

Lasers

Understanding the Basics

Coherent, Inc.

Although lasers range from quantum-dot to football-field size and utilize materials from free electrons to solids, the underlying operating principles are always the same. This article provides the basic information about how and why lasers work.

Over 50 years have passed since the first demonstration of a laser in 1960. After the initial spark of interest, lasers were for a while categorized as “a solution waiting for a problem,” but bit by bit, the range of their applications has expanded to encompass fields as diverse as DNA sequencing, consumer electronics manufacturing, or freezing the motion of electrons around atoms. Most of these applications simply would not have been possible without lasers. To grasp the relevance of lasers in physics, it is enough to note that no other man-made sources can generate pulses (of any type) as short as laser pulses - now below to 10^{-16} s - or tools to measure absolute frequencies with an accuracy of $\sim 10^{-15}$! Industrial manufacturing, microelectronics, and biomedical and instrumentation applications serviced by lasers are incredibly diverse and rely on unique capabilities, like producing features below the limit of light diffraction, modifying materials in their bulk while leaving the surface unaffected, or trapping and moving individual particles in mid-air.

All light sources convert input energy into light. In the case of the laser, the input, or pump, energy can take many forms, the two most common being optical and electrical. For optical pumping, the energy source may be a lamp or, more commonly, another laser. Electrical pumping can be via a DC current (as in laser diodes), an electrical discharge (noble gas lasers and excimer lasers) or a radio-frequency discharge (some CO₂ lasers).

In a conventional (incoherent) light source like a lightbulb, an LED or a star, each atom excited by input pump energy randomly emits a single photon according to a given statistical probability. This produces radiation in all directions with a spread of wavelengths and no interrelationships among individual photons. This is called spontaneous emission.




Figure 1. Spontaneous emission is a random process, whereas stimulated emission produces photons with identical properties.

Einstein predicted that excited atoms also could convert stored energy into light by a process called stimulated emission. This process typically starts with an excited atom first producing a photon by spontaneous emission. When this photon reaches another excited atom, the interaction stimulates that atom to emit a second photon (Figure 1). This process has two important characteristics. First, it is multiplicative - one photon becomes two. If these two photons interact with two other excited atoms, this will yield a total of four photons, and so forth. Second and most importantly, these two photons have identical properties: wavelength, direction, phase and polarization. This ability to “amplify” light in the presence of a sufficient number of excited atoms leads to “optical gain” that is the basis of the laser operation and justifies its acronym of Light Amplification (by) Stimulated Emission (of) Radiation. A wide range of solid, liquid and gas-phase materials have been discovered that exhibit gain under appropriate pumping conditions.

The laser cavity

The laser cavity, or resonator, is at the heart of the system. A single transit through a collection of excited atoms or molecules is sufficient to initiate laser action in some high-gain devices such as excimer lasers; however, for most lasers, it is necessary to further enhance the gain with multiple passes through the laser medium. This is implemented along an optical axis defined by a set of cavity mirrors that produce feedback (Figure 2). The lasing medium (a crystal, a semiconductor or gas enclosed in an appropriate confinement structure) is placed along the optical axis of the resonator. This unique axis with very high optical gain becomes also the direction of propagation of the laser beam. A somewhat different example of a uniquely long (and flexible!) gain axis is the fiber laser.




Figure 2. In the prototypical gas laser, the gain medium has a long, thin cylindrical shape. The cavity is defined by two mirrors. One is partially reflecting and allows the output beam to escape.

The simplest cavity is defined by two mirrors facing each other - a total reflector and a partial reflector whose reflectance can vary between 30 and close to 100 percent. Light bounces back and forth between these mirrors, gaining intensity with each pass through the gain medium. Photons that are spontaneously emitted in directions other than the axis are simply lost and do not contribute to the laser operation. As laser light is amplified, some of the light escapes the cavity, or oscillator, through the partial reflector (output coupler); however, at equilibrium (the so-called “steady state” or “continuous wave”), these “optical losses” are perfectly compensated by the optical gain experienced in the successive round trip of the photons inside the cavity. The output of the laser is exactly the part of the beam transmitted by the output coupler. In an ideal laser, all the photons in the output beam are identical, resulting in

perfect directionality and monochromaticity. This determines the unique coherence and brightness of a laser source.

Monochromaticity - A photon's energy determines its wavelength through the relationship $E = hc/\lambda$, where h is Planck's constant, c is the speed of light, and λ is wavelength. An ideal laser would emit all photons with exactly the same energy, and thus the same wavelength, and it would be perfectly monochromatic. Many applications are dependent on monochromaticity. For example, in telecommunications, several lasers at slightly offset wavelengths can transmit in parallel streams of pulses down the same optical fiber without crosstalk. Real lasers are not perfectly monochromatic because several broadening mechanisms widen the frequency (and energy) of the emitted photons. For example, free-running YAG lasers can have linewidths of hundreds of gigahertz, while stabilized diode-pumped YAG lasers can have a linewidth <1 kHz. The best known of these broadening mechanisms is the Doppler broadening, determined by the distribution of speeds in the collection of atoms making up active gas mediums.

Coherence - Besides sharing the same wavelength, the photons that make up a laser beam are all in phase (Figure 3), or "coherent," resulting in an electric field that propagates with a uniform wavefront. The ideal representation is a plane wave that propagates with a flat wavefront along a given direction and where each plane perpendicular to this direction experiences the same electric and magnetic field amplitude and phase at a given time. When two waves with such characteristics interact, they create interference patterns, as in Young's experiment. Real laser beams somewhat deviate from this ideal behavior, but they are still the sources that best approximate an ideal coherent plane wave, and they enable a host of applications that rely on optical interference. For example, the surface of precision lenses and mirrors is measured using laser interferometers, and so are the minute variations in the interference patterns of miles-long interferometers used to chase and detect gravitational waves.




Figure 3. Laser light differs from conventional light in that all the lightwaves are in phase with each other.

Brightness (or, more correctly, radiance) - The most strikingly visible difference between lasers and conventional light sources is that all the emitted light travels in the same direction as an intense beam. Radiance is defined as the amount of light leaving the source per unit of surface area and unit of solid angle. A star like the sun emits a large amount of radiation from a unit of surface area, but this is emitted in many different directions. On the contrary, a laser beam is highly directional, with the result that its brightness is much more intense than the sun's as experienced on the Earth's surface. For this reason, just 5 mW of power from a laser pointer is more "blinding" (and dangerous) for the eye than direct sunlight.

Because of its high radiance, a laser beam can be projected over great distances or focused to a very small spot. Well-designed lasers produce a beam of light that will expand ("diverge") only by the minimum amount prescribed by the laws of diffraction. For example, diffraction imposes that the minimum spot that can be produced by a laser beam is equal to about its wavelength.

Continuous-wave lasers

Lasers can be divided into three main categories: continuous wave (CW), pulsed and ultrafast.

As their name suggests, continuous-wave lasers produce a continuous, uninterrupted beam of light, ideally with a very stable output power. The exact wavelength(s) or line(s) at which this occurs is determined by the characteristics of the laser medium. For example, CO₂ molecules readily lase at 10.6 μm , while neodymium-based crystals (like YAG or vanadate) produce wavelengths in the range between 1047 and 1064 nm. Each laser wavelength is associated with a linewidth, which depends on several factors: the gain bandwidth of the lasing medium and the design of the optical resonator, which may include elements to purposely narrow the linewidth, like filters or etalons.

If a laser can simultaneously produce different lines, the first step in determining the operating wavelength is to use cavity mirrors that are highly reflective only at the desired wavelength. The low reflectivity of the mirrors at all the other lines will prevent these from reaching the threshold for laser action. However, even a single laser line actually covers a range of wavelengths. For example, laser diodes produce light over a wavelength range of several nanometers corresponding to their "gain bandwidth."

The specific wavelengths of the output beam within this gain bandwidth are determined by the longitudinal modes of the cavity. Figure 4 shows the behavior of a two-mirror cavity, the most basic design. To sustain gain as light travels back and forth between the mirrors, the waves must remain in phase and "reproduce" their wave pattern, which means that the cavity round-trip distance must be an exact multiple of the wavelength

$$N\lambda = 2 \times \text{Cavity Length},$$

where λ is the laser wavelength and N is an integer called the mode number. This is usually a very large integer, since the wavelength of light is so much smaller than a typical cavity length. In a high-power laser diode, for example, the IR output wavelength is 0.808 μm , yet the cavity length may be 1 mm, so that even in a very small laser resonator, N is ~ 2500 . Wavelengths that satisfy this resonance equation are called longitudinal cavity modes. The actual output wavelengths of the laser will correspond to the cavity modes that fall within the gain bandwidth, as shown in Figure 4 (bottom). This regime is called multilongitudinal-mode operation. Using the example of the high-power laser diode, the spacing between adjacent longitudinal modes is

$$\sim 150 \text{ GHz (equivalent to } \sim 0.3 \text{ nm in wavelength difference).}$$

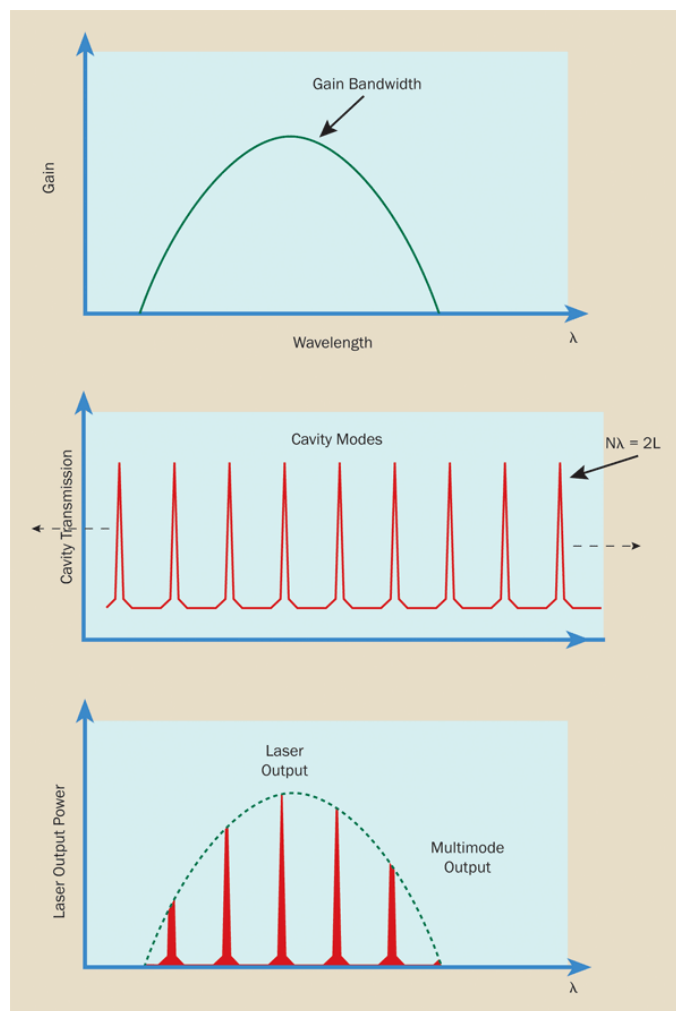


Figure 4. A resonant cavity supports only modes that meet the resonance condition, $N\lambda = 2 \times \text{cavity length}$. The output of a CW laser is defined by the overlap of the gain bandwidth and these resonant cavity modes.

If the laser diode operates on a 3-nm gain line, about 10 longitudinal modes, spanning 3 nm, will be able to oscillate. The resonator design also controls the so-called transverse modes, responsible for the intensity distribution on the plane perpendicular to the beam direction. The ideal laser beam has a radially symmetric cross section: The intensity is greatest in the center and tails off at the edges, following a Gaussian profile. This is called the TEM_{00} or fundamental output mode. Lasers can produce also many other TEM modes, a few of which are shown in Figure 5. Usually a round aperture placed inside the cavity is used to force the laser to operate in the fundamental mode. In multitransverse-mode operation, many modes are present at the same time, often resulting in a profile that appears to be Gaussian but in reality has degraded properties (higher divergence and lower radiance). Often the quality of a laser beam is specified using the M^2 (M squared) parameter. For example, a YAG laser operating exclusively in the TEM_{00} mode has $M^2 = 1$, while multimode laser diodes have an M^2 of hundreds. Different transverse modes also have slightly different frequencies; however, this difference is much smaller than the difference between adjacent longitudinal modes (~1 MHz compared with approximately hundreds of megahertz to hundreds of gigahertz).

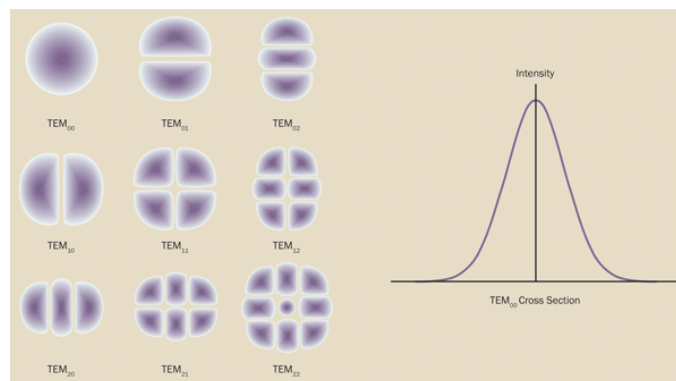


Figure 5. Lasers can emit any number of transverse modes, of which the TEM_{00} usually is most desirable.

A laser that produces multiple longitudinal modes has a limited coherence - different wavelengths cannot stay in phase over extended distances. Applications such as holography, which demand excellent coherence, benefit from using a single-longitudinal-mode laser. For some laser types with a narrow gain bandwidth, single-mode output is achieved with a very short resonant cavity; this makes the mode spacing larger than the gain bandwidth, and only one mode lases. Generally, though, a filtering element that preferentially passes only one mode is inserted into the cavity. The most common type of filter is called an etalon. Using a number of sophisticated design enhancements, it is possible to restrict the linewidth of a laser to less than 1 kHz, useful for scientific interferometric applications.

Some solid-state lasers have extremely broad bandwidths that extend to hundreds of nanometers. The most common example is the Ti:sapphire laser. Rather than being a disadvantage, this broad bandwidth enables the design of tunable and ultrafast (femtosecond and picosecond pulse width) lasers. Designing a tunable CW laser involves including an extra filtering element in the cavity - usually a birefringent (or Lyot) filter. A birefringent filter does two things: It narrows the bandwidth and, by rotating the filter, allows smooth tuning. This same type of filter is also used as a factory-set tool to lock the wavelength at a precise value, when broad-bandwidth lasers need to be preset at a specific application-dependent wavelength. This is typically the case with optically pumped semiconductor lasers (OPSLs) that can be set at the desired wavelength within their 5- to 10-nm operating range.

Most applications of CW lasers require that the power be as stable as possible over long time periods (hours or weeks), as well as over short time durations (microseconds), depending on the specific application. To ensure this stability also in the presence of varying environmental situations like temperature, vibration and the aging of the laser itself, microprocessor control loops are implemented. For example, a diode-pumped Nd laser will have servos to adjust temperature and output power of the pump diodes to maintain stable output power from the resonator. In addition, other servos may control the perfect alignment of the resonator mirrors.

Pulsed lasers

Some materials - like excited dimers (or "excimers") of a noble gas with a halogen, such as ArF and XeCl - sustain laser action for only a brief period of several nanoseconds. Other lasers, like Nd or Yb diode-pumped solid-state (DPSS) lasers, lend themselves to be operated both in CW or pulsed operation. Other lasers, like laser diodes or OPSLs, are not suitable at all for pulsed operations. Within this context, we define as "pulsed" laser devices that produce pulses of 0.5 to 500 ns. This regime is useful for time-resolved scientific experiments but especially for a vast range of manufacturing processes related to ablation or some other type of nonthermal materials modification. The most important characteristic of a nanosecond-pulsed laser is the capability to "store" and release energy very rapidly; i.e., on a nanosecond scale so that the laser output can achieve tens of kilowatts to megawatts of peak power. It is precisely this high peak power that enables the ablative processing of materials. In addition, the high peak power enables a number of so-called optical nonlinear processes; i.e., processes that rely on the interaction of more than one photon at a time with matter.

Operating a nanosecond-pulsed laser is substantially different from operating a CW laser. To build and produce each pulse, the light has time for very few round trips in the laser cavity, and the simple two-mirror cavity based on a partly transmissive mirror described so far cannot produce these energetic and short pulses. The key to producing these energetic pulses is storing energy from the pump in the atoms or molecules of the lasing medium by preventing the laser gain and the amplification process. Then, when the stored energy is at its maximum, lasing action is rapidly enabled: The stored energy results in an extremely high laser gain (amplification) that takes place during only a few round trips, during which a giant pulse builds up and gets coupled through the partly transmissive mirror. This regime is called Q-switched operation and can be conceptualized as a two-mirror cavity with an optical gate located between one of the mirrors and the laser medium (Figure 6). When the gate is closed and the laser medium is pumped, photons cannot circulate in the cavity, and the excitation of the atoms builds up; as soon as the gate is opened, photons start to build up via stimulated emission with a very large gain at each round trip; a fraction of them (~20-40%) get coupled by the partly transmissive mirror. The result is a pulse with a very sharp raising time and a slower falling time, with a typical duration of 1 to 200 ns. The pulse duration depends on several parameters: the type of gain medium and how much energy it can store, the cavity length, the repetition rate of the pulses and the pump energy, to mention the most important ones. Q-switched lasers commonly used in the industry can produce average powers up to tens or hundreds of watts and repetition rates as low as 10 Hz or as high as 200 kHz. Most industrial processes are in the kilohertz to tens-of-kilohertz regime.




Figure 6. Schematic showing the operating principle of a Q-switch in a solid-state laser.

The actual Q-switch device is an acousto-optical modulator or an electro-optical modulator (EOM). Both use crystals where an applied electric field produces some perturbation of the optical properties of the crystal. In the case of acousto-optical modulators, the applied electric field is a radio-frequency voltage that produces a high-frequency sound wave in the crystal. This sound wave diffracts the photons from the laser and prevents laser amplification. EOMs instead use an applied high voltage that modifies the crystal refractive index and alters the polarization of the incoming light; an appropriate combination of polarization-sensitive optics can be placed in the cavity to prevent light of altered polarization from circulating.

Other types of lasers, such as excimer lasers, do not require a Q-switch to produce nanosecond pulses but rather rely on a transient pump pulse: Excimer laser pulses are produced by exciting the noble gas/halogen mixture with a powerful and short electric discharge. Ti:sapphire lasers can also produce nanosecond

pulses if they are pumped with a nanosecond pulse of green light produced by a frequency-doubled, Q-switched YAG laser. This method is called gain switching because the cavity gain rather than the cavity loss is directly changed.

Apart from a huge number of industrial applications, Q-switched lasers have important applications in scientific research. One is pumping of Ti:sapphire ultrafast amplifiers (described in the following section) by using the frequency-doubled (green) output of a Q-switched Nd:YAG or Nd:YLF at 1-10 kHz. Another one is using the YAG or YLF laser to produce energies per pulse in the joule range at 1-100 Hz. These lasers are often used with nonlinear optical generators that can produce tunable wavelengths in the UV, visible and IR region, enabling time- and wavelength-resolved studies. Nowadays most YAG or YLF lasers operating at >100 Hz are diode-pumped, while high-energy 10-Hz systems require pumping with a flashlamp because diodes are not suitable for producing high-energy output pulses.

For some scientific applications, it may be desirable to have a narrow-linewidth Q-switched laser. In some cases, this can be accomplished using a combination of optical gratings and etalons; in other cases, the laser can be “seeded” with a low-power CW or Q-switched narrow-linewidth laser that is easier to control than the higher-power stage. This approach, called “injection seeding,” uses a MOPA (master oscillator, power amplifier), conceptually splitting the linewidth selection and the high-power generation into two stages that are optimally designed for the two purposes.

Ultrafast lasers

Ultrafast lasers are generally defined as lasers that produce pulses in the range of 5 fs to 100 ps (1 femtosecond = 10^{-15} seconds). CW lasers can produce many longitudinal modes; if all these modes can be locked in phase (mode-locking regime), the resultant overlap of the optical field of all the modes will not be a standing wave but rather a distribution corresponding to very short pulses spaced by the cavity round trip. This mode-locked regime results in a single pulse traveling back and forth inside the cavity: Every time this pulse reaches the output coupler, the laser emits a part of this pulse. The pulse repetition rate is determined by the time it takes the pulse to make one trip around the cavity.

To understand the concept of locking together these modes, one can think about the various modes as having a “zero” (i.e., null electric field), all at a certain position (function of time) in the resonator. It turns out that the more modes that interfere, the shorter the pulse duration (Figure 7). Since larger lasing bandwidths support a larger number of oscillating modes, the pulse duration is inversely proportional to the bandwidth of the laser gain material. This explains why materials used for broadly tunable lasers produce also the shortest mode-locked pulses. The most popular ultrafast laser material is titanium-doped sapphire or Ti:sapphire, thanks to its large bandwidth and broad tuning range; turnkey commercial Ti:sapphire lasers can deliver pulses as short as 6-10 fs (10×10^{-15} s). Ti:sapphire lasers are universally pumped using a green-wavelength CW pump laser, usually a frequency-doubled OPSSL or Nd-based laser, but the future availability of blue or green high-power diodes may shift this balance.

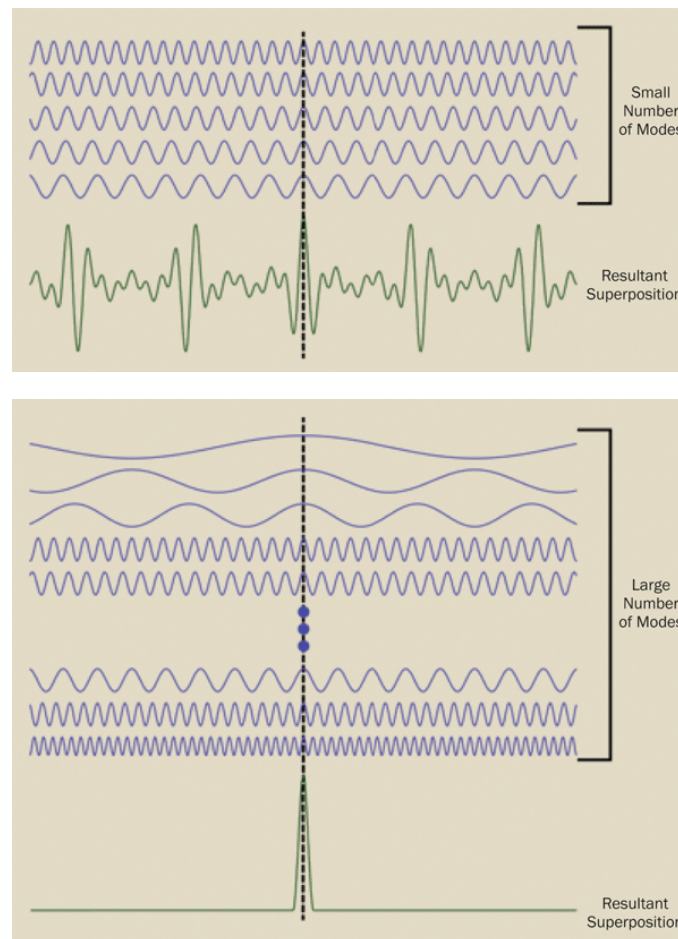


Figure 7. When a very large number of laser modes that all have a “zero” in the same position interfere, the resultant superposition is an extremely narrow pulse.

With cavity lengths of about 1.5 to 2 m, typical repetition rates are 75-100 MHz, and the average power can reach 4 to 5 W. The peak power of a mode-locked 4-W Ti:sapphire laser operating at 100 MHz and producing 100-fs pulses is then $4 / (100 \times 10^{-15} \times 100 \times 10^6) = 400$ kW.

It is important to note that there is simply no other man-made source that can produce events on such a short timescale, a timescale that is faster than the formation and dissociation of chemical bonds in molecules. Moreover, the high peak power of these lasers enables the generation of many nonlinear processes; that is, processes where two, three or more photons simultaneously interact with a target. Thanks to the short pulse duration and high peak power, the advent of femtosecond lasers in the 1990s enabled a true revolution in science that resulted in new areas of research as diverse as femtochemistry or nonlinear optical microscopy. In addition to scientific applications, pico- and femtosecond lasers are now employed in industrial materials processing applications that require ablation or materials modification without any residual thermal effect and/or on a submicron spatial scale.

In the past few years, new laser materials have emerged that can produce femtosecond-class pulses. These lasers are mostly based on Yb atoms in various host crystals (YAG, KGW, KYW, etc.). One important distinction between Ti:sapphire and Yb lasers is that Yb lasers are limited to producing pulses of ~100-200 fs, while Ti:sapphire can achieve <10 fs; secondly, Ti:sapphire lasers are tunable over as much as 400 nm (680-1080 nm), while Yb lasers can be tuned only over 10-20 nm at the most (1040-1060 nm). This means that ultimate tuning range and short pulse duration reside with Ti:sapphire lasers. On the other hand, the simplicity of directly diode-pumped Yb lasers and their capability to produce higher average powers (in excess of 100 W) means that industrial applications will be best addressed by the laser type.

Yb also lends itself for use as an active medium for ultrafast fiber lasers. These lasers use a diode-pumped optical fiber with an integrated mode-locking mechanism to produce 100-fs to 10-ps pulses with wavelengths around 1 μm , while fiber lasers based on Er produce wavelengths around 1550 nm and are now a common staple in optical fibers for telecom applications.




Figure 8. The basic elements and operation of a chirped pulse amplifier (CPA).

Mode-locked lasers produce pulses at a high repetition rate and with average powers of 1-100 W. The high repetition rate, necessary to support mode-locking, limits the energy per pulse to tens or hundreds of nanojoules. For some (mostly scientific) applications, it is desirable to have millijoule- or even joule-level energies with peak powers of hundreds of megawatts up to a petawatt. This can be achieved through chirped pulse amplification (CPA) (Figure 8). CPA is a technique that uses a mode-locked oscillator as a first building block. The output of this megahertz-repetition-rate oscillator (~20-100 fs) is stretched to become ~50-200 ps. A fraction of these 50- to 200-ps pulses is injected into an amplification stage, where the pulses are amplified to the millijoule level. At this point, the pulses are ejected from the amplification stage, usually with an EOM, and become recompressed close to the original pulse duration of 30-100 fs or less. The word “chirped” refers to the fact that the pulses are stretched to longer duration to avoid damage to the optical components of the amplifier stage caused by the high peak power. The amplifier of course requires optical pumping to sustain gain and amplification; in the case of Ti:sapphire, most CPA systems operate at 1-10 kHz or ~100 kHz. The lower-repetition-rate systems are energized by a Q-switched pump laser, while the higher-repetition-rate ones are pumped by CW green pump lasers, of the same type used to pump a mode-locked Ti:sapphire laser. Yb systems can also use a CPA scheme but, because of the diode pumping, they are more suitable to operation at higher repetition rates and lower energy per pulse than Ti:sapphire systems.

Another distinct class of mode-locked, ultrafast lasers has been developed for industrial uses. Scientific ultrafast lasers are typically designed to deliver cutting-edge performance in at least one of three parameters: namely, pulse duration, pulse energy or repetition rate. In contrast, industrial ultrafast lasers are usually designed to produce a balance between these characteristics that enables materials processing at practical rates, while also delivering high operational reliability at the lowest possible cost. Typically, these industrial ultrafast lasers have pulse widths in the 10-ps range and average powers that range from about 10 up to 100 W.

Frequency doubling and harmonic generation

Even with the broad choice of commercially available lasers, it is not always possible to find one that exactly matches the wavelength required by a specific application. Ti:sapphire lasers are broadly tunable, but in most cases, they are too complex for industrial applications and unable to reach the all-important UV region of the spectrum. OPSLs are simple and can be designed at many wavelengths in the 920- to 1160-nm region but are not ideal for pulsed operation. To achieve the desired wavelength in just about any regime of operation - CW, pulsed or ultrafast - the processes of harmonic frequency conversion and parametric generation provide wavelength flexibility when used in conjunction with the lasers described so far. All these processes are related and are called nonlinear phenomena since they depend nonlinearly on the laser peak power. That is, they are proportional to the square, third or higher power of the laser output power.

In simple terms, when an intense and/or tightly focused laser beam passes through a suitable crystal, its oscillating electric field interacts with the electrons of the crystal in several ways. One of these mechanisms distorts the electron cloud in the crystal, thereby polarizing the atoms at a frequency that is the same as that of the laser beam, but also at a frequency that is its double (nonlinear polarization). This frequency corresponds to a wavelength that is half that of the incoming laser. The nonlinear polarization is much smaller than the linear term, but it depends on the square of the laser power, therefore increasing more strongly in the presence of an intense laser pulse. It generates an optical field at double the frequency of the original laser beam, with the result that part of the incoming laser power will be converted to half the original wavelength (second-harmonic generation (SHG) or frequency doubling) (Figure 9). Since energy has to be conserved, any gain in the SHG beam is traded for a decrease in power of the original beam. In some cases, it is possible to achieve an almost total conversion of the original (“fundamental”) beam into its second harmonic. Common crystals for SHG are BBO, LBO and KDP. The most common example of SHG is the conversion of a Nd-based laser IR output at 1064 nm into a green output at 532 nm (green), constituting the most popular visible wavelength, used ubiquitously to pump Ti:sapphire lasers.

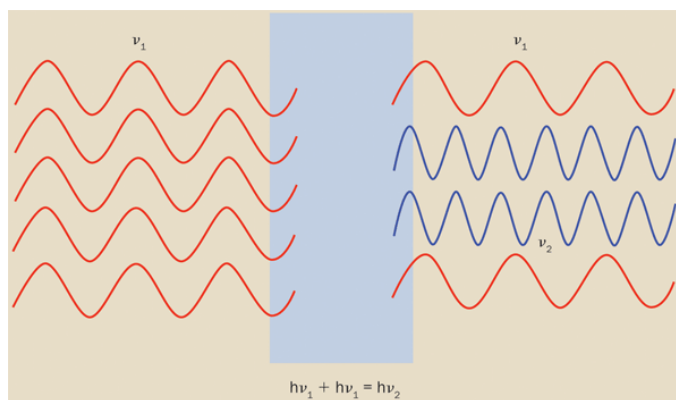


Figure 9. Basic functioning of a second-harmonic generating crystal.

Efficient SHG can be achieved only under a condition called “phase matching.” Under most conditions, the light at the new frequency would be reconverted to the original frequency and lost or simply not added in phase to create any sizable power. This difficulty is overcome by choosing a crystal temperature and orientation that create the so-called phase-matching condition where the phase velocities of the fundamental and second-harmonic light are the same. This happens by choosing a specific direction of propagation (usually a function of temperature and wavelength) in the crystal such that the two waves propagate at the same velocity.

Extensions of the SHG process are third-harmonic generation (THG), where the wavelength at one-third the incoming wavelength is created by the interaction of an SHG beam with its fundamental; and fourth-harmonic generation (FHG), where the SHG beam is frequency-doubled again. All of these harmonic processes can be generalized as frequency-mixing, where two coherent beams at different wavelengths are mixed to produce sum-frequency and difference-frequency generation (SFG and DFG, respectively).

Harmonic generation can be applied to CW, pulsed and ultrafast lasers, greatly expanding the range of available wavelengths. Pulsed or ultrafast lasers have enough peak power (kilowatt range) to achieve relatively high conversion efficiency in a single pass through the harmonic crystal. On the other hand, CW lasers usually do not produce sufficient power for efficient harmonic generation, so that the power in the crystal has to be enhanced by putting the nonlinear crystal inside the laser cavity (“intracavity doubling”), or building a cavity ad hoc around the crystal (“resonant doubling”) that matches the modes of the original CW laser cavity.

Optical parametric generation

In the SFM process described in the previous section, two photons at different frequencies interact to produce a single photon at a frequency that is the sum of the two initial frequencies. The opposite process is also possible: A single photon interacts with a suitable crystal and disappears, creating two photons of lower and different (called “nondegenerate”) energies. This process, called optical parametric generation, is useful because it results in the generation of two new tunable wavelengths, bound solely by the conservation of energy and momentum, and by the nonlinear crystal refractive indices n:

$$\nu_p = \nu_s + \nu_i$$

$$1/(n_p \lambda_p) = 1/(n_s \lambda_s) + 1/(n_i \lambda_i)$$

The subscripts p, s and i refer to the pump wavelength and to the two new wavelengths that are called (for historical reasons) “signal” and “idler,” with the signal wavelengths being the shorter of the two and both being longer than the pump wavelength. Phase matching takes place when the refractive indices of the three wavelengths n_p , n_s and n_i satisfy the above momentum conservation equation in the crystal. This can be accomplished by changing the temperature or angle of the crystal, or the pump wavelength (if the pump laser is tunable) so that phase matching takes place at the desired signal or idler wavelength.

While the wavelength of an SHG beam is unambiguously determined by the pump wavelength, the OPG process can generate an infinite set of wavelength pairs. Amplification of the desired wavelength pair requires not only phase matching in the crystal but also “jump starting” the process from the noise of the statistical distribution of wavelength pairs. This is exactly what happens in an optical parametric oscillator (OPO) or an optical parametric amplifier (OPA), both advanced laser accessories able to produce tunable outputs anywhere from the mid-UV to the mid-IR.

In an OPO, the signal and possibly the pump are resonated in a laser-like cavity, where the desired signal wavelength starts from the noise (random distribution of signal/idler pairs) and is amplified by passing through the crystal properly phase-matched for that wavelength at the same time as the pump during each round trip (Figure 10). In an OPA, a sapphire or YAG disk is pumped to produce a relatively bright beam of white light that, as the name suggests, contains all the wavelengths of the visible and near-IR spectrum. The unique wavelength pair that is phase-matched in the OPA crystal will then be amplified when the crystal is pumped by the pump laser.

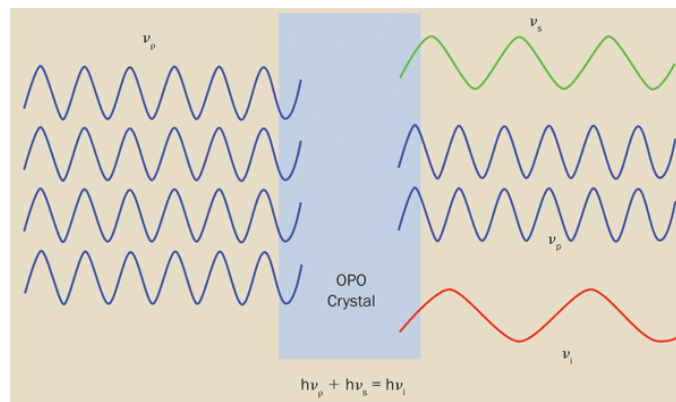


Figure 10. An optical parametric oscillator (OPO) converts an input photon into two photons having lower energy, and which conserve the energy and momentum of the input.

Nano-, pico- and femtosecond OPOs are complex devices that are implemented in conjunction with pulsed and ultrafast pump lasers. CW OPOs are equally, if not more, complex. OPAs are easier to design and build but require a more energetic pump to produce the white light and one-pass amplification in the crystal. For this reason, they are pumped by CPA pico- or femtosecond amplifiers producing at least several microjoules. The addition to an OPA/OPO of one or more stages of harmonic generations and mixing yields a range of wavelengths that can cover 200 nm to 20 μm .

Common laser types

For many years, the most common CW laser was the helium neon laser, or HeNe. These low-power lasers (a few milliwatts) use an electric discharge to create a low-pressure plasma in a glass tube; nearly all emit in the red at 633 nm. In recent years, the majority of HeNe applications have switched to visible laser diodes. Typical applications include bar-code readers, alignment tasks in the construction and lumber industries, and a host of sighting and pointing applications ranging from medical surgery to high-energy physics.

In fact, the laser diode has become by far the most common laser type, with truly massive use throughout telecommunications and data storage (e.g., DVDs, CDs). In a laser diode, current flow creates charge carriers (electrons and holes) in a p - n junction. These combine and emit light through stimulated emission. Laser diodes are available as single emitters with powers up to tens of watts, and as monolithic linear bars with numerous individual emitters. These bars can be assembled into 2-D arrays with total output powers in the kilowatt range. They are used in both CW and pulsed operation for so-called direct diode applications. But even more importantly, laser diodes now underpin many other types of lasers, where they are used as optical pumps that perform the initial electrical-to-optical power conversion.

For example, higher-power visible and UV CW applications were originally supported by argon-ion and krypton-ion lasers. Based on a plasma discharge tube operating at high current, these gas-phase lasers are large and inefficient, generating a large amount of heat that must be actively dissipated. The tube also has a finite lifetime and thus represents a costly consumable. In most former applications, the ion laser emitting at blue or green wavelengths was displaced by DPSS lasers. Here, the gain medium is a neodymium-doped crystal (usually Nd:YAG or Nd:YVO₄) pumped by one or more laser diodes. The near-IR fundamental at 1064 nm is then converted to green 532-nm output with the use of an intracavity doubling crystal.

The DPSS laser, in turn, has been challenged by several newer technologies, with the OPSL the most successful of these. Here the gain medium is a large-area semiconductor laser that is pumped by one or more laser diodes. The OPSL offers numerous advantages, most notably wavelength and power scalability. Specifically, these lasers can be designed to operate at virtually any visible wavelength, at last freeing applications from the restrictions of limited legacy-wavelength choices (i.e., 488 and 514 nm from argon-ion lasers and 532 nm from frequency-doubled YAG lasers). Indeed, OPSLs represent a paradigm shift in lasers because they can be designed for the needs of the application instead of vice versa.

OPSL is now a leading technology in low-power bioinstrumentation applications, most notably at 488 nm; the power scalability and inherent low noise of OPSL technology is now seeing multiwatt green and yellow OPSLs moving strongly into other applications, including scientific research, forensics, ophthalmology and light shows.

While YAG and other neodymium crystal hosts lend themselves to operation in CW, Q-switched and mode-locked operations, laser diodes, OPSL and ion lasers do not support Q-switched operation and are virtually not used in mode-locked regime.

At longer wavelengths, carbon dioxide lasers, which use plasma discharge technology, emit in the mid-IR around 10 μm . Most are CW or pseudo-CW, with commercial output powers from a few watts to several kilowatts. Another important technology is the fiber laser, which can be operated in CW, Q-switched and mode-locked formats. Here, laser diodes optically pump a rare-earth doped fiber, which typically emits at about 1 μm .

Nd:YAG, CO₂, fiber and direct diode lasers are the workhorses of industrial laser applications. Direct diode lasers predominantly service low-brightness applications, such as heat treating, cladding and some welding applications. This is because direct diode lasers offer the lowest capital cost of any industrial laser type, as well as the lowest operating costs, due to their high electrical efficiency. On the downside, high-power laser diodes or arrays cannot deliver anything close to the diffraction-limited beam provided by other laser types.

The advent of slab-discharge technology has allowed the size:power ratio of CO₂ lasers to be greatly scaled down, increasing their utility in subkilowatt applications. Low-cost waveguide designs also support a healthy market for CO₂ lasers with powers in the tens of watts, primarily in marking and engraving applications.

Sealed CO₂ lasers and fiber lasers have come to dominate metal cutting in the 2- to 4-mm thickness range. Sealed CO₂ is usually the first choice when both metals and nonmetals must be processed, while fiber lasers have proved quite successful in certain markets that can benefit from their combination of high repetition rate, low pulse energy and high brightness. They also excel at metal cutting and welding in the 4- to 6-mm thickness range, as well as some marking applications. Flowing gas CO₂ lasers still dominate the market for cutting thick metal (>6 mm).

Nd:YAG can deliver the high peak power for materials processing applications such as metal welding. In these heavy industrial applications, raw power is more important than beam quality, and for many years, these lasers were lamp-pumped. But the ever-increasing power and lifetime characteristics of laser diodes are causing these lasers to switch to diode pumping; i.e., DPSS lasers.

Conversely, lower-power Q-switched DPSS lasers are often based on Nd:YVO₄. These are usually optimized for high beam quality for micromachining and microstructuring applications with high repetition rates (up to 250 kHz) to support high throughput processes. They are available with powers up to tens of watts with a choice of near-IR (1064 nm), green (532 nm) or UV (355 nm) output. The UV is popular for producing small features in “delicate” materials because it can be focused to a small spot and minimizes peripheral thermal damage. Deep-UV (266 nm) versions are starting to be used in some applications, but their relatively high cost and the need for specialty beam delivery optics causes many potential applications to rely instead on 355-nm lasers optimized for short pulse duration, which can produce similar results in many materials.

Excimers represent another important pulsed laser technology. They can produce several discrete wavelengths throughout the UV; depending on the gas combination, emission ranges from 157 to 348 nm. The deep-UV line at 193 nm is the most widely used source for lithography processes in the semiconductor industry. The 308-nm wavelength is used for annealing silicon in high-performance displays. The same wavelength is also key to generating a unique long-wear surface on the cylinder liners of high-performance diesel engines. And finally, excimers have a unique ability to produce high pulse energies - up to one joule per pulse. This enables direct writing of low-cost electronic circuits for applications such as medical disposables.

Ultrafast lasers for scientific applications are dominated by Ti:sapphire, as already described. Ultrafast lasers are also a fast-growing technology for micromachining and other high-precision materials processing applications. While there is some diversity in the form and construction of commercially available industrial ultrafast lasers, they all utilize a certain basic configuration. Specifically, a passively mode-locked oscillator is used to generate output at the pulse widths of about 10 ps or shorter that are necessary to drive photoablation. However, most mode-locked oscillators produce relatively low energy pulses (in the nanojoule range) at repetition rates in the tens of megahertz. Best results for micromachining are achieved when the pulse-to-pulse overlap is in the range of 50 to 70 percent. In other words, the beam deflection mechanism moves the beam about one-third of the beam diameter before the arrival of the next ultrafast pulse. Consequently, a repetition rate in the range of tens of megahertz is too high to be used with existing scanning technology, so a pulse picker selects a fraction of these pulses. The energy of these pulses is then boosted in an amplifier to produce the final output.

Most commercial picosecond products are based on one of the following architectures:

- A fiber laser oscillator followed by a fiber- or rod-type amplifier
- A fiber laser oscillator followed by a free-space amplifier
- A diode-pumped solid-state oscillator followed by a free-space amplifier

The all-fiber (oscillator and amplifier) approach has the advantage of being relatively low cost and holds the potential of being robust. The big negative is that nonlinearities, scattering and other effects in the fiber amplifier limit the maximum per-pulse energy that can be attained to about 10 μ J (at a 10-ps pulse width). This level of pulse energy can cater to some applications, but a large number of applications are served with pulse energies in the 100- μ J range. One can use specialty fibers to increase the mode inside the fiber and thereby allow for larger pulse energies, but such fibers lead to limited bend radii and hence bring their own packaging limitations.

To achieve the higher pulse energies required for most applications, a fiber oscillator can be mated with a free-space amplifier. Because of the relatively low energy output from the fiber seed, a regenerative amplifier is often used. In a regenerative amplifier, a pulse undergoes a large number of passes and can therefore be amplified substantially. Regenerative amplifiers also have the advantage of being compact and delivering good beam performance.

The third approach is to use a diode-pumped solid-state oscillator (usually with Nd:YVO₄ as the gain medium), which can produce higher pulse energies than a fiber seed. This is followed by a free-space amplifier, typically in either a regenerative or multipass configuration. In fact, more than one amplifier stage can boost power to levels as high as 100 W.

And finally, many other types of niche and exotic lasers exist that are beyond the coverage of this overview article. Examples include Raman lasers used in telecommunications, quantum cascade lasers used in some gas-sensing applications, and chemical lasers, which tend to be limited to military programs.