# Chapter 9
# Boosting the Performance of a Pulsed Laser

## 9.1 Introduction

There are seemingly endless applications of lasers, such as medical, military, material processing, remote sensing, laser fusion, science with ultrahigh intensity light, etc., that call for coherent light of the kind of intensity a cw laser is incapable of delivering. As we have seen in Chap. 7, pulsing the laser can allow scaling up of the output power or intensity of a laser albeit over a limited period of time. Although specific applications will form the kernel of volume II of this book, we need to emphasize here that the requirement on the magnitude of intensity varies widely from application to application. For instance, while modest intensity over a relatively longer period may suffice for certain material processing applications, fundamental research aimed at unraveling the design of nature essentially calls for a flash of light of unthinkably high intensity. The very onset of stimulated emission, which is the key to the realization of a laser, restricts the peak power and, in turn, the intensity of the pulse it gives out. Practical limitation on the rate at which energy can be pumped into the active medium is yet another factor that also puts a brake on the maximum achievable peak power from a pulsed laser. These seemingly unsurmountable challenges have largely been overcome by compressing the duration of the pulse rather than enhancing its energy content.[1] The great advances made in the area of laser physics, not long after the first laser flashed its light, led to the conceptualization of techniques such as Q-switching, cavity dumping, and modelocking to achieve pulse compression. The practical implementation of these techniques was made possible by exploiting the remarkable advantages offered by saturable absorbers or devices such as electro-optic, acousto-optic, or optical Kerr media. From the modest

---

[1] Power, as we know, is the rate at which energy is delivered. Clearly therefore an ultrahigh power laser will be capable of delivering either enormously large energy over a modest length of time or a modest amount of energy over an extremely short duration.

D. J. Biswas, *A Beginner's Guide to Lasers and Their Applications, Part 1*,
Undergraduate Lecture Notes in Physics,
https://doi.org/10.1007/978-3-031-24330-1_9

beginning of the production of nanosecond optical pulses in the 1960s, we have now graduated through picosecond into the regime of femtosecond and even beyond. A femtosecond pulse straight out of the laser oscillator can accomplish seemingly unthinkable tasks. To name but a few, when focused on a tiny spot, it performs as a scalpel of extraordinary sharpness allowing a surgeon the luxury of performing a high-precision job such as a delicate eye surgery or drilling a tiny hole on the wall of the heart; tracking the minute details of all the intermediate species formed during a chemical reaction defeating the lightning speed of their formation, a concept that impacted the field of reaction dynamics so profoundly that its originator Ahmed Zewail (1946–2016), nicknamed the “father of femto-chemistry,” was awarded the 1999 Nobel Prize in Chemistry; offering seemingly limitless bandwidth in data communication; etc. The energy content of this femtosecond pulse from the laser oscillator, however, needs to be elevated leaps and bounds before it reaches a level when, upon focusing, it would mimic the extreme condition that prevails at the center of a star, a pathway to big new science for knowing the unknown. As we shall see later in this chapter, the power of this ultrashort pulse can barely be boosted by letting it pass through an amplifier as its growing power invariably causes optical damage to the amplifier itself. Gerard Mourou (b–1944) and his doctoral student Donna Strickland (b–1959) at Michigan University succeeded in circumventing this problem in 1985 when they cleverly employed the chirped pulse amplification (CPA) technique. This made possible for the kilowatt ($10^3$) ultrashort (~100 fsec) oscillator pulse to leap coolly to tens of terawatt ($10^{12}$), a colossal optical power that even dwarfs the output of all the power plants of the world put together, although just at the moment of its flashing! It is no wonder that the conceptualization and implementation of the CPA fetched Mourou and Strickland a Nobel Prize in Physics, albeit a little belatedly in 2018. In addition to providing deeper physical insight into the various techniques to achieve pulse compression, this chapter will also address the challenges of amplifying an ultrashort pulse in the conventional way and overcoming them by adopting the technique of CPA.

## 9.2 Q-Switching

We know by now that the peak power of the laser pulse is essentially determined by the maximum value of the population inversion and, in turn, the gain that can be realized in the operation. In conventional operation, the phenomenon of gain clamping pins the gain down to the loss line and consequently puts a check on the maximum peak power of the emitted pulse. This situation is schematically illustrated in the top trace of Fig. 9.1. Clearly, the inability of the process of pumping to push the gain up beyond the loss causes the laser to underperform with respect to the emitted power. A possible remedy to this would be to push the loss line a long way up by way of appropriately augmenting the cavity loss. This would allow gain to rise much beyond what is possible in the conventional operation of the laser. An extreme

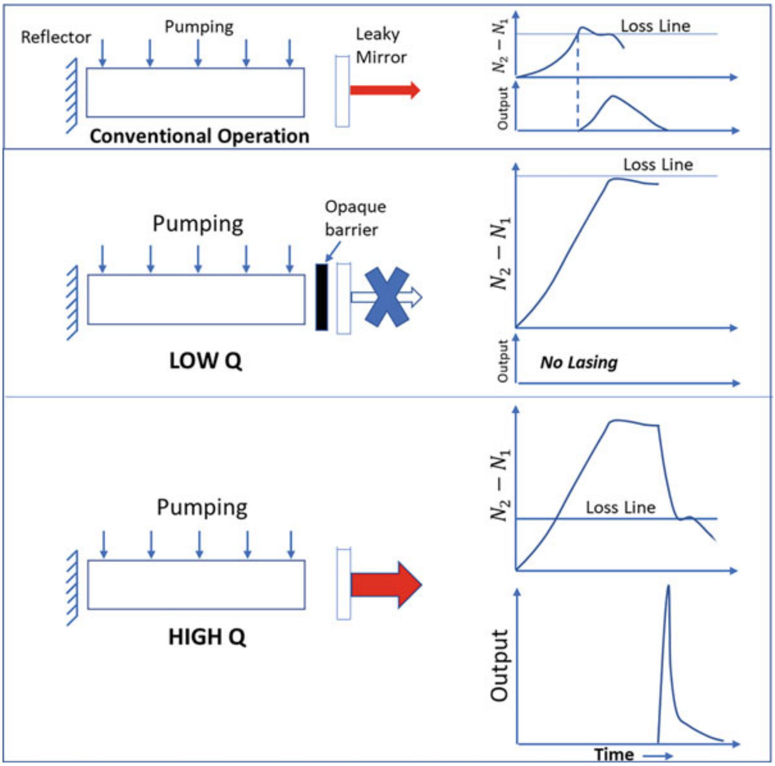

**Fig. 9.1** In the conventional operation of a laser, when the gain can barely cross the loss, the rate of stimulated emission stays low and consequently burns up the energy stored in the population rather slowly. This, in turn, limits the peak power of the pulse (top trace). If a mirror of the cavity is blocked, the loss becomes much too big making room for the gain to build up to significantly higher value. No possibility of lasing here as the gain doesn't cross the loss line (middle trace). Before the energy stored as population inversion can appreciably leak through spontaneous emission, the mirror is unblocked restoring the cavity Q, and the extremely fast rate of stimulated emission rapidly brings out the stored energy as a giant pulse with greatly increased peak power

situation has been depicted in the middle trace of this figure where the cavity has been destroyed by blocking one of its mirrors with an opaque barrier. The loss that always overrides gain here prevents the onset of lasing, and the pumped energy stays locked up as population inversion, i.e., as the internal energy of the active medium. The possibility of loss of a part of this energy over time through spontaneous emission, however, cannot be ruled out. A lossy cavity essentially means a cavity with a low Q value. Thus, we might as well say that in a cavity with low Q, it is possible to establish a gain much beyond what is possible in a conventional cavity. Once the gain or population inversion has built up to the desired value, the trick would be to restore the cavity Q to its original high value by unblocking the mirror in a flash. The lasing now begins with a value of gain far exceeding the loss, and consequently, the profoundly increased rate of stimulated emission burns up the population inversion extremely rapidly. This results in the emission of a giant pulse that is compressed in time and possesses remarkably high peak power.

Clearly, switching the cavity Q from very low to high will initiate the pulse by setting up stimulated emissions to use up the inversion and, in turn, releasing the stored energy into the cavity as light. This pulse will be terminated once the cavity becomes devoid of any light. The duration of the pulse will basically be governed by the combined time the energy takes to get out of the active medium into the cavity and eventually out of the cavity itself. The absolute value of the gain at the point of switching of Q will determine the rapidity of the stimulated emission. The higher the gain is, the quicker the population inversion burns and vice versa. For a high gain laser, the Q-switched pulse can be as short as tens of nanoseconds, while it can extend to several hundred nanoseconds in the case of a relatively low gain laser.

In a nutshell, therefore, Q-switching is a technique to extract a short pulse of remarkably high peak power from a pulsed laser by modulating its intracavity loss. When the loss is high and lasing is prohibited, the active medium should be capable of retaining the accumulated pumped energy as its internal energy. Upon lowering the cavity loss, lasing begins with an exceptionally high value of gain allowing the active medium to quickly release the vast quantity of stored energy as a giant pulse. Clearly, therefore, not all lasers are Q-switchable. Only lasers with an upper laser level lifetime long enough so that the stored energy does not leak out appreciably through spontaneous emission are conducive to Q-switching.

## *9.2.1 Q-Switching Techniques*

The Q-switch is essentially a device that should be able to switch the cavity loss from high to low quickly enough before the energy stored in the active medium leaks out substantially. Blocking and unblocking the laser cavity with an opaque barrier, considered earlier to elucidate the concept of Q-switching, may not be a very useful technique in reality; removing the barrier from the path of the oscillating light beam within the time light takes to complete one to and fro motion in a typical cavity seems quite impractical. The Q-switch devices, in operation today, can be broadly categorized into two types, active and passive. Active devices require power from an extraneous source for their operation and switching control, while a passive device is capable of operating on its own without any external help. The operation of various types of Q-switches is described in the following sections.

### 9.2.1.1 Active Q-Switching

The majority of the Q-switches that find applications are of an active nature. Some of the common active Q-switches include mechanical devices or optical modulators based on acousto-optic or electro-optic effects.

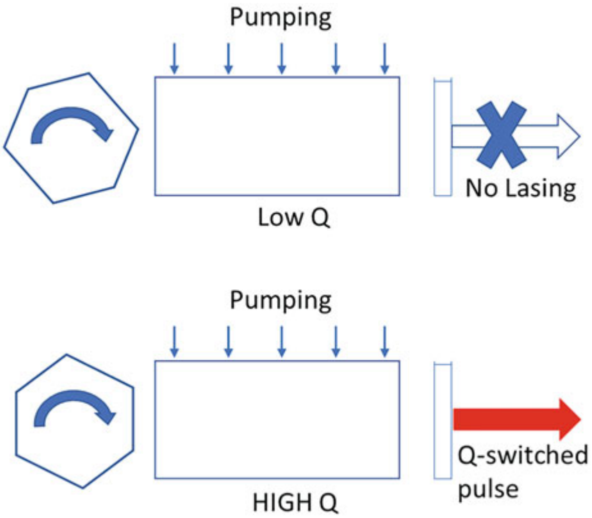

**Fig. 9.2** Operation of a rotating mirror Q-switched laser. Lasing is impossible whenever there is a cavity mismatch (low Q, top trace). Q-switched pulses emerge every time one of the reflective sides of the rotating mirror gets aligned with the other mirror (high Q, bottom trace)

### 9.2.1.1.1 Mechanical Q-Switches

In the early days of lasers, mechanical devices, such as shutters, choppers, rotating mirrors, or prisms, were used to effect the Q-switching operation. A mechanical Q-switch with a hexagonal-shaped spinning mirror is shown in Fig. 9.2. The cavity $Q$ is switched from low to high every time a surface of the rotating mirror gets aligned with the other mirror, albeit for a very short duration. This method, although robust and suitable for use with high power lasers, suffers from low switching speed and less reliability as it involves moving mechanical parts. It still finds applications, limited nevertheless, primarily due to its operational simplicity.

### 9.2.1.1.2 Acousto-Optic Q-Switches

The acousto-optic (A-O) effect, as the name suggests, essentially addresses the interaction of sound and light. It is well known that the sound wave, while traveling through a medium, creates alternating regions of compression and rarefaction within it in a periodic manner. In the optical context, these compressed and rarified regions can be identified as zones of high and low refractive index (r.i.), respectively. Thus, a sound wave passing through a medium causes periodic modulation of its r.i. along the direction of motion of the wave. This essentially renders the medium as a r.i. grating. When a beam of light passes through such a medium, this r.i. grating acts like a series of slits and diffracts light out of the main beam in multiple directions that manifest as different orders on a screen in the far field (Fig. 9.3). This means that

**Fig. 9.3** Schematic illustration of the acousto-optic effect. Passage of an acoustic wave through a medium creates a r.i. grating that behaves like a series of slits. Light upon passing through such a transparent medium gets diffracted into multiple orders

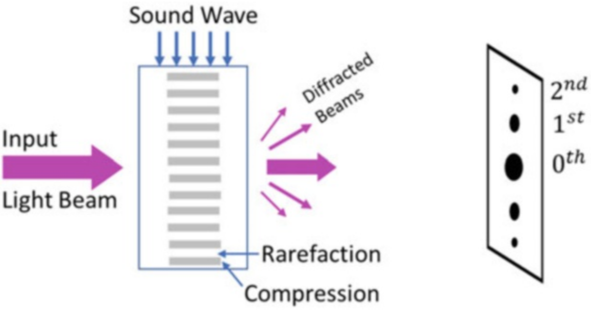

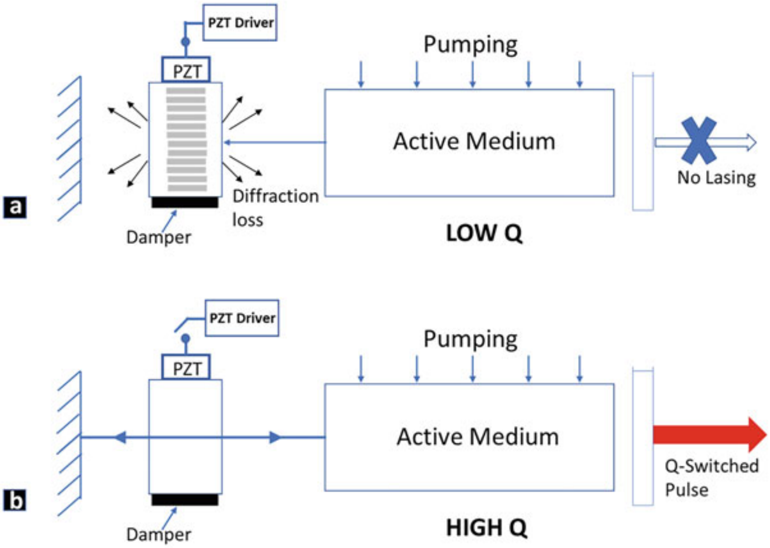

**Fig. 9.4** In order to effect Q-switching of a laser, the A-O modulator is placed between its active medium and one of the cavity mirrors. When the PZT is driven to send a sound wave into the modulator, it renders the cavity more lossy and prevents lasing (trace a). Disengagement of the PZT driver with the A-O device restores Q-value of the cavity prompting the laser to emit a giant pulse (trace b)

as the beam of light passes through such a transparent material, a part of it is spilled out as a result of the acousto-optic effect.[2]

The idea thus is to place an acousto-optic device into the laser cavity as shown in Fig. 9.4. As the gain begins to build, spontaneous emission also sets in. Application of an acoustic signal will initiate the acousto-optic effect resulting in the removal of a

---

[2] The diffraction of light into multiple orders as it interacts with an acoustic wave is called Raman-Nath effect [29]. There is yet another kind of acousto-optic modulator called the Bragg modulator [30] that diffracts light only in one direction. The Bragg modulator assumes importance when the light diffracted from the central beam constitutes the laser output as in the case of partial cavity dumping to be covered in a latter section of this chapter.

fraction of this spontaneous light from the intracavity photon flux and, in turn, lowering the cavity Q (Fig. 9.4a). As the acoustic signal is withdrawn and the acousto-optic effect disappears, the material will allow unhindered passage of light, switching the resonator Q to its original high value. It is pertinent here to discuss the feasibility of applying this concept to a practical laser system. For a visible laser, a transparent material such as quartz of cuboid geometry is often used as the acousto-optic device. The sound wave is created by applying a sinusoidal voltage to a piezoelectric transducer (PZT, introduced earlier in the preceding chapter, Section 8.7) and is launched into the quartz crystal by bonding the PZT to its one end (Fig. 9.4a). (In order to suppress the reflection of the acoustic wave from the opposite face of the crystal and rule out the possibility of its interference with the forward wave, a vibration damper is attached to its other end.) This allows gain to rise to a considerably higher value. Once the voltage applied to the PZT is withdrawn, it switches the cavity Q from low to high and the laser, in turn, gives out a giant pulse (Fig. 9.4b).

As the A-O modulator diffracts only a part of the intracavity light, it is therefore capable of introducing only a marginal loss. Not surprisingly therefore, an A-O Q-switch is not very effective in the case of a very high gain laser that by virtue of its high gain can overcome the small loss imposed by the A-O modulator and continue to lase. As will be seen later in this chapter, this disadvantage of the A-O Q-switch, however, can be exploited in its application as a partial cavity dumper. The A-O Q-switch also suffers from a low switching speed caused largely by the modest speed with which sound travels inside a material. Rapidity of the switching of Q from low to high basically depends on the time sound will take to get out of the path of the intracavity beam of light once the voltage to the PZT is withdrawn. For a typical beam diameter of ~1 cm and speed of sound in glass (~3000 m/s), the switching speed will work out to be several microseconds. Thus, the smaller the cross section of the laser beam is, the faster the switching speed.

#### 9.2.1.1.3 Electro-Optic Effect and Birefringence vis-à-vis Q-Switching

Another class of switches that takes advantage of the electro-optic (E-O) effect has also emerged as a popular technique of Q-switching. This offers fast switching at the nanosecond timescale and is also capable of switching a high gain laser. It is imperative here to gain physical insight into the electro-optic effect before we discuss its exploitation to effect the Q-switching of a laser.

Certain materials exhibit a change in refractive index upon application of an electric field, and the magnitude of change is linearly[3] proportional to the applied field. This phenomenon was discovered by Friedrich Pockels (1865–1913) in 1906

---

[3] There is yet another effect called Kerr effect or quadratic electro-optic effect where the change in refractive index is proportional to the square of the electric field. Kerr effect is, however, generally much weaker than Pockels effect.

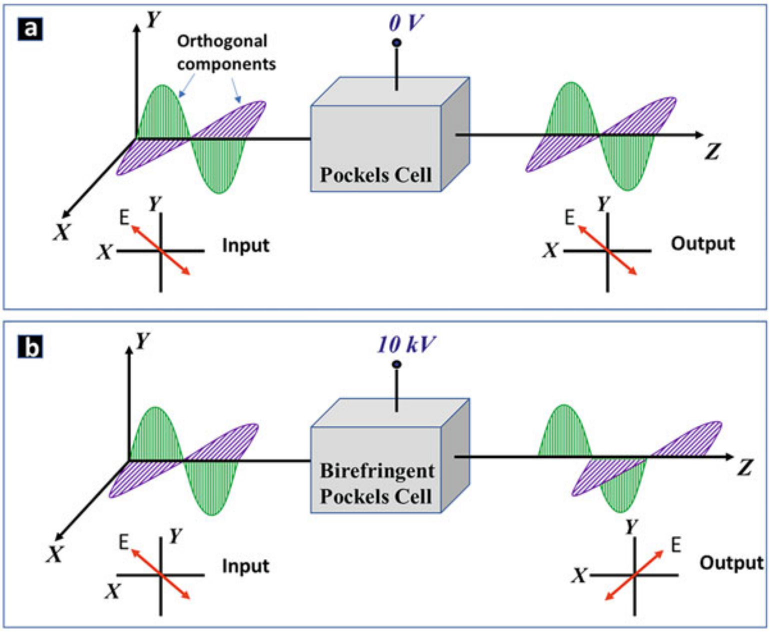

**Fig. 9.5** When the Pockels cell is unbiased, there is no birefringence and light consequently passes through with unchanged polarization (a). When the Pockels cell is biased, the two components of polarization of the incident light travel at different speeds, and the light thus emerges with a rotated polarization

and is also known as the Pockels effect. We shall focus our attention here on those materials that, in addition to displaying the Pockels effect, also exhibit the phenomenon of birefringence[4] upon application of an electric field. The amalgamation of the E-O effect with birefringence leads to a remarkable prospect of achieving a controllable rotation of the polarization of light as it passes through a suitably biased Pockels cell of appropriate length. This fact has been illustrated in the traces of Fig. 9.5 that show propagation of light, linearly polarized on the X-Y plane, through both unbiased and biased Pockels cells. Having studied the “Behavior of Light” (Chap. 2, Section 2.8), we already know that the electric field of linearly polarized light can be split into two in-phase orthogonal components. When the Pockels cell is unbiased, it behaves like an ordinary transparent material through which both components travel at the same speed. The emerging light will therefore have the same polarization as the incoming light. However, when the Pockels cell is biased, the E-O effect sets in and alters the refractive index, and if the biasing also induces birefringence, the two in-phase polarization components of the incoming light will travel at different speeds. They will no longer be in phase anymore as one will fall

---

[4] Birefringence signifies a phenomenon where the refractive index of a material becomes dependent of the polarization of light passing through it. Such materials are called birefringent.

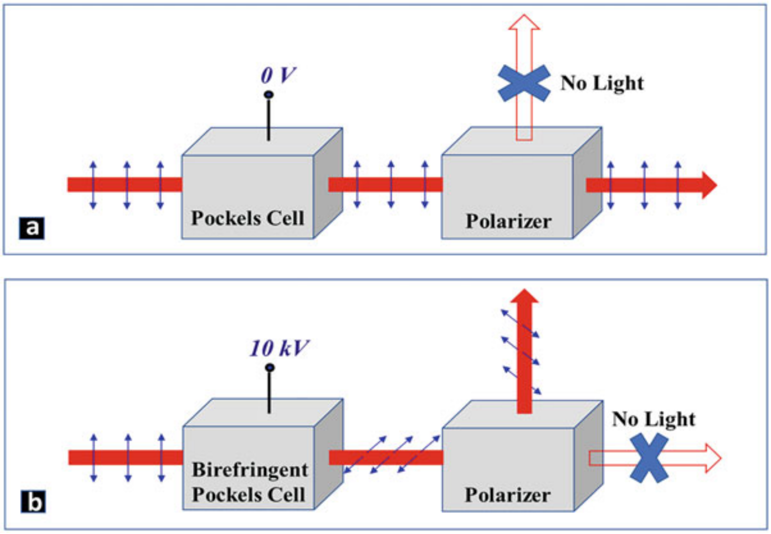

**Fig. 9.6** Light passes unhindered through the polarizer when the Pockels cell is unbiased (trace a). When a biased Pockels cell rotates the polarization of light by 90° the Polarizer blocks the light (trace b)

behind or move ahead of the other. The extent of biasing and the length of the cell can be so adjusted as to make the phase of the two components differ exactly by 180°. The vectorial addition of the electric fields of these two exactly out of phase components establishes beyond doubt that emerging light is now orthogonally polarized with respect to the light entering the Pockels cell.

### *9.2.1.1.3.1 Pockels Cell as a Light Switching Device*

The ability of the Pockels cell to rotate the polarization of light by 90° is the key to its functioning as a device to switch light when used in conjunction with a polarizer[5] as illustrated in Fig. 9.6. When the Pockels cell is not biased, light travels through it with unrotated polarization, and a polarizer, placed next, transmits this light to the other side (Fig. 9.6a). However, when the Pockels cell is biased and passes the light by rotating its polarization through 90°, the polarizer blocks this light completely (Fig. 9.6b). Clearly, a Pockels cell and a polarizer together are capable of destroying as well as restoring the Q-value of a resonator cavity and therefore, as described below, can be employed to Q-switch a laser.

---

[5] A polarizer, as we have seen earlier in Chap. 2, is an optical device that passes light on one polarization and blocks light on an orthogonal polarization.

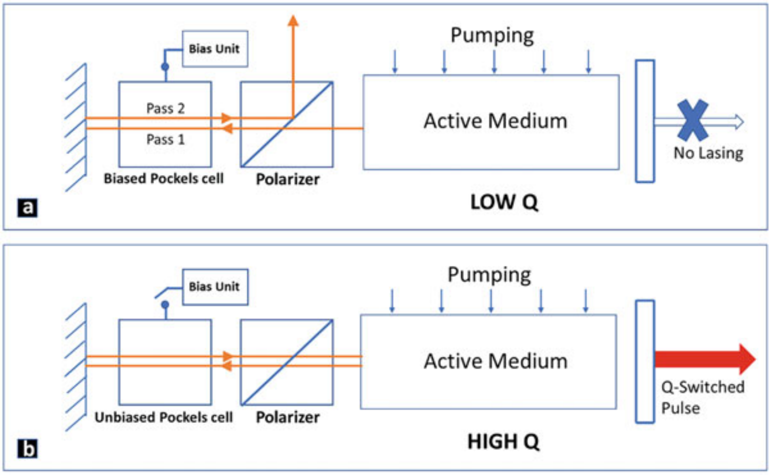

**Fig. 9.7** When the Pockels cell is appropriately biased, the polarizer rejects any light from the active medium that reaches it after two passes through the Pockels cell preventing lasing (trace a). When the voltage to the Pockels cell is withdrawn, light passes through the polarizer unhindered, and cavity Q is restored to its original high value (trace b)

### *9.2.1.1.3.2 Electro-Optic Q-Switches*

An electro-optic Q-switch that has two elements, namely, a Pockels cell and a polarizer, is placed between the active medium and one of the cavity mirrors as depicted in Fig. 9.7. Light from the gain medium passes through the Pockels cell twice before returning to the polarizer. For the polarizer to reject this returning light and prevent lasing, the Pockels cell should be so biased as to rotate the polarization by 45° in each pass (Fig. 9.7a). This allows the gain to rapidly build up until the biasing unit is switched off and the cavity Q is restored when the entire stored energy is emitted as a giant Q-switched pulse (Fig. 9.7b).

The general practice is to place the electro-optic Q-switch in front of the back mirror rather than the front mirror. With the back mirror being more reflective, a larger fraction of the light reaching here from the active medium is returned back to the polarizer boosting the performance of the Q-switch. The electro-optic Q-switch offers a very high switching speed as there is no moving part here as it happens to be in the case of a mechanical switch or a slow-moving sound wave for an acousto-optic switch. Furthermore, as the polarizer rejects the orthogonally polarized light in its entirety, this switch is capable of switching a high gain laser as well. The E-O switch is thus the most effective among all the active Q-switching devices, and no wonder it has emerged as the most popular Q-switching technique. However, electro-optic materials are not only very expensive but also require a high bias voltage for their operation. The rapid switching of such high voltages requires state of the art electronics.

### 9.2.1.2 Passive Q-Switching

The heart of passive Q-switching, also often referred to as self-switching, is a saturable absorber.

#### 9.2.1.2.1 Working of a Saturable Absorber

It is imperative to gain insight into the working of a saturable absorber before we can appreciate its role in Q-switching a laser. To this end, we perform a simple experiment, the arrangement of which has been depicted in Fig. 9.8. The goal of the experiment is basically to monitor the transmission of light through a saturable absorber as a function of its intensity. The monochromatic light is derived from a laser, preferably a *cw* one, the frequency of which ($\nu$) has an overlap with one of the absorption features ($\nu_0$) of the absorber (inset of Fig. 9.8). The intensity of light transmitted by the absorber, $I_{out}$, as recorded by the detector is shown in Fig. 9.9 as a function of the intensity of incident light, $I_{in}$. In the beginning when the incident

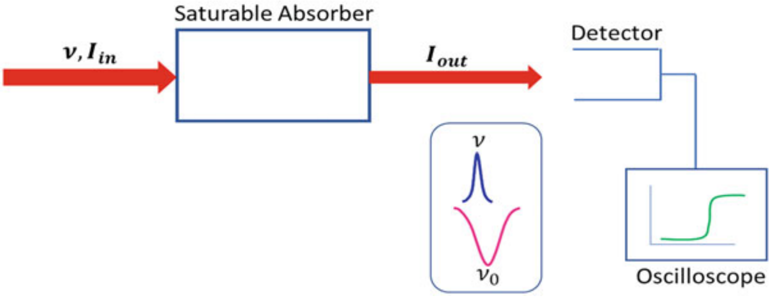

**Fig. 9.8** A typical experimental arrangement to monitor the transmission of a saturable absorber as a function of the intensity of the incident light. The inset illustrates the overlap of the frequency $\nu$ of the incident light with the absorption center $\nu_0$ of the absorber

**Fig. 9.9** Beyond a threshold value of the incident intensity, the transmission by the absorber exhibits a steep rise

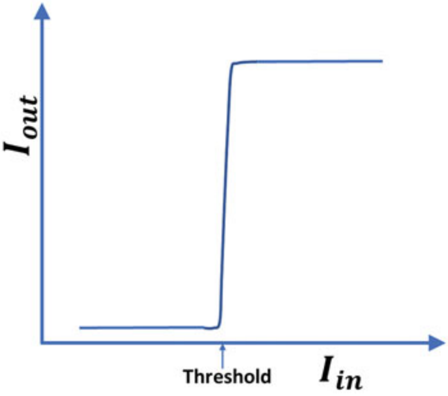

intensity is low, only a modest number of photons shine on the absorber and cannot escape absorption. With increasing intensity of the incident light, absorption also increases, and consequently the population of the corresponding excited level of the saturable absorber gradually swells. The absorber continues to prevent any leakage of the incident light to its other side until its intensity reaches a threshold value at which the excited level population equals that at the ground level. The obvious fall out of this is the saturation of absorption, and the medium's inability to remove any more photons from the incident beam of light results in a sudden rise of the output intensity beyond this point. (It is worthwhile to take another look at Fig. 4.5, Chap. 4, and the related text if you are hazy on this.)

#### 9.2.1.2.2 Q-Switching with a Saturable Absorber

This remarkable ability of a saturable absorber to block light up to a certain intensity can be exploited to Q-switch a laser, and the arrangement of the resonator cavity to achieve this is shown in Fig. 9.10. The absorber is placed between the active medium and one of the cavity mirrors. In the beginning, when pumping has just commenced, the population inversion and, in turn, the gain is low, and the spontaneous light emitted from the active medium is weak and blocked by the absorber. This prevents the onset of lasing and allows the gain to rise rapidly, and light emitted by the gain medium soon becomes intense enough to pass through the absorber with practically no loss leading to the formation of a Q-switched pulse. Ideally the recovery time of the SA should be more than the duration of the pulse but less than the time interval

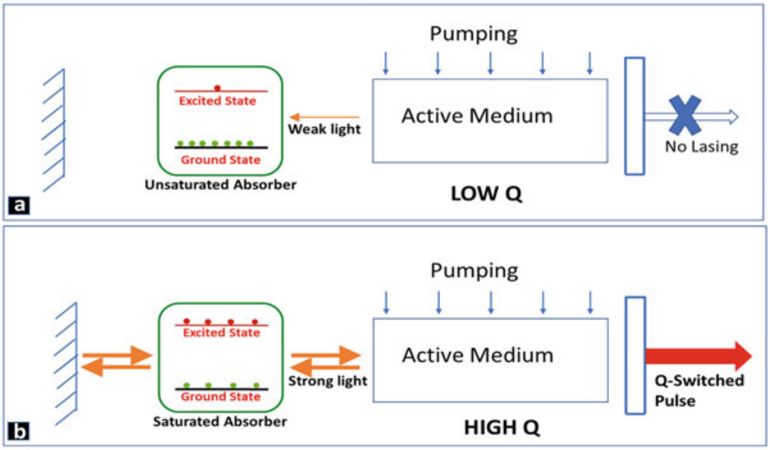

**Fig. 9.10** In the beginning when the gain is low, the weak light from the active medium is blocked by the absorber and lasing is prevented (trace a). This allows gain to build up rapidly, and consequently the stronger spontaneously emitted light from the active medium saturates the absorption and passes through the absorber almost unhindered restoring the cavity Q and, in turn, forming the Q-switched pulse (trace b)

between two successive pulses. It is also desired that a SA be able to perform by absorbing only a small fraction of the energy of the generated pulse and thus will have a negligible effect on the power efficiency of the laser. Certain special dyes are known to perform satisfactorily as passive Q-switch devices. Although degradation of the dye over time is a major drawback, they find frequent applications primarily because of their inherent simplicity and cost-effectiveness.

## 9.3 Cavity Dumping

We know now that lasers with relatively short upper-level lifetime cannot be Q-switched. This is because leakage of the stored energy through spontaneous emission prevents any appreciable buildup of the laser gain, an essential prerequisite for effecting Q-switching. An alternative approach, called cavity dumping, can be employed for such lasers to achieve pulse compression. The short upper-level life may prevent the pumped energy from being stored as internal energy of the active medium, but it does not forbid its storage as radiation energy in the cavity. The scheme takes advantage of an extremely high Q cavity, with practically no leakage, to store photons emitted through stimulation. Once the number of intracavity photons has swelled to the required level, the cavity is destroyed, and these great many photons emerge in a flash forming a giant pulse of duration not exceeding the cavity round-trip time. The principle of operation can be understood by referring to Fig. 9.11, which depicts a rudimentary illustration of the experimental arrangement to effect cavity dumping. The cavity here is formed with two fully reflective mirrors. With practically no loss, the cavity $Q$, in sharp contrast to the case with Q-switching, is initially very high. Loss being very small, lasing begins almost in synchronism with the onset of pumping, and the generated light bounces back and forth between

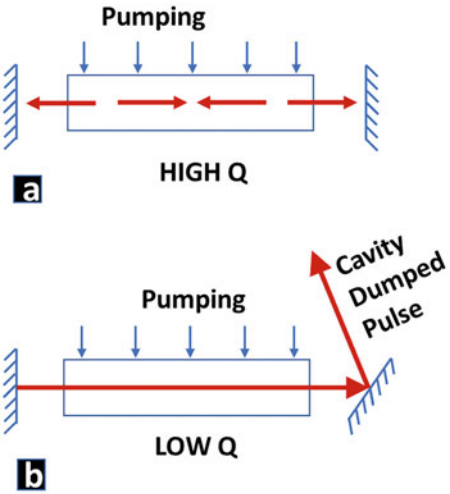

**Fig. 9.11** With the onset of pumping, stimulated emission begins to manufacture light that oscillate back and forth inside the cavity. When both the mirrors are fully reflective, the circulating light, with no escape route, rapidly builds up (a). Just at the instance of the termination of the pump pulse, a slight tilt of one of the mirrors will deflect all the light out in a flash as an extremely intense beam of light (b)

the mirrors and grows as long as the pump continues (Fig. 9.11a). The energy being dissipated by the pump is first stored briefly in the active medium as its internal energy before stimulated emission quickly realizes this as light. With no transmission loss, the intensity of light circulating inside the cavity can therefore reach an exceptionally high value by the time the pump pulse extinguishes. Destruction of the cavity at this point will result in the extraction of the entire accumulated light out of the cavity forming what can be termed as "cavity dumped pulse." Tilting of one of the cavity mirrors, as shown in Fig. 9.11b only for the purpose of comprehending the concept, will obviously reflect the beam out within a time that equals the cavity round-trip time of light. The first photon to emerge will be the one that happened to shine on the mirror just at the point of its tilting, while the last to emerge is the one reflected back into the cavity at that very instant. Obviously therefore, the duration of the cavity dumped pulse will be the time this last photon takes to complete one cavity round trip. Imparting a mechanical tilt to a mirror in a timescale commensurate with the rapidity of the pulse buildup seems quite impractical. An electro-optic switch, which we studied in-depth in the context of Q-switching, with the capability of transmitting or reflecting the light incident upon it, fits the bill quite well.

### *9.3.1 Cavity Dumping with Electro-Optic Switch*

The experimental arrangement to employ an electro-optic switch to store electromagnetic energy in a high $Q$ resonator cavity and then extract it in the form of an intense pulse by dumping the cavity is schematically illustrated in Fig. 9.12. When the Pockels cell is unbiased, the polarizer, as we already know, allows the passage of light in both directions. This makes possible to and fro lossless oscillation of light inside the cavity and, in turn, its rapid growth upon application of the pump pulse (Fig. 9.12a). In synchronism with the extinction of the pump pulse, the Pockels cell is biased, and consequently the polarizer passes light in one direction (pass 1) and reflects the return light (pass 2) out of the cavity producing the desired result (Fig. 9.12b). The cavity dumped laser, as a matter of fact, is just yet another kind of Q-switched laser where switching of cavity $Q$ occurs in the reverse order, namely, from high to low $Q$. As the energy dissipated in the active medium is stored as optical energy in the cavity, the operation of a cavity dumped laser is independent of the lifetime of the energy levels involved in the process of lasing. Thus, the applicability of cavity dumping is not only limited to lasers with short upper-level lifetime that are unsuitable for Q-switching but also capable of driving Q-switch compatible lasers. As we have seen, cavity dumping allows extraction of the entire light stored in the cavity within one cavity round-trip time, and this is independent of the time required for building up the intracavity optical flux. The requirement placed on the switch is that it should be capable of switching in a time faster than the cavity round-trip time of the laser and the electro-optic devices offering nanosecond switching speed are thus compatible for cavity dumping operation. The time required by the pump source to supply energy into the active medium basically

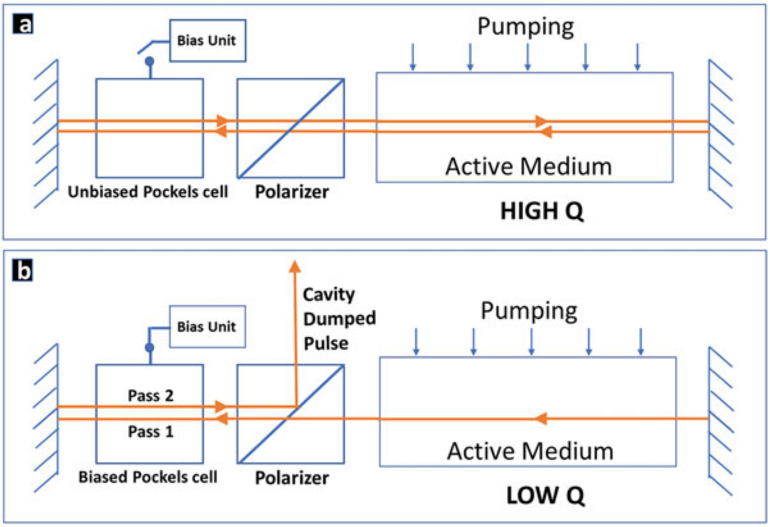

**Fig. 9.12** The operation of a laser with an electro-optic cavity dumper. An unbiased Pockels cell allows to and fro oscillation and growth of circulating light (**a**). Biasing the Pockels cell lets the entire light escape through reflection at the polarizer forming an intense cavity dumped pulse of light (**b**)

governs the achievable repetition rate as the cavity dumped pulse is extremely short-lived. In contrast, a Q-switched pulse is much too long as it is essentially governed by the slow decay rate of the high *Q* cavity. This, in turn, limits the maximum achievable repetition rate of a Q-switched laser.

### 9.3.2 Partial Cavity Dumping with Acousto-Optic Switch

At the point of extraction of a cavity dumped pulse, the cavity is required to be prohibitively lossy, a condition that is generally met when an electro-optic switch is used. An acousto-optic switch, on the other hand, is capable of introducing a moderate loss to the cavity as it only deflects a fraction of the cavity light away and is therefore not a choice for effecting cavity dumping. However, in applications where only a part of the intracavity light is to be dumped out, the acousto-optic switch comes handy as a partial cavity dumper. Of particular interest in this context is the case of ion lasers that generally have a very short-lived upper level unsuitable for the storage of energy even for a reasonable length of time prerequisite for pulsed operation. Partial cavity dumping of such lasers when pumped in the *cw* mode can produce repetitive intense pulses. The typical experimental arrangement to partially cavity dump such a laser with an acousto-optic switch is shown in Fig. 9.13. The resonator cavity is understandably formed by two fully reflective mirrors, and the active medium is pumped continuously. Stimulated emission burns the gain, and the

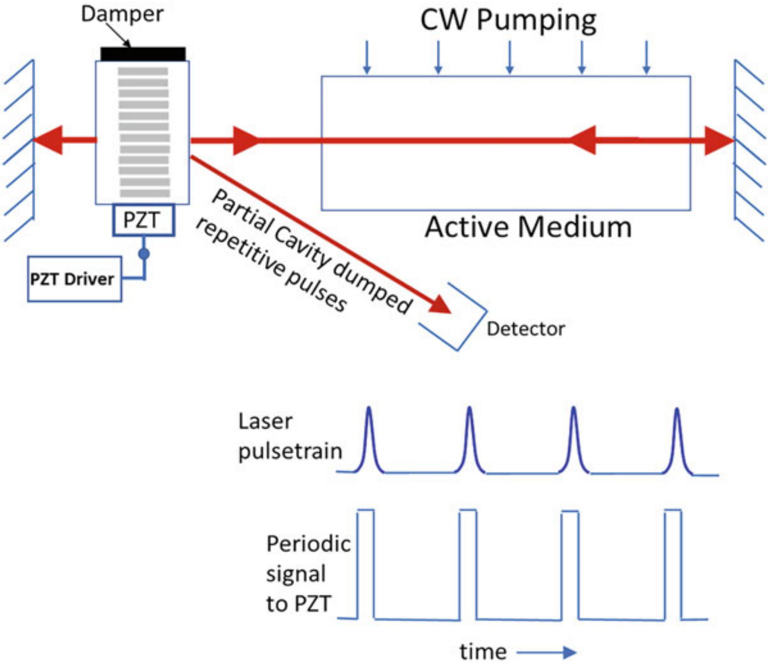

**Fig. 9.13** A typical experimental scheme to partially cavity dump a cw laser with an acousto-optic dumper. A train of signal into the modulator and the corresponding laser pulses are also shown

intensity of intracavity light builds up rapidly with seemingly no escape route. Whenever a signal is impressed onto the PZT, the ensuing acousto-optic effect deflects a part of this circulating light as an intense pulse of coherent light. As this pulse is made up of only a fraction of the intracavity light and the energy is being pumped continuously into the active medium, it may not take much too long before the intracavity light is replenished, and the laser is ready for the next dump. Thus, upon partial cavity dumping, a laser, which otherwise is not conducive for repetitive operation, becomes high repetition rate compatible. A snapshot of the repetitive pulsing of the PZT and the corresponding laser pulse train is also shown in Fig. 9.13.

## 9.4 Modelocking

The Q-switched pulse, depending upon the typical laser parameters, can vary from tens of nanoseconds to hundreds of nanoseconds. The duration of a cavity dumped pulse, on the other hand, equals the time light takes to complete a round trip of the laser cavity and thus is much shorter and can be as small as a nanosecond for a moderately long cavity. A dramatic reduction in the pulse duration by up to three orders of magnitude is, however, possible when a multimode laser operates under

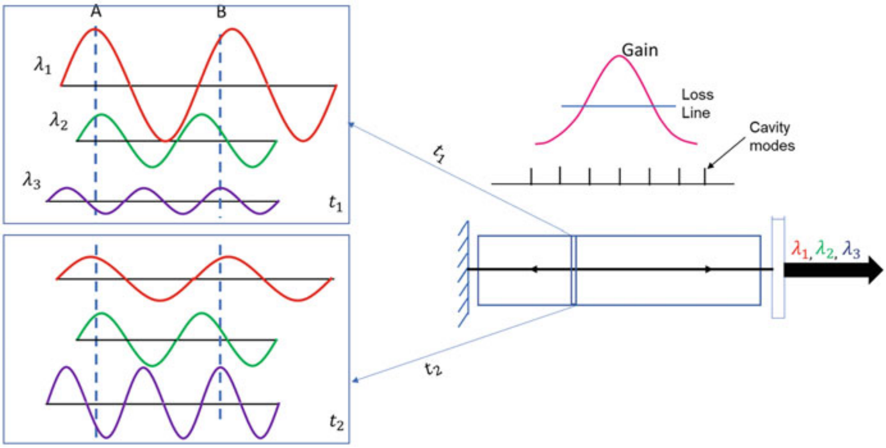

**Fig. 9.14** Lasing is possible on three longitudinal modes of wavelengths $\lambda_1$, $\lambda_2$, and $\lambda_3$ for which the gain exceeds loss. A magnified view of these modes at an arbitrary location of the active medium is shown at two different times $t_1$ and $t_2$. For clarity, the modes are shown as spatially displaced in the transverse direction. For a free running laser, the location of the interfering peak changes with time

modelocked conditions. As we shall see below, the modelocking is required to be forced upon in the operation of a multimode laser that in the free running conditions operates with modes of randomly varying phases. Toward gaining physical insight into the phenomenon of modelocking, let us consider the operation of a laser wherein three longitudinal modes of wavelength $\lambda_1$, $\lambda_2$, and $\lambda_3$ ($\lambda_1 > \lambda_2 > \lambda_3$) are able to participate in lasing (Fig. 9.14). The electric field associated with all these three modes, which share the same gain volume, will interfere with each other to produce a resulting spatial intensity pattern. If the three modes are always in phase, then the profile of the spatial intensity pattern will remain invariant in time. In reality, however, phase of the oscillating modes of a free running laser will continuously vary in time and so would be the resulting spatial intensity pattern. This will manifest as irregular pulsation in the temporal emission of a multimode laser. To bring out this point, vital in the contest of modelocking, in a more revealing manner, we magnify an infinitesimally narrow strip of the gain medium and examine the interference pattern here at two instants of time, namely, $t_1$ and $t_2$. The magnified view of the three waves is spatially displaced in the illustration of Fig. 9.14 for clarity. As seen at time $t_1$, the three waves will interfere to produce a maximum field at the location A of the strip. As the phase of the waves varies randomly, at time $t_2$, the point of the maximum interference will shift to another location *B*. If the phase of the waves were to be somehow locked, then location *A* will continue to have the highest intensity. Considering now the entire length of the gain medium, there would be one location where the combined intensity of the three modes or equivalently the number of photons present will reach a maximum. If we ignore the tiny effect of dispersion, then the three waves would move at the same speed, and the photons in this packet

**Fig. 9.15** The short modelocked pulse of light burns gain and rapidly grows as it moves along

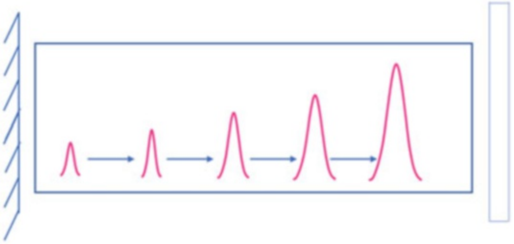

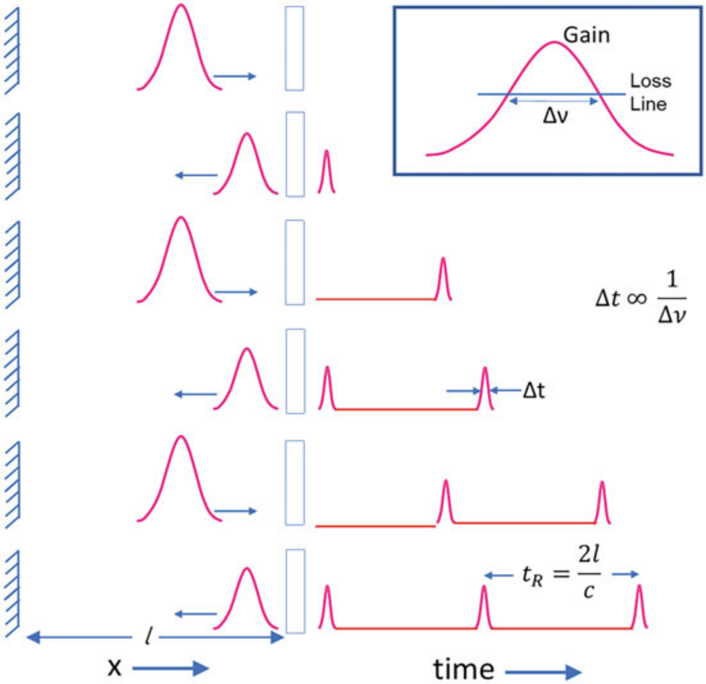

**Fig. 9.16** The to and fro intracavity oscillation of the modelocked light pulse manifests in the emission of ultrafast train of pulses by the laser. Temporal width Δt of the pulses equals the inverse of lasing bandwidth $\Delta\nu$, while the periodicity of their emission is essentially the cavity round trip time $t_R$

too will move along at this speed and rapidly grow in numbers as the light circulates within the cavity (Fig. 9.15). Needless to say, we have considered here only the photons present at the location where interference yielded the highest intensity of light as they will outperform the photons present anywhere else in burning the gain through stimulated emission. Every time this bunch of intracavity photons bounces off the output mirror, a fraction of it leaks out and constitutes the modelocked train of pulses with a periodicity of cavity round-trip time, namely, $2l/c$ (Fig. 9.16). The duration of the pulse $\Delta t$ is inversely proportional to the lasing bandwidth $\Delta\nu$. Thus, the modelocked pulses progressively shorten as an increasing number of modes participate in the lasing. It is therefore safe to say that the greater the broadening of the gain is, the higher the possibility of manufacturing shorter pulses through modelocking. For instance, while a Nd:YAG laser, with moderate broadening of gain, yields modelocked pulses of tens of picosecond duration, a Ti-sapphire laser, on the other hand, with enormous bandwidth can produce pulses well into the femtosecond regime.

It is straightforward to conclude now that while optical energy prevails across the entire length of the cavity in the Q-switching or cavity dumping cases, energy spreads only over a tiny fraction of the resonator length in the case of modelocking. A seemingly intuitive approach to effect modelocking is therefore to place a shutter into the cavity in the immediate vicinity of one of its mirrors that opens only for the desired pulse duration and remains shut for the rest of the cavity round-trip time. Although such a shutter is far from reality in a true sense, both active and passive switches, normally used to Q-switch a laser, can mimic such a shutter and facilitate the production of modelocked pulses. Like Q-switching devices, all the laser modelockers in operation today can be broadly categorized into two types, active and passive.

## *9.4.1 Active Modelocking*

Active modelocking can be achieved with both electro-optic and acousto-optic switching devices. Here, we examine their functioning as modelockers.

### 9.4.1.1 Modelocking with Electro-Optic Switches

As we know, an unbiased electro-optic switch allows the passage of light and blocks it upon biasing. Clearly if this switch has to perform the role an optical shutter plays in the operation of a modelocked laser, the application and withdrawal of biasing voltage must be in sequence as shown in Fig. 9.17. Thus, the switch opens only once very briefly per round trip, which basically compacts the optical energy into a modelocked pulse with a spatial stretch shorter than the resonator cavity itself. The only light that can oscillate back and forth without any hindrance and burn the population inversion is the light that is packed inside this modelocked pulse. As illustrated in Fig. 9.16, this circulating pulse of light, which grows from short to shorter with the participation of an increasing number of modes in lasing, produces the modelocked train of pulses emitted by the laser.

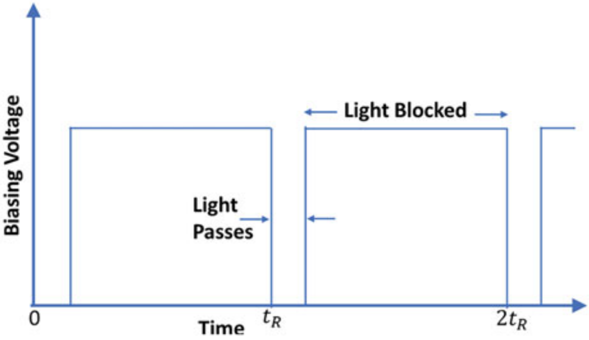

**Fig. 9.17** The sequence of biasing the Pockels cell that would make the electro-optic switch favorable to effect modelocking

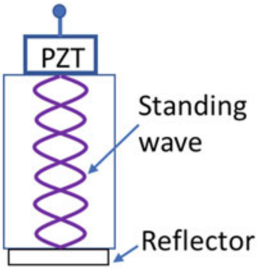

**Fig. 9.18** An acousto-optic modelocker has a sound reflector at the rear

### 9.4.1.2 Modelocking with Acousto-Optic Switches

Unlike the electro-optic switch, which can completely block a beam of light, an acousto-optic device is able to slash only a fraction of light from the incident beam. Consequently, this device is unsuitable to perform as a modelocker in its conventional operation wherein the formation of an acoustic grating diffracts away a part of light from the incident beam. The trick here is to place a reflector at the other end of the transparent medium to send the transmitted acoustic wave back into the medium (Fig. 9.18). This, as we know, is in total contrast to the conventional operation of this device wherein a vibration damper is usually made use of to get rid of this transmitted sound wave. (It may be a good idea to refresh our memory on the working of an acousto-optic device by taking another look at Fig. 9.3 and the related text.) The reflected backward wave now interferes with the forward wave to create a standing wave akin to what is produced with a plucked string of a musical instrument. To understand the role standing waves play in the working of an acousto-optic modelocker, we take a closer look at the standing wave pattern formed in a string. The fundamental standing wave that has the longest wavelength is shown in the top trace of Fig. 9.19. Clearly the shape of the stretched string changes as it continues to oscillate between the two extremities. As seen, the string, bow shaped at the two extreme locations *A* and *B*, becomes perfectly straight in midway, represented as *C*, and this momentary state of straightness occurs twice in each cycle. Thus, it is prudent to believe that the string behaves as though it is not stretched at all twice per cycle. Drawing an analogy, we may now assume that the elastic medium of the

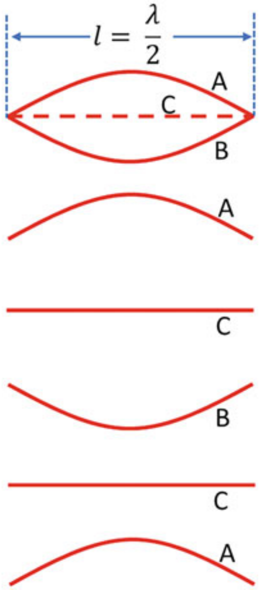

**Fig. 9.19** This picture depicts a standing wave (top) followed by the shape of the string at a few specific instances during an entire cycle

acousto-optic modulator that contains the stationary wave is also not perturbed twice per cycle. When the medium remains unaffected by the sound wave, the pulse of the modelocked light then passes through it unhindered without suffering any diffraction loss. Obviously if this device has to perform as a modelocker for a laser, then it must allow passage of a light pulse once per cavity round-trip time ($t_R$). If a signal of frequency $f$ is applied to the PZT, then the standing wave developed in the medium will allow passage of light $2f$ number of times per second. This means that the modulator does not diffract light once every $1/2f$ second. The necessary condition for the synchronization of the arrival of the modelocked pulse at the modulator just at the instance of its opening is.

$t_R = \frac{1}{2f}$

$\rightarrow \frac{2l}{c} = \frac{1}{2f}$, where $l$ and $c$ are the cavity length and speed of light, respectively.

$\rightarrow f = \frac{c}{4l}$

A typical laser with a cavity length of *1.5 m* can thus be modelocked with a standing wave acousto-optic modelocker when a signal of frequency 50 MHz is applied to its transducer.

### *9.4.2 Passive Modelocking*

This can generate pulses that are usually much shorter than what is possible with active modelocking. This technique is also more popular as the shutter performs without any extraneous intervention, has long operational life, and is not required to be serviced. A saturable absorber, the working of which is described earlier in this

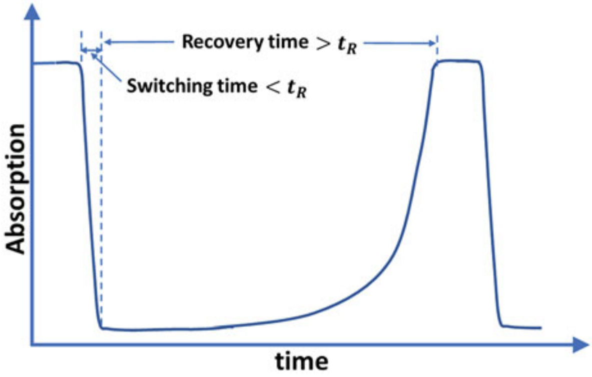

**Fig. 9.20** The recovery time of a saturable absorber Q-switch is too big compared to its switching time that is lesser than the cavity round trip time $t_R$

chapter (Sect. 9.2.1.2) and can passively Q-switch a laser, can also perform the role of a passive modelocker. To appreciate the performance of a saturable absorber as a modelocker, it is worthwhile to take a closer look at its functioning as a Q-switch, as illustrated in Fig. 9.20. The switching time (or alternatively the time for absorption to drop to a negligible value) must be smaller than the cavity round-trip time $t_R$ to facilitate the rapid buildup of the Q-switched pulse. The recovery time of the switch, i.e., the time taken for the absorption to return to its original value, needs to exceed the duration of the Q-switched pulse lest the absorptive loss eats away a part of the pulse energy. A Q-switched pulse usually exceeds the cavity round time by many multiples; this essentially means that if a saturable absorber has to perform as a Q-switch, then its recovery time must far exceed $t_R$. In sharp contrast to Q-switching, modelocking, as we know, leads to the formation of a train of temporally sharp pulses separated from one another by $t_R$ in time. Thus, for realization of modelocking, the saturable absorber must recover in a time quicker than $t_R$ and be ready to be saturated again with the arrival of the next pulse of the circulating light. As we shall see below, an appropriately chosen saturable absorber can perform remarkably well as a modelocker for both linear and ring cavity lasers when it is located correctly inside the respective resonators.

#### 9.4.2.1 Saturable Absorber: Modelocking of a Linear Cavity Laser

A schematic of saturable absorber assisted modelocking of a linear or Fabry-Perot cavity laser is shown in Fig. 9.21. The saturable absorber is placed in the immediate vicinity, often in optical contact, to one of the cavity mirrors. This is to ensure that the pulse of light can complete its to and fro journey comfortably through the absorber before its recovery. Although a saturable absorber is capable of modelocking both pulsed and *cw* lasers, we consider here only the case of *cw* lasers for ease of depiction and understanding. In order to examine the temporal evolution of the modelocked pulse, we consider the cases of both fast and slow absorbers as depicted in Fig. 9.22. A fast absorber can recover sufficiently quickly and almost as

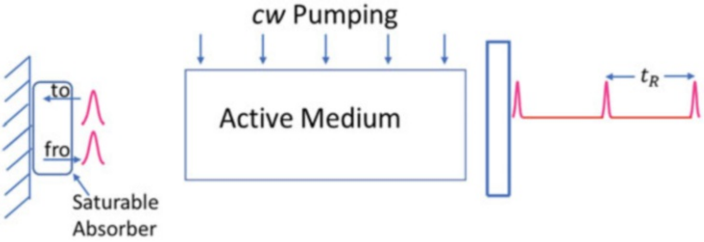

**Fig. 9.21** In a linear cavity, the saturable absorber is placed close to one of the cavity mirrors, often the rear one, to ensure unhindered to and fro passage of the modelocked pulse through it

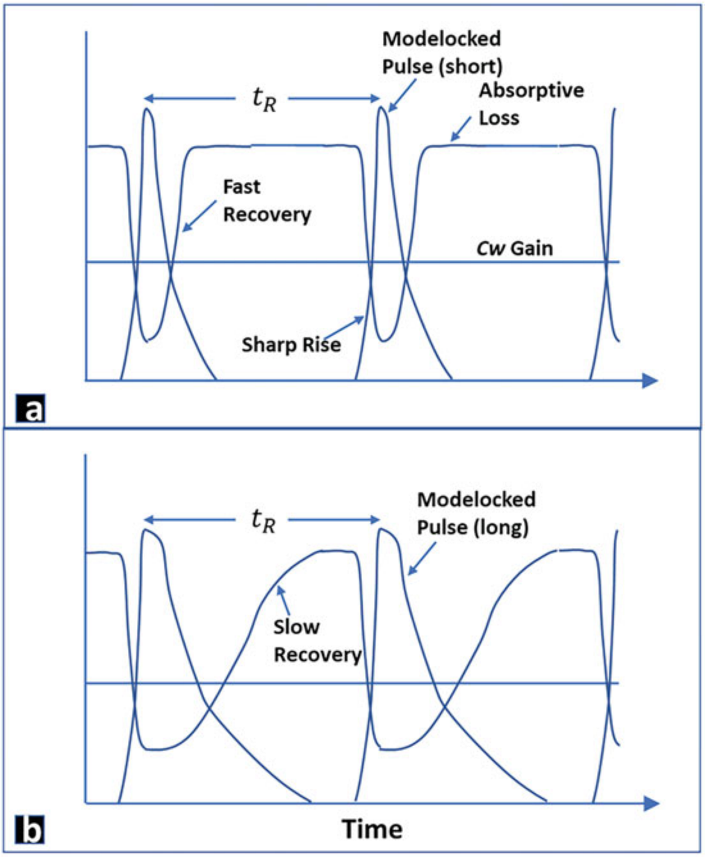

**Fig. 9.22** The loss, gain, and power of optical pulse as a function of time for a fast (**a**) and slow (**b**) saturable absorber. The faster modulation of loss offered by a fast absorber makes the pulse shorter

fast as it takes to break into transmission, while the recovery of a slow absorber can extend almost up to the cavity round-trip time. The leading edge of the circulating light understandably suffers absorption in both cases which manifests in sharpening the pulse rise time. In the case of a fast absorber, the trailing edge of the pulse is also attenuated as it encounters appreciable absorption (Fig. 9.22a). In the case of a slow absorber, on the other hand, the trailing edge experiences much less absorption and consequently exhibits a gentler descent (Fig. 9.22b) and, in turn, tends to extend the pulse. The gain remains practically unaffected as the energy content of the modelocking pulses is too modest to burn it appreciably. In a nutshell, the fast absorber, by virtue of its ability to impose faster modulation of loss, tends to shorten the modelocked pulse. The decided advantage of colliding pulse modelocking, which is intrinsic in a ring laser and will be described below, essentially stems from this fact.

#### 9.4.2.2 Saturable Absorber: Colliding Pulse Modelocking of a Ring Cavity Laser

One major hurdle encountered in modelocking a linear (also known as Fabry-Perot) cavity laser with a saturable absorber, in particular the fast one, is the problem in fabricating the absorber cell in optical contact with the cavity mirror. As we shall see now, the usage of a ring cavity readily overcomes this difficulty as it places no requirement of locating the absorber cell in the immediate proximity of cavity optics. The configuration of a saturable absorber assisted modelocked ring cavity laser is shown in Fig. 9.23. The ring cavity of length *4 L* comprises four mirrors with each of its arms being of length *L*. Although both the gain and absorption cells are shown here to be placed in one of the arms, the operation is possible by locating them

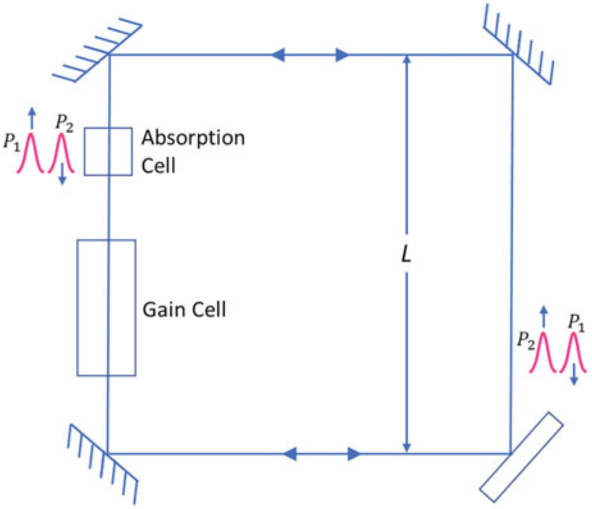

**Fig. 9.23** A four mirror ring cavity containing the gain and absorption cells in one of its arms. The co- and counterpropagating pulses of the circulating light will always cross inside the absorber

**Fig. 9.24** When the gain cell is spaced by quarter of the cavity length from the absorber, the pulses arrive alternately in the gain cell once per half of the cavity round trip time

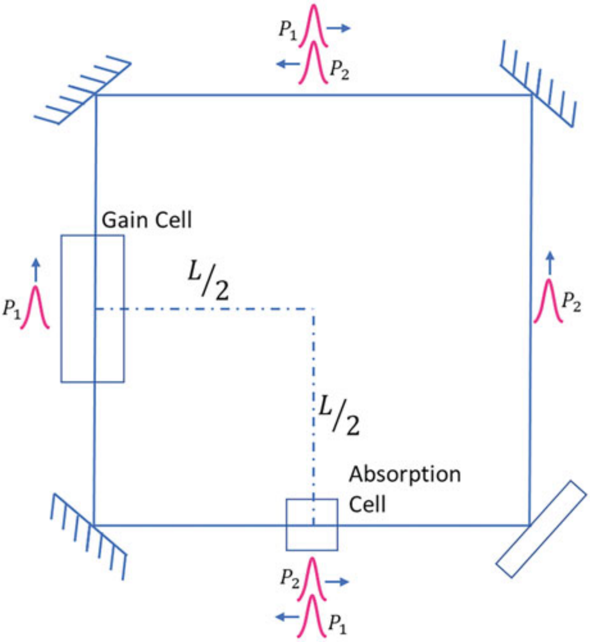

anywhere in the cavity. As we know in a ring cavity, two light pulses propagate in clockwise and counterclockwise directions. The saturable absorber parameters are so optimized that it would saturate only during the time the two pulses make an overlap inside the absorber. The rest of the nonoverlapping light of both pulses will be blocked by the absorber. The enhancement of light intensity at the instant of overlap causes faster modulation of loss that manifests in the creation of two ultrafast pulses $P_1$ and $P_2$ moving in opposite direction in a manner that they always cross inside the absorber once per round trip.[6] The synchronized arrival of the pulses in the absorber is automatic here and is independent of the location of the absorber with regard to the gain cell. However, the appearance of these pulses in the gain cell depends on its location with respect to the absorber. For instance, when the gain cell is located very close to the absorber as shown in Fig. 9.23, $P_2$ clearly arrives at the gain cell almost immediately upon the passage of $P_1$ through it. $P_1$, on the other hand, always reaches the gain cell much longer after the passage of $P_2$ through it. This temporal asymmetry in the time the two pulses take in reaching the gain cell after their respective passage through the absorber can be overcome upon locating the absorber a quarter of the cavity length away from the gain cell, and such a situation is depicted in Fig. 9.24. As seen from this figure, the two pulses now arrive at the gain medium

---

[6] The interference of the two counterpropagating waves, although not explicitly mentioned here but inevitable in a bidirectional ring cavity, results in manyfold increase of light intensity at the nodal plane. This, in turn, results in even faster modulation of loss yielding better synchronization, higher stability, and shorter pulses.

with the same time delay corresponding essentially to half of the cavity round-trip time. This allows the gain to always recover to the identical condition before the arrival of a pulse in it.

## 9.5 Chirped Pulse Amplification vis-à-vis Manufacturing Extreme Light

Exploitation of the phenomenon of modelocking, as we know now, has a decided advantage as far as generation of an ultrashort pulse is concerned. Furthermore, as the two successive pulses are temporally separated by the cavity round-trip time, a modelocked laser can thus give out pulses at an epically high repetition rate. This, however, comes at the expense of the energy contained in a pulse that essentially is much too modest compared to what is possible to be extracted from a Q-switched or cavity dumped laser. To drive this point home, we consider a typical modelocked laser with a cavity length of *1.5 m* delivering 100 watt average output power in the form of a train of 100 *fs* pulses. With the cavity round-trip time of this laser being 10 *ns*, it gives out one pulse every 10 ns or alternately 100 million pulses in a second. The 100 joule of energy that this laser emits in 1 second will therefore be carried by all 100 million modelocked pulses. The energy share of a lone pulse is thus just a mere one millionth of a joule, a ridiculously tiny magnitude to serve any meaningful purpose. Fortunately, the story does not end here though. True, a millionth of a joule is indeed too small, but the fact that the pulse is compacted in a time that lasts over just a tenth of a millionth of a millionth second[7] makes its peak power swell to an incredible magnitude of ten million watts. It is no wonder that these ultrashort pulse trains, although energetically insignificant, are powerful enough to fire the imagination of both scientists and technocrats across the globe. This has led to the proliferation of numerous applications, over the years, distinct enough to profoundly impact our everyday life and will form an important part of volume II of this book. Furthermore, the built-in focusability of these superpowerful ultrashort pulses to a tiny spot allows ablation of even the hardest material surface with unprecedented precision and control. This has led to the realization of microsurgeries on extremely sensitive and vulnerable parts of the human body that are unconceivable by a conventional tool or heat source.

There is yet another rapidly emerging area of research aimed at studying matter at extreme physical conditions such as that which exists at the core of a star. Simulation of such a condition in the confines of a laboratory is possible by focusing a light pulse of power exceeding a trillion watts to a tiny spot. This means that the million-watt power of the so-called ultrafast pulses must be scaled up by a factor of at least another million before they qualify to recreate such extreme conditions. Raising the power of this pulse even to a modest level by passing it through a light amplifier

---

[7] Please note that 100 *fs* or $100 \times 10^{-15}$ second is equivalent to $10^{-1} \times 10^{-6} \times 10^{-6}$ second.

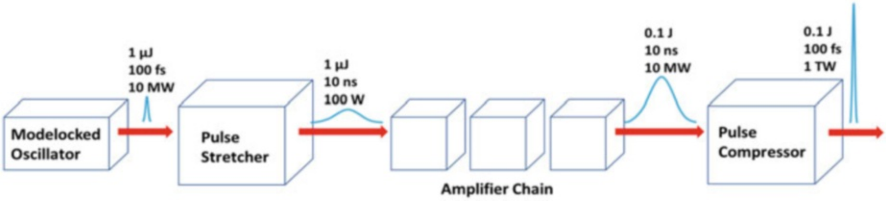

**Fig. 9.25** Block diagram representation of a CPA cycle. The performance of the technique is exemplified here for a modelocked pulse of 1 μJ energy and 100 fs duration. The energy and duration of the modelocked pulse, however, can vary as long as the amplified power does not set off the detrimental nonlinear optical effect inside the amplifier

poses unsurmountable difficulties, let alone a millionfold boost. As the intensity of the pulse begins to grow inside the amplifier, unwanted nonlinear optical effects such as self-focusing[8] set in, impairing both the amplifying medium and the optics. It is worth mentioning in this context of another approach to develop a massive laser facility comprising a large number of independent oscillator-amplifier systems each capable of delivering a KJ of energy in pulses of nanosecond duration. The multiple outputs when combined into one raise the optical power beyond the terawatt level and are central to inertial confinement of fusion (ICF) research program. Notwithstanding the recent advances made toward ICF at the National Ignition Facility located at Livermore, California, this approach of raising power by building ever larger lasers did not gain popularity largely owing to their gigantic foot print and exorbitantly high cost. The seemingly insurmountable difficulty in realizing the passage of a femtosecond pulse through a light amplifier has eventually found a perfect solution in a magical concept now popular as chirped pulse amplification (CPA). The sequence of events that together represent the CPA is represented as a block diagram in Fig. 9.25. As seen, the femtosecond modelocked pulses are first stretched[9] by a factor that can be as high as 100,000. This operation understandably brings down the power of the pulse by the same factor. The standard light amplification technique can then be employed to raise its power. Clearly the power can be boosted by the same factor by which the pulse was initially stretched. Finally, a pulse compressor squeezes the pulse back to its original duration and, in turn, raises the power a hundred thousand times beyond the amplifier's capability. A typical

---

[8] Certain nonlinear materials, with intensity dependent r.i., behave like a lens and progressively focuses an intense beam of light with nonuniform spatial intensity profile, such as a Gaussian beam, passing through it. Self- focusing is a nonlinear optical phenomenon and will be described in detail in V-II of this book.

[9] A pulse stretcher basically is a dispersive element like a prism or grating that spreads the constituent colors of the femtosecond pulse spatially. Once the color components are spatially separated, they can be made to travel different distances and thus stretched in time. A compressor is again a dispersive element that negates the effect of stretcher and recombines the individual colors both in time and space. The literature is quite rich on CPA, and the interested readers may refer to the literature for developing a thorough understanding on the working of a pulse stretcher or compressor.

example will make this easier to understand. We once again begin with a seed pulse of a microjoule energy lasting over 100 femtoseconds. Upon stretching it by a factor of 100,000, the power of the pulse reduces from 10 MW to 100 W. Amplification of the pulse now by the same factor of $10^5$ raises its energy from 1 μJ to $10^{-1}$ J, and finally compressing the pulse back to 100 femtoseconds will make its power swell to the terawatt level. Generally speaking, the light, with regard to its purity and intensity, has, to date, undergone two revolutions; first is of course the birth of laser in 1960, and the second is none other than the emergence of the CPA technique in 1985 to manufacture light of TW power or more from a table top laser. Indeed, after its birth in 1960, the implementation of CPA technique in 1985 brought about a youthful exuberance to the laser at 25, and, incidentally, even today when it is 60 plus, the exuberance is enduring.