicated procedures. In some cases lasers may need an hour of warmup time before they are operational. Well-built lasers also have many clever devices built in to avoid accidental damage of the laser rod and optics by the user. It is well worth the extra cost to get those safeguards.
- *Installation.* Are there necessary facilities in the laboratory for laser operation? High average power lasers require special powerlines. They may also need cooling water to operate, or exhaust hoods to remove or deal with dangerous substances.
- *Environmental impact.* Dye laser operation involves dealing with organic solvents and dyes, some of them known to be carcinogenic. Therefore, safety precautions in handling the dyes and solvents and provisions for proper disposal of the waste material are necessary. Excimer lasers use gases that are highly toxic and sometimes flammable. Extreme care must be taken to prevent leaks in the circulation system, and the waste gas should be disposed of properly.
- *Safety features.* Are there shutters built into the laser? What about emergency shutoff? Are there high-quality laser safety goggles available for the laser?
- *Other equipment.* A stable laser table is a must. High-damage-threshold laser optics are needed for high-peak-power lasers. One also has to prepare the mounts, power meters, detectors, spectrometers, and other accessory equipment. Make certain that you can test the operating specifications in your laboratory. Frequent examination of the specifications can ensure proper laser operation, early damage detection, and savings in maintenance and repairing cost.

6.4 Builder's Guide

6.4.1 Introduction

This section describes how to design, assemble, and operate various forms of nanosecond-duration pulsed tunable dye lasers. We provide a step-by-step guide to building these dye lasers, from the simplest superradiant laser to a high-performance single-mode dye laser. Many of the basic concepts we introduced in the first section are discussed in detail here in a practical setting. It is our hope that by going through this exercise the reader can understand the issues that are representative of those that come up in the design of almost all lasers.

All of the dye lasers to be discussed in this section are pumped by direct optical excitation. This means that another light source (usually another laser) is used to create population inversion in the gain medium (i.e., an organic dye in solution) as required for laser action.

The pumping light source can be any of several different lasers, including a nitrogen laser (3371 Å), a frequency-doubled Nd:YAG laser (5320 Å), or an XeCl excimer laser (3080 Å). The only requirements are that the pumping laser must be of a shorter wavelength (i.e., higher frequency) than that of the dye laser, and that the gain medium is able to absorb light at the pumping laser wavelength. All of the examples presented in this section are obtained using a relatively inexpensive and low-powered nitrogen laser having an output pulse energy of about 1 millijoule in a 10-nanosecond pulse. The pump pulse is in the form of a rectangular-shaped beam of area equal to roughly 1 cm^2. This beam is focused using spherical and cylindrical lenses to a line on the face of a high-quality optical cell. (The cell and lenses must be made of a material that is transparent at the pump laser wavelength. For 3371 Å it is necessary to use a material such as fused silica (quartz)). This beam of light having a cross-section in the form of a line is absorbed in a fraction of a millimeter by the organic dye solution in the cell. Penetration depth and excited dye volume height are shown in Fig. 6. The peak power of the pump laser is about 1 mJ/10 nsec/1 cm^2 or 100 W/cm^2, and it operates at a pulse repetition rate up to 100 pulses per second.

Given that this section concentrates on the laser as a practical tool, most of the discussion is concerned with the hows and whys of various optical configurations and advice on how to select optical elements. The discussion deals with issues that must be addressed and understood if one is to obtain a useful instrument. The practical bent also means that there are forms of the instrument that are not ideal in the sense that laser designers who are purists would not likely consider them as desirable. However, if the goal is to obtain a light source that is matched to one's needs, then there is no reason to build a better laser than one has to. This is especially true with respect to aspects such as spectral width and beam quality. In fact, in many applications the lower-quality instrument may be preferable in that

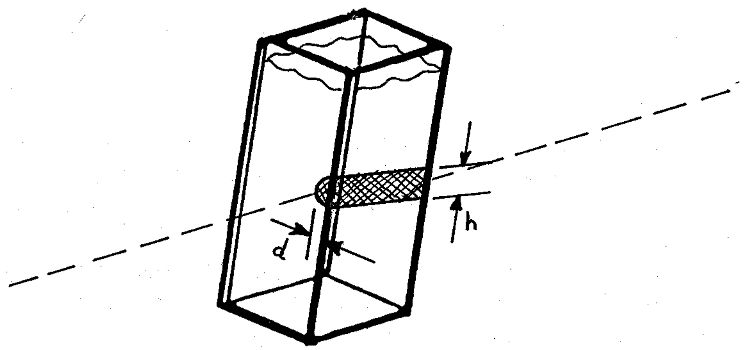

FIG. 6. Excited dye volume in a liquid-filled cell. Penetration depth and excited volume height are shown. (The cell is shown rotated 10° clockwise about the pumping laser propagation axis, and the long axis of the excited dye volume is rotated 3° clockwise about the pumping laser propagation axis. The cell is square in cross-section, and the central axis of the exit light from the excited volume lies in the horizontal plane.)

it might have other properties that are beneficial. For example, a lower-quality tunable laser might be more efficient, or have a broader tuning range, or have less amplitude jitter from pulse to pulse. Some of the lasers described in this section can be constructed out of parts that are extremely inexpensive and readily available in many laboratories.

6.4.2 A Laser Without Mirrors: Amplified Spontaneous Emission

It is surprising to many that a laser does not require mirrors, or a cavity, or feedback, in order to operate. A laser requires only that light be amplified by the process of stimulated emission. Amplification will occur in an optically active medium anytime that the population is inverted, meaning that there are more quantum particles (e.g., atoms, molecules, ions, electrons, electron–hole pairs) in an excited state than there are in a lower state. The system is "inverted" in the sense that a quantum system in thermal equilibrium would not exhibit such a feature. Inversions require effort to create, although they are known to exist in nature under special nonequilibrium conditions. (One frequently noted case involves the inversion of ammonia vapor in the atmosphere of Jupiter, which is excited by strong microwave emissions from the planet's surface.)

Evidence of laser action can be obtained from observing the spatial distribution of light emanating from a nonspherical volume of excited quantum particles. For example, if the excited volume has a long aspect ratio, let us say an elongated cylinder in the shape of a pencil, then spontaneous light emitted at one end of the pencil in the solid angle subtended by the remainder of the pencil-like volume will be amplified (see Fig. 6). In contrast, light emitted to the side will not be amplified to the same extent because the pathlength through the medium is shorter. Recall that the curve of growth in intensity is exponential as described by Beer's law. The

effect of all of this is that the light from the volume in total will be bunched along the long axis of the volume. As one turns up the population inversion, this bunching effect becomes more pronounced. Axial amplification of spontaneous light is known as *Amplified Spontaneous Emission*. It is equivalent to single-pass lasing, and the bunching or collimation of light is an indicator of laser action. The greater the bunching, the higher the gain of the medium. (If the medium is weakly interacting due to the low density of scatterers or low oscillator strength or equal populations in the upper and lower states, then the light emitted will be spherically symmetric far away from the volume. The light emitted will simply be the sum of the light from each independent voxel (i.e., volume element). This is the expected distribution from a dye medium that is not very strongly pumped. If the medium is strongly interacting but not inverted (i.e., absorbing), then a dark beam will be observed along the long axis of the pencil-shaped volume, that is, the light along the axis will be absorbed and preferentially scattered to the side.

To observe single-pass lasing in the laboratory is easy if one uses a highly stable and fluorescent dye such as rhodamine 6G. This dye is one of the most efficient UV converters known, which is also why it is used as an attention-getting coloration for highway reflectors. (This dye was also the first one used by Peter Sorokin, the IBM researcher who invented the dye laser. He selected it on its fluorescent properties as listed in a handbook of fluorescent materials. To this day it is the most commonly used dye in lasers.) One simply fills a 1 cm × 1 cm × 5 cm quartz spectrophotometer cuvette (i.e., dye cell; dimensions are not critical, although the flatness of the inner face and the sharpness of the inner frontal corners is critical; high-quality dye cells can be obtained from Nippon Scientific Glass or Helma) with a $1–5 \times 10^{-4}$ molar solution of rhodamine 6G in ethanol. Other solvents such as methanol or water are also usable as long as one is able to achieve an adequate concentration of dye so that the pump light is absorbed completely in a depth of less than 1 mm as it enters the cell.

The purity of the solvent is not important when working with rhodamine 6G, although it is important when working with other dyes that are at risk of breaking down following exposure to the UV light from the pump laser. The breaking down or "bleaching" effect of UV light is a problem with many dyes. This is the same effect that causes coloring agents in paints and dyes to fade in sunlight. For dyes that are prone to fade such as the Oxazine IR dyes, one must take care to use very pure (reagent grade) solvents and be prepared to replenish the dye solution periodically.

The pump light is focused into the cell using spherical and cylindrical lenses. The dye cell is slightly rotated (5 to 15°) from the vertical about the propagation axis of the pump light so that reflections from the faces do not contribute to the lasing action. This is needed because we are concentrating on single-pass lasing. In later forms of the laser we will continue to use this geometry. Unwanted reflections can interfere with the action of feedback elements as used in later forms

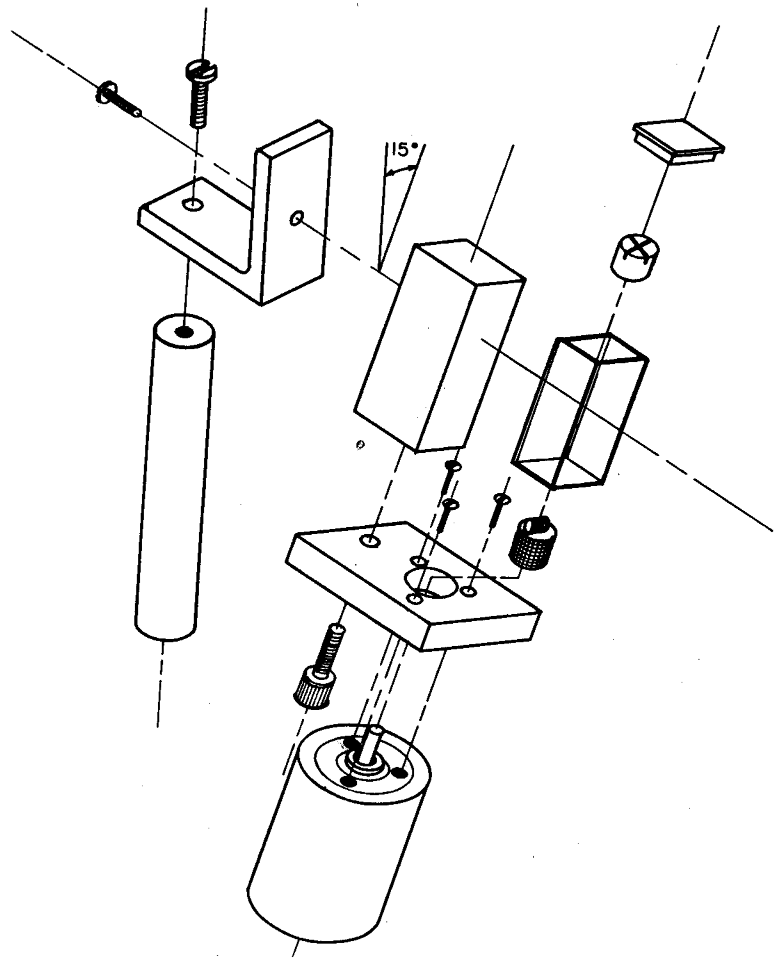

FIG. 7. Dye cell mounting at a 15° angle to the vertical.

of the laser and are to be avoided at all costs. It is useful to rotate the line axis of the focused beam in the same direction as the tilt of the dye cell but by an amount that is less by a factor roughly equal to $(1 - 1/n)$, where n is the index of refraction of the dye solution. The cell rotation and its mounting is shown in Fig. 7. The rotation of the axis of the focused beam allows for light coming from the long axis of the excited dye volume to be refracted in such a way that the ASE lies in the horizontal plane (i.e., the plane of the tabletop). Take the extra time that it takes to get the axis right. Optical misalignment out of the tabletop plane will cost you time in the long run.

An He–Ne alignment laser on the tabletop is a very useful tool in this regard. For example, to determine the correct angle to rotate the line focus of the pump laser in order to obtain table-planar operation, a planar He–Ne alignment beam can be sent into the cell. As it passes through the dye medium it will leave a streak that can be used as an indicator of the correct angle. After the laser is operating you can make additional angle corrections by adjusting the cylindrical lens rotation angle.

The dye cell is stirred with the help of a magnetically coupled stirrer in the cell. In stable dyes such as rhodamine, stirring is necessary for pulse repetition rates in excess of 10 Hz. In dyes that are bleached by the UV pump light stirring is always needed. Since most of the pump energy is converted to heat following the pulse (even the best dye lasers have an efficiency of at most 30 or 40%), it is necessary to gently swirl the dye or to flow it. This brings in fresh dye for each shot and makes the excitation volume more optically uniform. If one does not do this, the beam may appear distorted. This is the same effect that causes optical distortion if one observes an object along a path in air over a hot source such as a radiator or a stove.

A word about safety. These types of dye laser do not appear to be very intense until one reaches pulse energies on the order of 10–100 mJ per pulse. Because of this, many researchers get sloppy and allow primary beams and reflections to bounce around the laser laboratory. This is to be avoided at all costs. The average power of these lasers is indeed low, but peak power can be very high. Eye damage is the result of a mixture of peak power and average power. One shot directly in the eye from a 100-μJ dye laser is enough to burn a hole in your retina. A 100-μJ pulse in 1 cm^2 when observed on a white surface looks relatively dim, even though its peak power is 10 kW/cm^2. One reason for this perceptual error is that the eye integrates light over roughly a 40-msec period (a 25-Hz scan rate) in conditions of normal room illumination. These nanosecond dye laser pulses are much shorter in duration than the eye's integration period, so that when we compare them to other light sources that are on for the entire integration period the laser light can look unimpressive. Do not be fooled. The same warning is relevant when one works with lasers outside our normal spectral response. An 8000- or 4000-Å laser may not look very bright since our retinal photosensitivity is low at these wavelengths; however, if one gets this light in the eye it can cause a great deal of damage. Eye damage due to UV may not show up right away. We recall an incident at MIT where a fellow using a quartz microscope and a mercury vapor light source (a UV source) burned his retina. He did not know it until several hours after exposure, by which time he was in a great deal of pain.

Laser goggles are recommended when working with high-powered lasers. For the UV, inexpensive UV goggles may be obtained from a variety of sources (e.g., the Cole–Palmer Instrument Company). Check the absorption factor when obtaining laser goggles. Goggles are rated by their absorption. For example, a goggle might absorb 99.9% of the incident light in the UV, which corresponds to an optical density (OD) of 3.0. Goggles for the dye lasers are more problematic because dye lasers operate at wavelengths over the entire visible spectrum. It may be necessary to have many sets of goggles for different ranges. Goggles with very high optical density can be very expensive, often in excess of a thousand dollars each. Also, do not get goggles that make the laboratory appear so dim that it increases your chances of having another type of accident. This is especially important if one works near rotating machinery or around high-voltage instruments. In fact, the single greatest hazard in a laser lab is the high-voltage power supply used to run most lasers. Laser power supplies are lethal.

Be reasonably careful not to let the UV light hit your skin. UV in high intensities can cause burns and is carcinogenic. The 3371-Å light in the nitrogen laser is not very dangerous in terms of skin damage, however. Shorter-wavelength UV light can be very dangerous. Particularly the fourth harmonic of the Nd:YAG at 2660 Å, which is sometimes used in pumping UV dyes, is in the dangerous wavelength band for producing skin cancer.

The procedure for obtaining single-pass lasing is straightforward. Turn on the pump laser to low power. Make sure that the magnetic stirrer is operating. The UV beam can be made visible by letting it hit a white business card. Business cards appear "whiter than white" because of an organic dye (stilbene) that they are soaked in. This is the same whitener that is used in laundry detergent to make laundered clothes appear bright. (The clothes are actually being dyed.) These whitening agents absorb UV light and reradiate in the visible range, where we can see it. Focus the light to a line, as discussed previously. Observe the dye light emitted from the excitation volume on each side of the dye cell. Do this by using white cards placed about 10 cm away from the dye cell at right angles to the optic axis of the laser and on either side of the cell. Notice that the UV pump light is absorbed as soon as it enters the dye medium. Also look at the lenses and the cell face to see if you see a bluish-white streak where the UV has passed through the transparent material. If you have selected, in error, a cell or lenses made of pyrex or some other type of glass, then the UV beam will be visible as it goes through the material. Quartz or fused silica will not have this bluish-white appearance.

Turn up the UV pump laser power. This can be accomplished, for example, by turning up the operating voltage on the pump laser. Watch the beams on the white cards as the power is increased. You will notice that the light emitted along the optic axis of the dye laser will become progressively brighter and more bunched along the long axis of the cell. This is evidence of the single-pass lasing that was discussed earlier. You will likely also observe another bright spot due to a reflection off the cell wall. This spot will appear more well-defined and might even be very bright. It is the result of two passes through the gain medium. The gain in the case of single-pass lasing can be very high. Given that a 100-μJ pulse of rhodamine 6G light (energy of about 3 eV per photon) contains about 2×10^{14} photons, depending on the aspect ratio the gains may be as large as 10^{10} or higher. (In order to calculate the exact gain, it is necessary to consider the factors that determine the probability that a spontaneous photon win be emitted and will fall in the solid angle defined by the geometry. When the gains are so high, the likelihood that many spontaneous photons contribute to the final emission is high, and as a result the level of coherence of this light source is low.)

To convince yourself that the gain is high, try turning down the pump laser so that the dye laser output is only weekly bunched along the axis, and then bring another business card close to one side of the dye cell so that some of the reflected light is bounced back in the direction of the cell. You should notice that on the screen on the other side of the cell a spot due to this small amount of reflected light will be visible. The laser is able to amplify any light that passes through it during the time that the population is inverted, even if it comes from a diffuse source.

Under some circumstances vertical diffraction "stripes" may be observable in the dye laser beams on each of the screens. This will be especially pronounced if the dye concentration is high. A high dye concentration means that the penetration into the cell is minimal. The shorter the penetration, the larger the diffraction angle from the confined emitting volume. Think of this as the diffraction from a slit where the diffraction angle in radians is roughly the wavelength divided by the slit width. Given that the screen is 10 cm away from the cell and that the wavelength of the dye laser is about 0.5 μm, a penetration depth of 50 μm will result in stripes of separation equal to 10 cm × 10^{-2} radian or 1 mm. Thus, if one observes multiple stripes separated by more than 1 mm it means that the penetration depth is very small and that the dye concentration is probably too high.

This discussion about the penetration depth also suggests the need for a high-quality dye cell. Since all of the "action" in the dye laser is happening in a very narrow region just behind the front face of the dye cell, it should be obvious that the surface quality of the inner face of the dye cell needs to be high (at least $\lambda/4$ or better) and that the inner corners of the dye cell need to be sharp and without curvature.

A few comments about the laser wavelength are in order. If the output of the dye laser is analyzed as a function of the pump power, it will be noted that the spread in wavelength will reduce dramatically as the pump power is increased. This occurs as the result of gain narrowing in which the central wavelength of the band for which the gain is high will be preferentially amplified over wavelengths that are in the "wings" of the gain profile. This is a consequence of the exponential curve of growth for intensity and the fact that the gain factor, which is in the exponent, is frequency-dependent.

Many dyes other than rhodamine 6G can be used to obtain light from this single-pass laser. In fact, the first tunable dye laser was tuned by changing dyes and solvents and by modifying dye concentrations. (Peak wavelengths of a given dye can be affected by choice of solvent, and concentration.) Rhodamine dyes operate in the red–yellow–green range, Oxazine dyes work in the IR–red range. Coumarin dyes operate in the blue–near-UV range. DCM operates in the blue–green range. Other specialty dyes are available for use in certain narrow bands. A good source of laser dyes and information is Eastman Kodak (see Fig. 8). Sometimes it is possible to mix dyes to get coverage in a hard-to-reach band. This was the case in the 7500-Å range, but this problem range has been covered by the emergence of a new laser material, titanium:sapphire, which uses a titanium ion embedded in a sapphire host crystal. This outstanding tunable solid-state material has virtually eliminated the use of the dye lasers in the spectral region to the red of 7500 Å.

In general, one wishes to select dyes that absorb strongly at the wavelength of the pump laser. Stated another way, one should select a pump laser wavelength that best matches the absorption properties of the dye that one wishes to use. Organic dyes have a broad spectral range of usefulness because they are quite heavy (molecular weights of 500 are typical) and as a consequence have many modes of oscillation. This high degree of freedom is what allows them to be used as a tunable medium. Lighter and simpler molecules have a regular quantum line and band structure that has many "holes" in the absorption/emission spectrum. Atoms are even worse, in that they have a spectrum that is typically sparse. Thus, a helium–neon laser can oscillate at only certain wavelengths that correspond to optical transitions in neon. A carbon dioxide molecular gas laser similarly will oscillate at only certain wavelengths corresponding to transitions in the carbon dioxide molecule. Of course, since it is a molecule, there are a large number of nearby lines corresponding to quantized rotations of the molecule. Heavy molecules like organic dyes have so very many modes of oscillation and vibration that their spectrum is smooth over a wide range, and it is for this reason that they are useful as a tunable gain medium.

The strongest "spring" in the organic dye molecule is the carbon–carbon double bond. One vibration of this bond is about 1200 cm^{-1} of energy. The decaying electronically excited dye molecule will typically drop down to a lower electronic state in the lowest unoccupied vibrational state (usually the first excited vibrational state). This vibrational state is at an energy that is 1200 cm^{-1} below that corresponding to the peak absorption energy. Thus, rhodamine 6G, which has a peak emission at 6000 Å (16,666 cm^{-1}) has a peak absorption at 5500 Å (18,000 cm^{-1}). Thus, we see that, while it is possible to use UV to pump this dye, it would be far more efficient to use light at 5500 Å since this corresponds to the peak of dye absorption. UV pump lasers are common in many laser labs because most of the fluorescent dyes will absorb the light. A single pump source such as a nitrogen laser or an XeCl excimer laser can be used to cover the entire visible spectrum merely by changing dyes. Atomic and molecular spectroscopists especially take advantage of this feature. Sometimes the bleaching effects of a UV pump make it

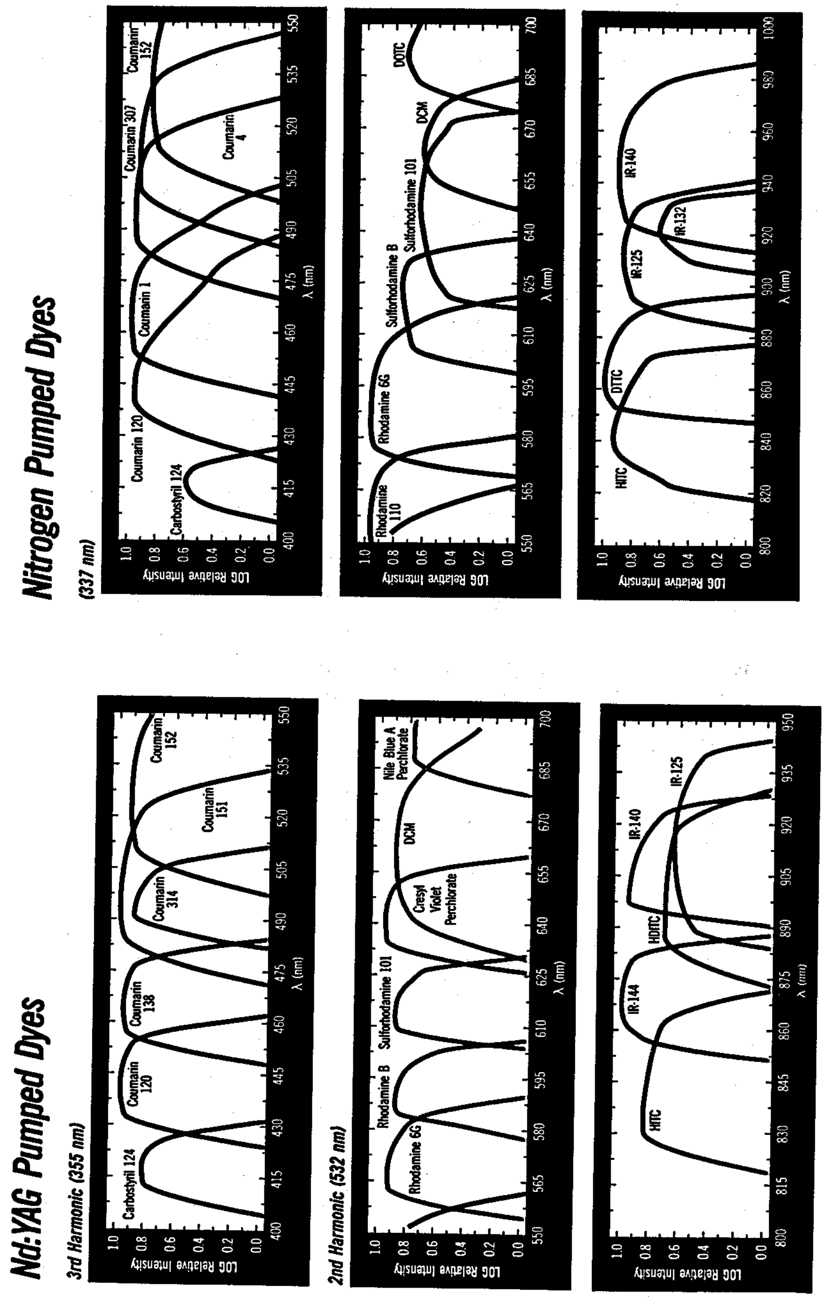

FIG. 8. Dyes for use in dye lasers: nitrogen and Nd:YAG laser pumping sources. From Eastman Kodak [8]. Reprinted with permission from the Eastman Kodak Company.

necessary to use the light from one dye laser to pump another. This is especially useful with the fragile IR dyes that are easily bleached by the UV pump light.

An excited dye cell is usable as a light amplifier. It can amplify light from a pulsed laser oscillator or from a CW laser. The acronym MOPA refers to "master oscillator/pulsed amplifier." MOPAs are used routinely in many laboratories where many millijoules of tunable light are required. The usual configuration uses one or two pulsed amplifiers of the side-pumped geometry. In many cases, the last stage of amplification is a longitudinal amplifier cell. For pumping using a laser with a rectangular output beam such as many excimer lasers, it may be convenient to incorporate a Bethune cell in the last stage of amplification. The Bethune cell is a cleverly bored-out prism created by Don Bethune, formerly of the IBM Research Laboratory. Dye flows through the cylindrical bore. (Bethune cells are available commercially from the Santa Ana Laser Company.) Light from a pump source having a rectangular beam shape, typical of high-powered excimer lasers, is internally reflected radially to excite the dye from several sides. The dye concentration is lower in this stage, so that the pump light can penetrate further into the dye volume. This kind of amplifier is needed because, as the peak power of the dye laser light increases, damage to optical elements becomes more of a problem. Longitudinal amplifiers are used to obtain hundreds of millijoules of tunable dye light. The concept of increasing the size of the pulse amplifier as the intensity builds in each stage is routine in the laser field. The rule of thumb is to keep the dimensions as small as possible given the constraints imposed by the strength (in this case, the optical damage power threshold) of materials.

Side-pumped pulsed amplifiers are also useful for increasing the power of CW lasers. One configuration that is sometimes used involves a tunable CW dye laser that is pulse-amplified by a series of several side-pumped pulsed amplifiers. The gain of a chain of amplifiers can be very high. Unfortunately, this type of tunable oscillator–pulsed amplifier system is very expensive, large, and difficult to maintain.

6.4.3 The Advantage of Mirrors: Multiple Passes Through the Gain Medium

It is usually the case that one wishes to have a laser light source that has many desirable properties. Properties such as high power, good beam shape, low beam divergence, narrow spectral width, tunability, and spatial/temporal coherence are often required. In the superradiant laser discussed in the previous section, many of these properties are absent.

The use of mirrors to recirculate the amplified light for further amplification is extremely useful in helping to obtain a laser that has many of the beneficial properties mentioned above. Recirculation may be accomplished in a standing-wave cavity, as shown in Fig. 9, or in a ring cavity [3]. Historically, the standing-

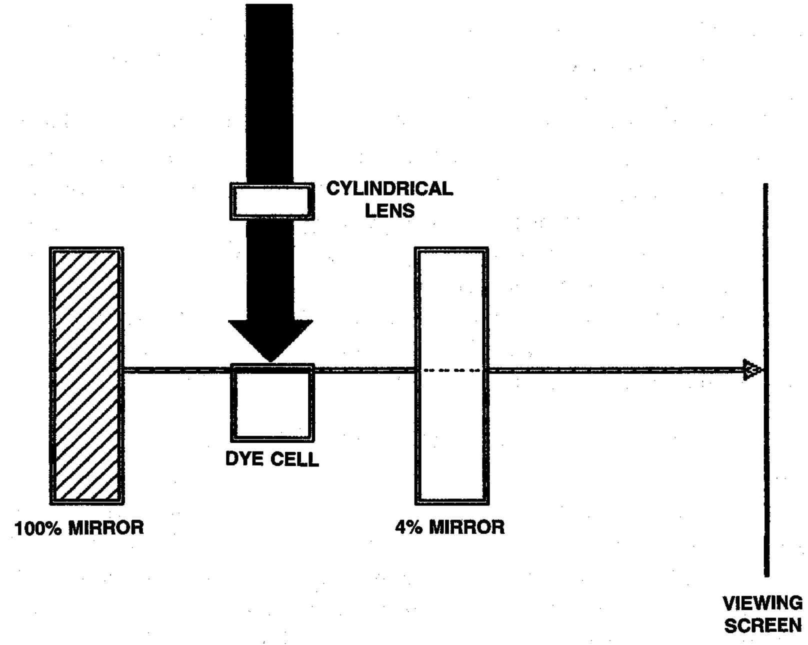

FIG. 9. Standing-wave laser cavity.

wave cavity has received the greatest use, and it is the type of cavity that will be considered here.

Let us consider the excited dye cell of the superradiant laser in the previous section when it is placed between two mirrors, one fully reflecting and one partially reflecting and separated by 30 cm. Since these two mirrors are assumed to be aligned perfectly, the gain medium plus reflectors may be replaced conceptually by a series of dye gain cells. Light travels at a speed of roughly 30 cm/nsec, so that the 10-nsec duration of the excited medium translates into a series of 10 gain cells of 1 cm pathlength and separated by 30 cm. In essence we have formed a chain of dye gain cells that is 300 cm long. The gain of each cell does not have to be anywhere as high as it was in the superradiant laser of the previous section in order to extract light efficiently, as long as the exponential curve of growth of the dye light is given the opportunity to build to the same value as before.

The beam divergence of this multiple-pass laser is seen immediately to be much less than the single-pass laser. The single-pass laser beam divergence was set by geometry, as is the multiple-pass version, except that the multiple-pass version has a much smaller aspect ratio and therefore a much smaller divergence.

Since each stage of the amplifier chain is pumped less than the laser in the previous version and since the solid angle is smaller, the number of spontaneous photons that initiate the laser to oscillate is greatly reduced. In the previous case, many "seed" photons were amplified, so that the output was in the form of a collection of many amplified incoherent photons. In this case, we have fewer "seed" photons that are amplified to a greater extent, with the result that the coherence of the multiple-pass laser oscillator is far better than that of the single-pass laser. Note that, since the laser is still operating on many wavelengths, coherence is not high. Coherence of macroscopic measure (i.e., lengths of centimeters or times of nanoseconds) is only possible if we additionally add narrow spectral filters to the laser cavities. By coherence we mean that the phase of the electromagnetic wave at another point in space or time is known given that we know it at a point we use as an origin. Coherence is a necessary ingredient of interference. Interference effects will be considered later in this section.

6.4.4 Intracavity Spectral Filters for Narrow-Band Operation and Tuning

When dye lasers were first invented, it was not appreciated that the laser frequency could be tuned with the help of intracavity optical filters. At first, the suggested method for tuning was to change dyes or the dye chemical environment (e.g., pH, solvent, temperature). In 1971, Strome and Webb discovered that the use of a spectrally dispersing intracavity prism would allow for a reduction in laser bandwidth [13]. It was also discovered that the operating wavelength could be changed by adjusting the angle of the dispersing prism. An aperture in the cavity also helped to reduce the linewidth. In essence, the laser became an *active* prism spectrometer. The years that followed this discovery led to many cavity configurations that included an enormous variety of intracavity tunable optical filters whose purpose was twofold: to narrow-band the regenerative oscillator and to tune it. Among many intracavity elements, researchers used Lyot (birefringence) filters, prisms, gratings, and interferometers. In this chapter, we consider only the use of reflecting diffraction gratings and interferometers for narrow-banding and tuning.

Diffraction gratings are desirable as narrow-band filters because they are highly dispersive. A grating, if used properly, has a $d\lambda/d\theta$ that may be an order of magnitude larger than that of a typical prism. Also, since they are reflective, it is possible to use gratings far into the UV. Their use, however, does introduce some difficulties. One problem is that diffraction gratings are easily damaged by high-powered laser beams. Another is that they are quite lossy in comparison to other filters, such as birefringent filters and prisms. For these reasons, gratings are not used often in lasers of low gain (gain per pass must exceed loss per pass, if the laser is to oscillate, i.e., regenerate), and they are not used in oscillators in which the intracavity power is high. What defines high power, however, is not easily stated because it depends on the geometrical configuration in which the grating is used in the laser.

The diffraction grating obeys the grating following equation:

$$m\lambda = d(\sin\theta_i + \sin\theta_d),$$

where d is the grating period, θ_i and θ_d are the incident and diffraction angles as measured with respect to the grating normal, m is the order of diffraction, and λ is the wavelength. In the special case where $\theta_d = -\theta_i$, the diffracted light is retroreflected. At this wavelength, the grating acts like a plane mirror at right angles to the optic axis of the laser. For other wavelengths, the light is not retroreflected. If one places a small aperture or slit in the cavity, the off-axis light will not reach the gain medium and therefore will not be amplified. This configuration is known as the Littrow arrangement. For the side-pumped geometry that has been discussed so far, the narrowly defined gain medium in the form of a thin filament along the front face of the dye cell serves the purpose of both gain medium and slit. Thus, if one replaces the highly reflecting mirror in the laser oscillator of the previous section with a retroreflecting and dispersing grating, the laser will oscillate at the retroreflected wavelength. Tuning of this laser is then readily accomplished by rotating the grating relative to the surface normal, as shown in Fig. 10. The output from the laser can be either the exit beam off of the grating (the zeroeth-order reflection) or the beam exiting from the partial reflector. The beam through the partial reflector is usually used since its direction does not change as the grating is rotated. It should also be obvious that, since there are two exit beams, the design is wasteful of light.

This laser is impressive in its ease of construction and its efficiency (except for the extra exit beam that was just mentioned). It is not especially impressive with respect to its linewidth, however. The linewidth of this laser is typically larger than 0.2 nm. For comparison, the Doppler-broadened atomic absorption linewidth of sodium vapor at low pressures is 0.002 nm. The reason for the large bandwidth of the laser is that only a very small portion of the grating is being used for dispersion. The linewidth can be expressed as:

$$\delta\lambda \approx \lambda^2/L,$$

where λ is the laser wavelength, and L is the linear dimension of the portion of the grating that is being used for dispersion. The laser spot size is typically 1 mm in this kind of set up, which makes the linewidth about 0.2 nm for an operating wavelength of 500 nm.

It should be obvious from this discussion that, if one could increase the intracavity beam size so that the entire surface of a large (>5 cm) grating could be used, it might be possible to achieve a very narrow band source. In fact, this was accomplished by T. Hänsch in 1972 [14] by means of an intracavity beam-expanding telescope. This arrangement significantly narrows the bandwidth. However,

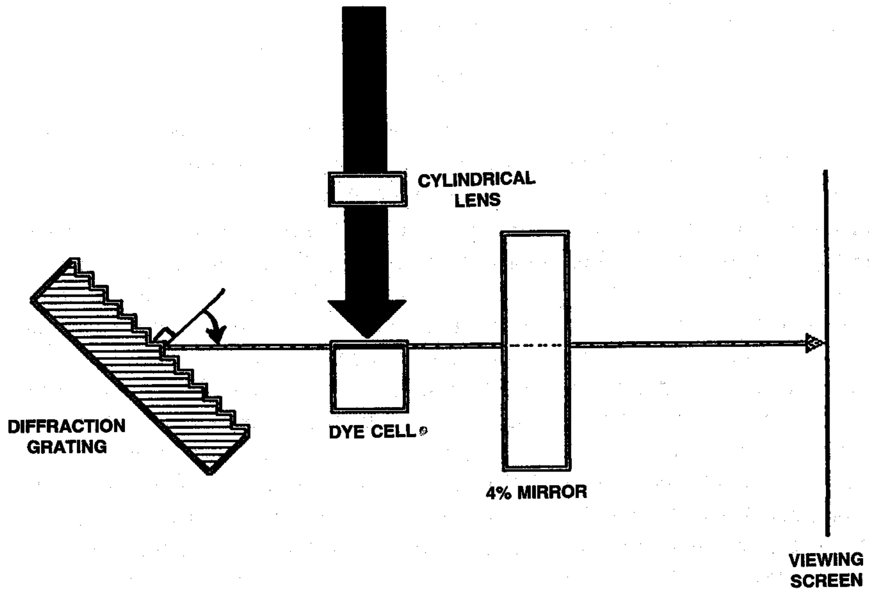

FIG. 10. Top view of a grating tuned dye laser.

using refractive optics in a laser cavity has some drawbacks: intracavity optics give rise to surface reflections that can compete with the feedback from the tunable filter; an intracavity optic has to be of high quality in order for it not to distort the spatial wavefront. High-quality optics are expensive. Refractive materials are also dispersive, which means that the beam expander can only be collimated at one wavelength. Other wavelengths will experience either focusing or defocusing, which has the result of broadening the laser linewidth, and it leads to a linewidth that changes as the laser is scanned in wavelength. Lastly, the requirement of exact alignment demands high-quality mounts and stages.

All of these problems are averted if the grating is used in grazing incidence, as shown in Fig. 11. In this case, the grating is held fixed and the tuning mirror is rotated about the grating. This configuration is known as the grazing incidence dye laser and it was developed in 1976–77 independently by Littman and Metcalf and Shoshan, Danon, and Oppenheim [15].

In the grazing-incidence configuration, the grating angle is steep enough that the intracavity optic beam illuminates the entire width of the grating. The grating diffracts the beam at an angle given by the grating equation. The diffracted light comes off as a streak. A mirror is used to retroreflect the light back to the grating, where it is rediffracted back to the dye cell. Only the wavelength that corresponds

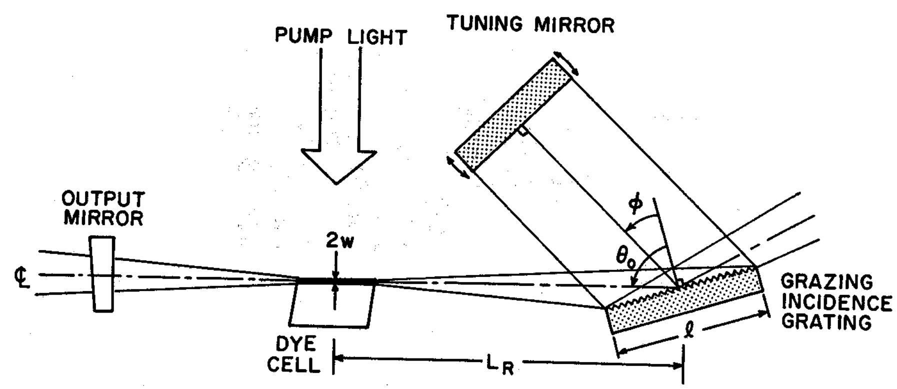

FIG. 11. Top view of a side-pumped grazing-incidence dye laser.

to the value given by the grating equation for specified angles θ_i and θ_d is reflected back to the dye cell for amplification. Other wavelengths are mapped to return angles that miss the dye gain medium and therefore are not amplified. This laser is tuned by varying the angle of the tuning mirror relative to the grating.

The major advantages of the grazing-incidence approach is that it is easy to implement with relatively inexpensive parts, it does not need careful focusing or alignment, and it allows for the cavity to be made very compact.

6.4.5 Single Longitudinal Mode Tunable Dye Laser

The above setup can be further improved to achieve single longitudinal mode operation (see Fig. 12). Unlike the previous ones, the new laser is end-pumped. The pumping light is focused into a dye cell by a spherical lens with a 75-cm focal length. The dye cell has a 2-mm pathlength and is polished on two sides. The end mirror is a fully reflecting mirror with a high damage threshold [16].

Since the dye cell's pathlength is only 2 mm, the gain per pass is small. As a result, the two-pass spot we used for alignment before is too weak to be useful. Therefore, we bring in an He–Ne laser beam pointed opposite to the dye laser output beam for alignment purposes. Care is taken to make sure that the alignment beam properly strikes the surfaces of the grating, end mirror, and tuning mirror. Alignment is completed when a single beam from the dye laser cavity is retroreflected back to the He–Ne laser. The pump laser is then focused onto the dye cell to a spot that is defined by the He–Ne laser. The pump light and the He–Ne light should form an angle of about 5°.

Note that at this time the angle of the tuning mirror corresponds to backreflection of the He–Ne wavelength (628 nm). Therefore, to make the laser oscillate, the tuning mirror is rotated to an angle that corresponds to a wavelength that is within

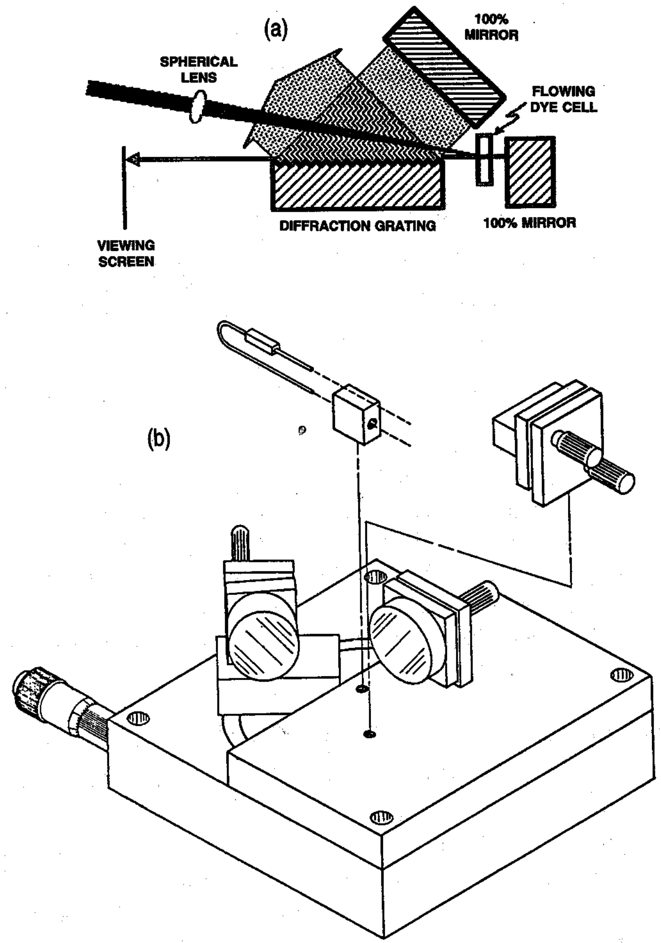

FIG. 12. Single-mode longitudinally pumped grazing-incidence dye laser: (a) top view; (b) exploded view.

the laser gain range. Once this is achieved, the laser will be oscillating, although it might be doing so in several longitudinal modes. Fine adjustment of the tuning mirror pitch angle, the position of the focus lens in the pump beam, and the grazing incidence angle of the grating is made while the output is monitored by a 20-GHz free-spectral-range Fabry–Perot etalon, until single-mode operation is achieved. The laser is continuously tunable by rotation of the tuning mirror about a unique pivot axis that is defined by the intersection of the surface planes of the grating, end mirror, and tuning mirror [17].

6.5 Summary

This chapter is not intended to be a complete review or full coverage of the subject of pulsed lasers. It is a practical guide to researchers who have had only limited exposure to the subject but are interested in building, using, and understanding lasers as tools for research. We first discussed the general concepts of laser operation and briefly introduced the readers to various kinds of pulsed lasers that are important in applications. A checklist was then developed for potential laser buyers. Finally, we provided a step-by-step illustration on how to build dye lasers with various levels of performance, so that the readers can get some hands-on experience with building lasers they can use and gain better understanding of the important issues in laser operation.

References

1. Maiman, T. H. (1960). *Nature* **187**, 493.
2. Fork, R. L., Brito Cruz, C. H., Becker, P. C., and Shank, C. V. (1987). *Opt. Lett.* **12**, 483.
3. Siegman, A. (1986). *Lasers*, University Science Books, Mill Valley, CA.
4. Hellwarth, R. W. (1961). In *Advances in Quantum Electronics*, p. 334, Columbia Univ. Press, New York.
5. Geusic, J. E., Marcos, H. M., and Van Uitert, L. G. (1964). *Appl. Phys. Lett.* **4**, 182.
6. Macklin, J. J., Kmetec, J. D., and Gordon III, C. L. (1993). *Phys. Rev. Lett.* **70**, 766; L'Huillier, A., and Balcou, P. (1993). *Phys. Rev. Lett.* **70**, 774.
7. Park, Y. K., Giuliani, G., and Byer, Robert L. (1984). *IEEE J. Quantum Electronics* **QE-20**, 117.
8. Sorokin, R. P., and Lankard, J. R. (1966). *IBM J. Res. Develop.* **10**, 162; and Eastman Kodak Laser Dye Technical Note and Catalog.
9. Moulton, P. F. (1982). Solid State Research Report No. DTIC AD-A124305/4, Lincoln Laboratory, MIT, Lexington, MA, pp. 15–21.
10. Rines, Glen A., and Moulton, Peter F. (1990). *Opt. Lett.* **15**, 434.
11. Basov, N. G., Krokhin, O. N., and Popov, Y. M. (1961). *JETP* **40**, 1320; see also Thompson, G. (1980). *Physics of Semiconductor Laser Devices*, Wiley, New York.
12. *Laser Focus World*, September 1992, p. 15.
13. Strome, F. C., and Webb, J. P. (1971). *Appl. Opt.* **10**, 1348.
14. Hänsch, T. W. (1972). *Appl. Opt.* **11**, 895; Schawlow, A. (1973). In *Fundamental and Applied Laser Physics*, M. Feld, A. Javan, and N. Kurnit (eds.), Wiley, New York.
15. Littman, M. G., and Metcalf, H. J. (1978). *Opt. Lett.* **3**, 138; Shoshan, I., Danon, N. N., and Oppenheim, U. P. (1977). *J. Appl. Phys.* **48**, 4495.
16. Littman, M. G. (1984). *Appl. Opt.* **23**, 4465.
17. Liu, K., and Littman, M. G. (1981). *Opt. Lett.* **6**, 117.