Нелинейности 3-го порядка

$\chi^{(3)}$ Processes.

Now we will begin to consider χ^3 processes. Let's return to the induced dipole polarisation equation:

$$\underline{P} = \varepsilon_0 (\chi^{(1)}\mathbf{E} + \chi^{(2)}\mathbf{E}^2 + \chi^{(3)}\mathbf{E}^3 + \ldots\ldots)$$

You can see we have a set of terms in: $\chi^{(3)}\mathbf{E}^3$, we can now come back to our substitution of $E = E_0\cos(\omega t)$ to obtain our $\chi^{(3)}$ terms:

$$E_{\chi 3} \propto E_0{}^3\chi^{(3)}\cos^3(\omega t)$$

Generating terms: $\frac{\chi^{(3)}}{4} E_0^3\left(3\cos(\omega t) + \cos(3\omega t)\right)$

Refractive index.

Let's start by considering linear refractive index.

Remembering that we can write:

$\mathbf{D}(\omega) = \varepsilon_0\mathbf{E}(\omega) + \mathbf{P}(\omega)$ and $\mathbf{P}(\omega) = \varepsilon_0\chi^{(1)}\mathbf{E}(\omega)$

So $\mathbf{D}(\omega) = \varepsilon_0(1+\chi^{(1)})\,\mathbf{E}(\omega) = \varepsilon_0\varepsilon_r\mathbf{E}(\omega)$

We can define the refractive index $n = \sqrt{\varepsilon_r} = \sqrt{1+\chi^{(1)}}$

[In general χ can be a complex function of frequency, so we should really write: $n(\omega) = \sqrt{1+\mathrm{Re}\left(\chi^{(1)}(\omega)\right)}$]

The refractive index then is independent of the amplitude or intensity of the light (but does vary with ω.)

Non-linear Refractive Index

Now consider a centro-symmetric medium (ie $\chi^{(2)}=0$)

$$P(\omega) = \varepsilon_0\chi^{(1)}\mathbf{E}(\omega) + \varepsilon_0\chi^{(3)}\mathbf{E}^3(\omega)$$
$$=\varepsilon_0\mathbf{E}(\omega)\,(\chi^{(1)} + \chi^{(3)}\mathbf{E}^2(\omega))$$

$$\text{So: } \mathbf{D}(\omega) = \varepsilon_0\mathbf{E}(\omega) + \varepsilon_0\mathbf{E}(\omega)\,(\chi^{(1)} + \chi^{(3)}\mathbf{E}^2(\omega))$$
$$=\varepsilon_0\mathbf{E}(\omega)\,(1 + \chi^{(1)} + \chi^{(3)}\mathbf{E}^2(\omega))$$

Note that we can write the intensity, $I(\omega) = \mathbf{E}^2(\omega)$

So we now have: $\varepsilon_r = (1 + \chi^{(1)} + \chi^{(3)}I(\omega))$

Therefore: $n = \sqrt{(1 + \chi^{(1)} + \chi^{(3)} I(\omega))}$

Allowing us to approximate: $n \approx n_0 + \frac{1}{2n_0}\chi^{(3)} I(\omega)$

Where n_0 is the linear refractive index = $\sqrt{1+\chi^{(1)}}$

Now we can write: $n_2 = \frac{1}{2n_0}\chi^{(3)}$

n_2 is known as the ***non-linear refractive index*** for the material.

So we now have a new expression for the refractive index that is a combination of linear and non-linear components:

$$n(I) = n_0 + n_2 I(\omega)$$

What does this mean in practice? The refractive index now changes with beam intensity giving rise to a phenomenon known as ***self-focusing***. As the intensity of a laser beam is not generally uniform, different parts of the beam see different refractive indicies generating a lens like behaviour within the medium. Allowed to go unchecked this can lead to major optical damage due to the continuous increase in central intensity arising from this phenomenon.

The effect is also known as the **Optical Kerr Effect**. Where the kerr effect is a change in a material's refractive index in response to an electrical field.

Self Phase Modulation – SPM

A related non-linear effect to self focussing is self phase modulation. A change in refractive index changes the velocity of the pulse in the medium giving rise to instantaneous phase change in the pulse itself.

As the pulse begins to move through in the higher index environment, so after travelling a distance L, the instantaneous value of the phase will change, $\Delta\phi$:

$$\Delta\phi(t) = \int_0^L \frac{\omega_0}{c}\Delta n(t)dz = \frac{\omega_0}{c}Ln_2I(t) = \frac{2\pi}{\lambda_0}n_2I(t)L$$

So after propagating a distance L:

$$\phi_L = \phi_0 + \Delta\phi$$

$$So:$$

$$\phi_L = \omega_o t + \frac{2\pi}{\lambda_0} n_2 I(t) L$$

We can define instantaneous frequency, $\omega(t)$: $\omega(t) = \frac{d\phi(t)}{dt}$

So we see a frequency shift: $\Delta\omega = \omega(t) - \omega_0$

So we can see that: $\Delta\omega$ = - const x dI(t)/dt

The negative sign arises from the fast increase in phase resulting in a fast increase in optical path length of the medium and according to the optical Doppler effect, the frequency will become lower.

In practice this effect acts to broaden the range of frequencies contained within the pulse. The front edge of the pulse (gradient of I(t) +ve) has its frequencies decreased (red shifted) whilst the trailing edge of the pulse (gradient of I(t) –ve) has its frequencies lengthened (blue shifted.) This change in frequencies is often known as ***CHIRP***. Taking the Fourier transform gives a characteristic broadening of the output spectrum of the pulse.

So a pulse of light that travels through a $\chi^{(3)}$ medium undergoes self focussing (the optical Kerr effect) and has the range of frequencies broadened (through self phase modulation.) Often these are unwanted consequences – eg. in optical telecommunications, SPM can give unwanted dispersion. As discussed above, self-focussing can cause catastrophic damage to non-linear elements. However in the right environment, the blend of SPM and self-focussing can give some spectacular effects.

The Kerr Lens Modlocked Ti:Sapphire Laser

In 1989, a laser was demonstrated in St Andrews that appeared to generate fs pulses spontaneously. This was entirely unexpected and has paved the way for a whole range of practical fs laser applications. Initially, the process by which the pulses were generated was unknown and was known as ‘magic modelocking”, however with SPM and self-focussing, we can explain the behaviour.

A 'Typical' Ti:Sapphire Laser

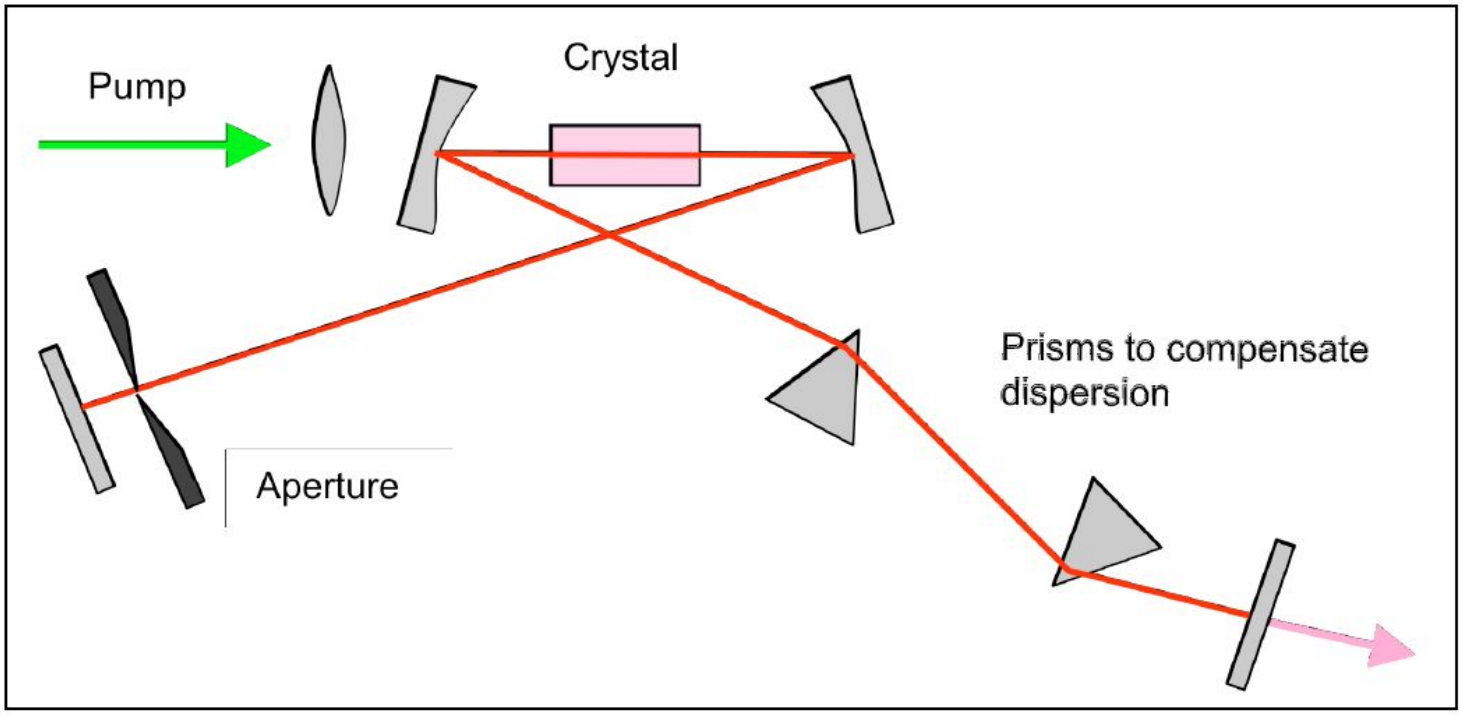

The key to operation is the aperture. This is set up so that when no pulses are formed, it slightly clips the laser mode. This is a loss mechanism. When the laser operates pulsed however, self focusing takes place within the crystal. This then slightly focuses the beam emerging from the end of the rod and the aperture is adjusted, so that the focused beam does not touch it. This gives a lower round trip loss to laser pulses (their higher intensity causes the self focus.)

Self phase modulation plays a vital role due to the bandwidth requirement of ultrashort pulses. Whilst the CW linewidth of the laser can be very small (~MHz), the linewidth of an ultrashort pulse can be much bigger. For example a 100fs pulse at 800nm needs ~12nm of optical bandwitdth whilst the shortest of pulses (~5fs) need many 100's nm bandwidth.

The laser pulse is initiated from CW operation by introducing a noise spike into the laser. This noise spike induces a small non-linear focusing effect, generating a lower loss and undergoing some SPM broadening. As lasers tend to run where losses are minimised, this pulsed operation becomes favoured and the effect continues until a balance is reached between available gain bandwidth and the effects of dispersion that can be corrected using a prism.

This operating arrangement is known as ‘hard aperture’ modelocking. A further operation regime ‘soft aperture’ modelocking is also observed when the aperture is removed from the cavity. In this mode, the self focussing causes and enhanced overlap between the laser and pump mode providing a higher gain for pulsed operation. This can however be more complicated to set up.

Numerical example

n_2 Ti:Sapphire at 800nm ~$3 \times 10^{-16} cm^2/W$

n_0 Ti:Sapphire at 800nm 1.76

Calculate the intensity: $n_2 I = n_0$?

$I = n_0 / n_2 = 6 \times 10^{15} Wcm^{-2}$

Typical peak intensity of a Ti:Sapphire pulse focused to 10μm?

P=1W, Rep rate = 80MHz, τ = 100fs

Peak power = $1/(80 \times 10^6 \times 100 \times 10^{-15}) = 125kW$

Focal area ~ $\pi r^2 = 3 \times 10^{-8} cm^2$

Intensity = $4 \times 10^{12} W/cm^{-2}$

Non-linear refractive index modulation = $3 \times 10^{-16} \times 4 \times 10^{12} = 1.2 \times 10^{-3}$!

This is small, but significant – take home message – non-linear effects need very high intensities!

3.1 Laser Spiking and relaxation oscillations

Two types of transient behaviour are frequently observed in the output of laser systems:

1. **Laser spiking:** This can occur when a laser is first turned on, and comprises large amplitude, somewhat irregularly-spaced, pulses of laser radiation.

2. **Relaxation oscillations:** These are more regular oscillations in the output power of the laser, typically with an exponentially-decaying amplitude. Relaxation oscillations can occur in the tail of a series of laser spikes, or can follow a small perturbation to the laser system such as disturbance of a cavity mirror, or a change in the pumping power.

The two phenomena are related; both arise from the fact that the population inversion density and the intra-cavity laser intensity are coupled together via stimulated emission and feedback from the cavity mirrors, and that the inversion density and intra-cavity intensity respond to changes in the operating conditions of the laser on different time scales.

Consider first the case of laser spiking. Immediately after the pumping is switched on, the population in the upper laser level will grow from zero in a time of order the upper level lifetime τ_2. At this point the density of photons in the cavity will be very low since the photons arise from spontaneous emission only. At some point the population inversion density will reach the threshold for laser oscillation, $N_{\rm th}^*$, but laser oscillation will not commence since the density of photons in the relevant cavity mode will still be low. As pumping continues, the population inversion grows beyond the threshold value, and the density of photons in the laser mode will build up by stimulated emission. Eventually the density of photons in the laser mode will be sufficiently high to burn the population inversion down towards $N_{\rm th}^*$, further increasing as it does the rate of stimulated emission and the photon density. Inevitably, the population inversion will be driven well below the threshold value, leading to reduction in the photon density and the cessation of lasing. The dramatic growth and rapid decline of the intra-cavity photon density correspond to a spike in the output power of the laser. As pumping continues, the population inversion will start to grow again and the cycle can repeat. Laser spiking of this type will not usually continue indefinitely, even in continuously-pumped lasers, owing to the fact that the population inversion and photon densities do not decay to zero after the spike. Consequently, after each spike these quantities tend to end up closer to their steady-state values, in which case the spiking will decay. Figure 3.1 shows a typical example of laser spiking from the first ever ruby laser.

It should be noted that laser spiking is often associated with rapid jumps in the longitudinal and/or transverse mode of the laser. This occurs for essentially the same reasons as spatial hole-burning in lasers operating on homogeneously broadened transitions: immediately after a laser spike, significant population inversion may remain in regions near nodes of the longitudinal nodes, or in regions where the transverse profile of the mode of the previous pulse was low in intensity. This gain can be exploited by other modes which, as pumping continues, may therefore reach threshold sooner than the mode (or modes) of the previous spike.

Laser spiking tends to occur when the time scale on which the population inversion responds to the pumping, typically of order τ_2, is slow compared to that in which the photon density changes in the cavity, which is of order the cavity lifetime $\tau_{\rm c}$. Hence we expect laser spiking to occur if,

$$\tau_2 \gg \tau_{\rm c} \quad \text{Approximate condition for laser spiking.} \tag{3.1}$$

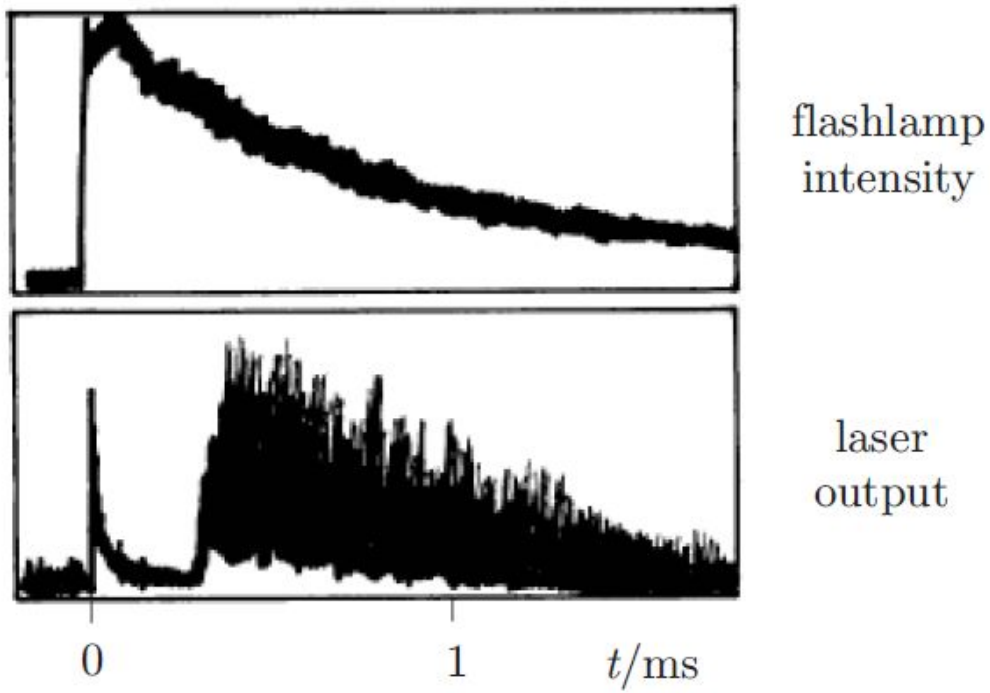

Figure 3.1: Laser spiking in the first ruby laser. The top trace shows the flashlamp pulse, the lower the output from the ruby laser as recorded by a photodiode. Note that the sharp feature at $t = 0$ is caused by flashlamp light reaching the photodiode.

3.1.1 Rate equations

For a homogeneously broadened laser the rate equations for the upper and lower laser level may be written simply as[1],

$$\frac{dN_2}{dt} = R_2(t) - N^*(t)\sigma_{21}\frac{I(t)}{\hbar\omega_L} - \frac{N_2(t)}{\tau_2} \tag{3.2}$$

$$\frac{dN_1}{dt} = R_1(t) + N^*(t)\sigma_{21}\frac{I(t)}{\hbar\omega_L} - \frac{N_1(t)}{\tau_1} + N_2(t)A_{21}, \tag{3.3}$$

where R_i is the pump rate of level i, and ω_L and $I(t)$ are the angular frequency and total intensity of the oscillating laser mode.

Notice that in the absence of stimulated emission ($I(t) = 0$) the population of the upper laser level grows in a time of order τ_2.

The rate per unit volume at which photons are stimulated into the oscillating mode is simply $N^*\sigma_{21}I/\hbar\omega_L$. Hence if we suppose that the fraction of the volume of the oscillating mode lying within the gain medium is equal to f_c, the rate equation for the density of photons is simply,

$$\frac{dn}{dt} = f_c N^*(t)\sigma_{21}\frac{I(t)}{\hbar\omega_L} - \frac{n(t)}{\tau_c}. \tag{3.4}$$

The last term gives, in terms of the **cavity lifetime** τ_c,the rate of loss of photon density by coupling through the output mirror (plus any other losses). We derive an expression for τ_c in Section 3.1.2.

This equation may be simplified by noting that $I(t) = n(t)\hbar\omega_L c$. Then,

$$\frac{dn}{dt} = \left(\frac{N^*}{N^*_{th}} - 1\right)\frac{n}{\tau_c}, \tag{3.5}$$

where,

$$N^*_{th} = \frac{1}{f_c\sigma_{21}c\tau_c} \tag{3.6}$$

and we have dropped the time dependence of $N^*(t)$ and $n(t)$ to avoid clutter. We see that $N^*_{\rm th}$ is the threshold population inversion density since we must have $N^* > N^*_{\rm th}$ for the photon density to grow (by stimulated emission). Equation (3.5) also shows that the time scale in which the photon density changes is $\tau_{\rm c}$ divided by a factor which depends on the extent to which the population inversion exceeds $N^*_{\rm th}$. As such we expect the photon density to change in a time of order $\tau_{\rm c}$, as stated earlier.

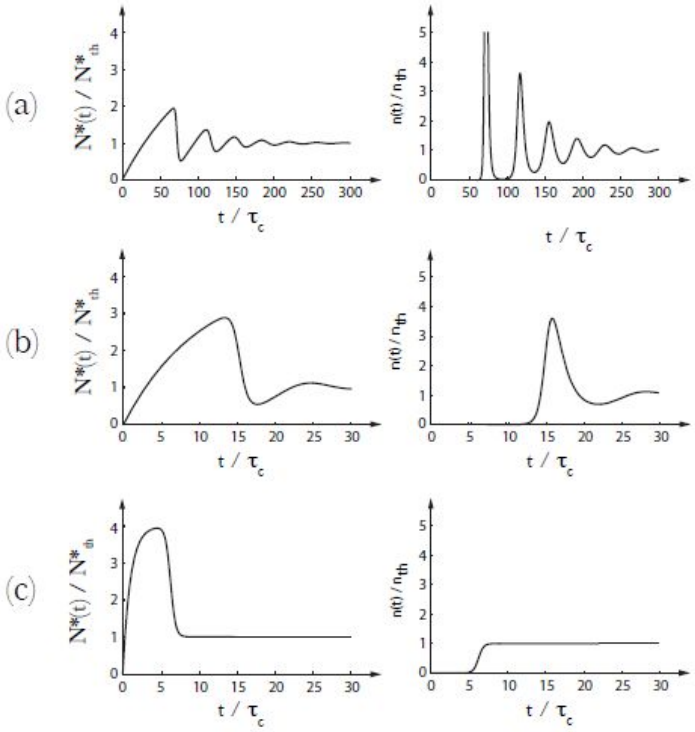

Figure 3.2: Numerical analysis of laser spiking in an ideal four-level laser with a filling factor $f_{\rm c}$ of unity pumped with an over-pumping ratio $r = 4$ and: (a) $\tau_2/\tau_{\rm c} = 100$; (b) $\tau_2/\tau_{\rm c} = 10$; (c) $\tau_2/\tau_{\rm c} = 1$. In each case the plots show the temporal evolution of the population inversion density and the photon density normalized to the steady-state value, $n_{\rm th}$.

Fig. 3.2 shows numerical solutions of these rate equations for a system with a filling factor f_c of unity and constant pumping at a rate corresponding to an over-pumping ratio (see Section 3.1.4) $r = 4$. In each of these simulations the pumping would, in the absence of stimulated emission, produce a population inversion density which was a factor of r greater than the threshold value. However, once the population inversion rises above N_{th} the rate of stimulated emission exceeds the rate of loss of photons from the cavity, increasing the density of photons within the cavity, and causing in turn the population inversion to burn down towards the threshold value. In all the cases shown the population inversion and photon densities tend to their steady-state values N_{th} and n_{th} respectively at long times after the pumping is switched on at $t = 0$. Substantial laser spiking is observed when $\tau_2/\tau_c \gg 1$, consistent with our discussion above, and the number and amplitude of the spikes in the laser output is seen to decrease as the ratio τ_2/τ_c is reduced. In particular, for the case $\tau_2/\tau_c = 1$ there is essentially no spike in the photon density. Instead the population inversion density photon density reaches rN_{th} before the smoothly-rising photon density burns the population inversion down to the steady-state value.

3.2 Q-switching

Q-switching is a technique that generates pulses of laser light with a peak power very much greater than the mean power that can be produced in steady-state operation. The high peak power is particularly useful in applications that require a high peak intensity, such as: laser cutting or drilling; or in pumping other laser systems such as Ti:sapphire amplifiers.

In a Q-switched laser the cavity is deliberately spoiled — imagine, for the moment, that one of the cavity mirrors is blocked — whilst the pumping builds up the population inversion. Since the losses of the modified cavity are so high, the population inversion can build up to much greater levels than the threshold inversion for lasing in the normal cavity without any lasing occurring. Once the population inversion has grown to a large value the mirror is unblocked, thereby switching the 'Q' of the cavity from a low value to a high value. High rates of stimulated emission then cause the density of photons in the cavity to grow very rapidly, burning the population inversion density well below the threshold value. The output from the laser takes the form of a giant pulse of radiation — essentially a single laser spike of the type discussed above.

Note that the pumping in Q-switched lasers may be pulsed or continuous. With pulsed pumping the ratio of the peak population inversion to the threshold value can be larger than can be achieved with continuous pumping, leading to a higher peak output power (although the mean output power may be similar).

3.2.1 Techniques for Q-switching

A wide variety of techniques have been used to implement Q-switching.

Rotating mirror

Historically one of the earliest techniques was to rotate one of the cavity mirrors with a motor. The cavity would then be completed briefly whenever the mirror was rotated into alignment. It is best to use a roof prism as shown schematically in Figure 3.3, rather than a plane mirror, since any ray incident in a plane perpendicular to the axis of the roof is reflected in the anti-parallel direction to the incident ray. Hence for the system illustrated alignment in the vertical plane is automatic — this is useful given the inevitable vibration in a mechanical system. Alignment in the horizontal plane is achieved once per rotation.

The disadvantages of this approach are obvious. The switching is rather slow and the timing of the output pulse has a relatively high uncertainty. In order to estimate the switching time, we suppose that the mirror has to be aligned to an angular tolerance about equal to the divergence of the oscillating cavity mode. This divergence will typically be of the order of 1 mrad. Hence, for a motor running at 24 000 r.p.m. the switching time is about 400 ns. The disadvantages are such that this technique is now used only very rarely.

Electro-optic switching

A more modern approach is the use of an electro-optic switch, as illustrated schematically in Figure 3.4. The refractive index of an electro-optic material changes when it is subjected to an external electric field. In the **Pockels effect**[3], observed in non-centro-symmetric materials, the induced refractive index varies linearly with the applied electric field. Suitable materials are KD*P or lithium niobate for visible to near infrared wavelengths, or cadmium telluride for the mid-infrared. In the **Kerr effect**[4] the change in refractive index is proportional to the square of the electric field. We note that all materials exhibit the Kerr effect to some extent. Since the Kerr effect arises from higher-order processes than in the Pockels effect (via the third-order susceptibility rather than the second-order), larger electric fields are usually required. Since these must be switched quickly for Q-switching to be effective, it is more usual to use the Pockels effect in this application.

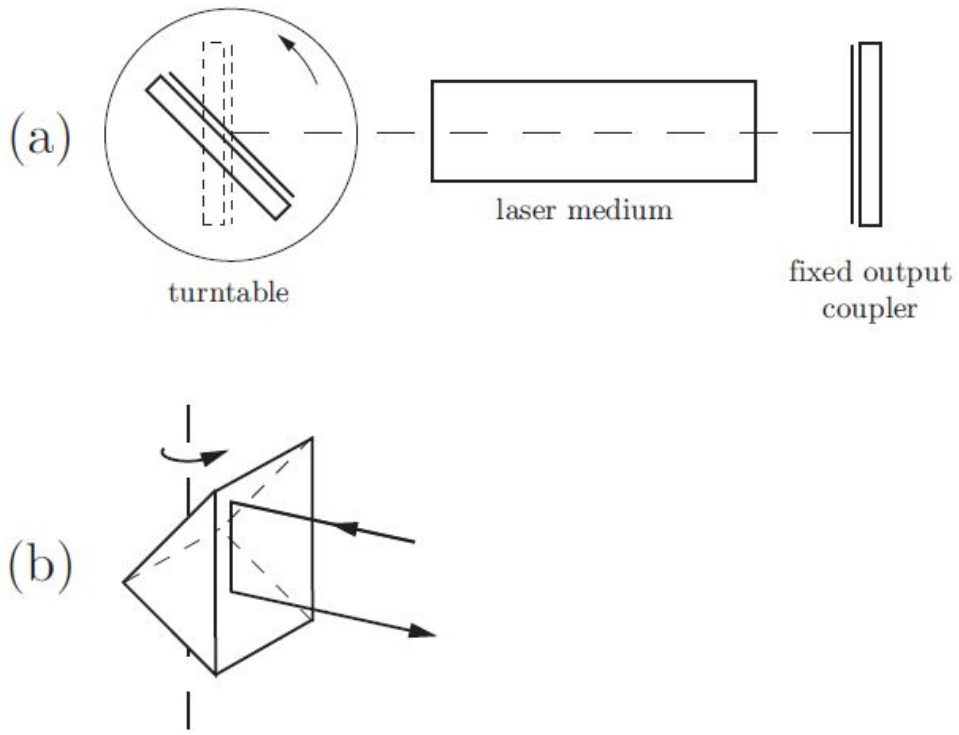

Figure 3.3: Q-switching with a rotating mirror: (a) schematic diagram of the layout of the optical cavity; (b) use of a roof prism as the rotating element. A roof prism has the advantage of automatically ensuring alignment of the reflected beam in one plane (in this case the vertical plane).

We can treat electro-optic devices of this type as wave plates with a birefringence which depends on the strength of the applied electric field. In the arrangement shown in Figure 3.4 a Pockels cell is oriented with its axis at 45^o to the axis of a polarizer. The cavity is held in a low-Q state by applying a certain voltage (the 'quarter-wave' voltage) to the Pockels cell, which then behaves as a quarter-wave plate. Consequently vertically-polarized light transmitted by the polarizer becomes circularly polarized after it transmission through the Pockels cell. Upon reflection from the end mirror of the cavity the handedness of the circular polarization changes, and hence on passing through the Pockels cell again it is converted to horizontally-polarized light and so reflected out of the cavity by the polarizer. Hence, with the quarter-wave voltage applied to the Pockels cell the end mirror of the cavity is effectively blocked.

The mirror can be unblocked by switching the voltage across the Pockels cell to zero, which reduces the birefringence of the Pockels cell to zero[5]. Vertically-polarized light can now pass through the polarizer and Pockels cell with essentially zero loss, and the Q-switched pulse can build up. A Kerr cell would be employed in essentially the same way.

Electro-optic switching has the advantages of fast switching times (of the order of 10 ns) and a high hold-off ratio, which allows the population inversion to be built up to many times the threshold value for the high-Q cavity. The disadvantages are that the technique requires several optics to be placed in the cavity (each with some insertion loss and subject to possible damage), and the devices and associated electronics are relatively expensive. However, the advantages are such that the technique is very widely used.

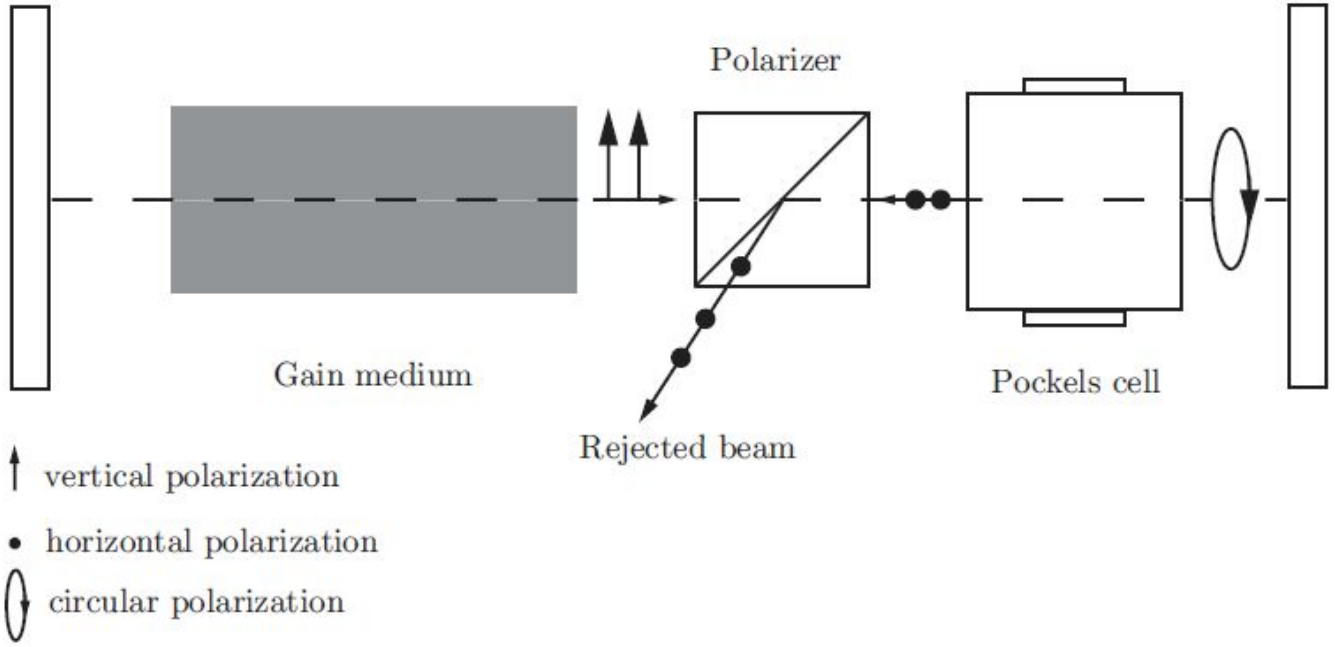

Figure 3.4: Electro-optic Q-switching. The system is shown with the Pockels cell voltage set so as to reject the beam reflected from the right-hand mirror.

Acousto-optic switching

If an acoustic wave is launched through an acousto-optic crystal the refractive index is changed slightly at the peaks and troughs as a result of the local expansion or contraction of the crystal lattice. The resulting periodic variation in the refractive index will cause incident radiation to be diffracted. We can think of this as Bragg reflection from the high- and low-index planes established in the crystal.

A wide variety of acousto-optic materials are available, including: fused silica, tellurium dioxide and lithium niobate for visible wavelengths; gallium arsenide for use in the infrared; and silicon for radiation in the 10 μm region.

Figure 3.5 shows an acousto-optic switch employed in a Q-switched laser. With the radio-frequency (RF) voltage applied to the crystal an acoustic wave propagates along the length of the crystal and is damped at the opposite end, the damping often being achieved by cutting the end of the crystal to form a wedge. Some fraction of the laser radiation incident on the crystal will be Bragg reflected away from the cavity axis and hence lost from the cavity mode. With the RF voltage applied, then, the acousto-optic crystal increases the losses of the laser cavity. Turning off the RF voltage removes this additional loss, and sets the cavity to a high-Q condition.

The advantages of acousto-optic switching are that it is relatively cheap, and insertion losses may be made low by orienting the crystal at Brewster's angle. However the hold-off ratio is low, and hence the technique is usually employed in continuously-pumped Q-switched lasers for which the over-pumping ratio is relatively modest. The switching times achieved are also relatively slow, the switching time corresponding to the time taken for the acoustic wave to propagate out of the region of the cavity mode. Taking the speed of sound in the crystal to be 5 $\mathrm{km\,s^{-1}}$, we can estimate the switching time to be 200 ns for a 1 mm diameter beam. The switching time can be decreased by focusing the cavity mode through the acousto-optic crystal.

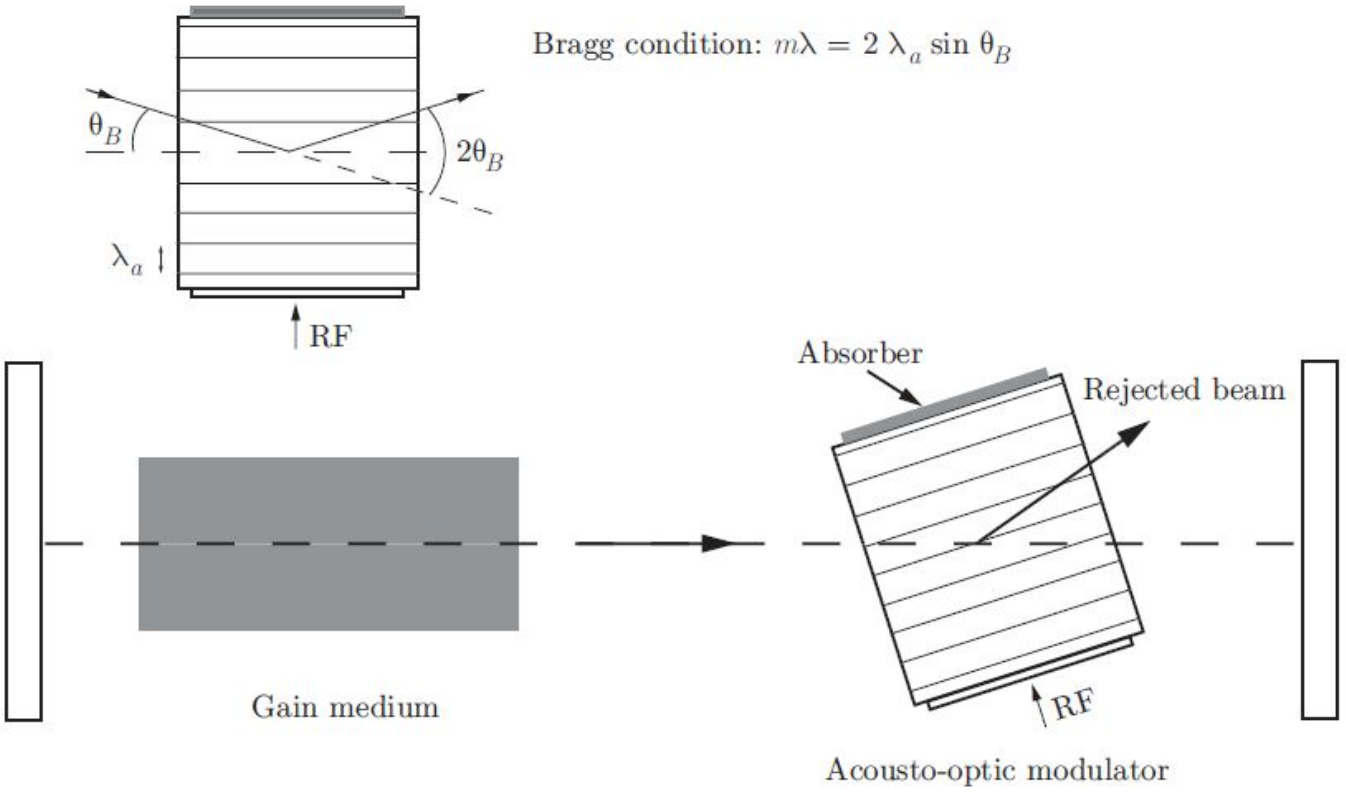

Figure 3.5: Acousto-optic Q-switching. Note that the Bragg angle θ_B has been greatly exaggerated. A value closer to 1° would be more typical. The insertion loss of the acousto-optic crystal can be eliminated (for one polarization) by rotating the crystal out of the plane of the paper through Brewster's angle.

Saturable absorbers

A very simple, and quite widely used, method for achieving Q-switching is to insert a saturable absorber into the laser cavity. During the pump pulse the absorber will cause the laser mode to experience a high loss and lasing is prevented. However once the population inversion reaches a high value the absorber will be subjected to an ever increasing flux of spontaneous and weak stimulated emission. As population is transferred to the upper level of the absorbing species the absorption starts to saturate, increasing the feedback from the cavity mirror and so increasing the rate of stimulated emission. This positive feedback leads to rapid saturation (or 'bleaching') of the absorber until, at full saturation, it represents essentially zero loss. The cavity is now in a high-Q state, and a giant pulse develops. It is clear that the absorber must have a strong absorption transition at the laser wavelength and also that the lifetime of the upper level must be reasonably long so that the bleached state is maintained for the duration of the Q-switched pulse. On the other hand, the upper level lifetime should be short compared to the desired time between Q-switched pulses. Many types of saturable absorbers have been used, such as solutions of dyes, and doped crystals such as Cr^{4+}:YAG.

Saturable absorbers are a simple (there is no external drive circuit) and cheap way of achieving Q-switching. The relatively slow turn on corresponds to many cavity round-trips, which allows good mode discrimination. As a consequence it is quite straightforward to achieve Q-switched lasing on a single transverse mode of the cavity. The disadvantages are that the shot-to-shot fluctuation in the output energy and in the timing of the output pulse can be relatively high, and that the absorber can degrade over time.

Modelocking

4.1 Introduction

Modelocking is an extremely important technique that is used in many laser systems. Its value lies in the ability to generate pulses of laser radiation which have a short duration — typically picoseconds (1 ps $\equiv 10^{-12}$ s) to femtoseconds (1 fs $\equiv 10^{-15}$ s) — and a peak power many orders of magnitude greater than the mean power that can be extracted from the laser medium. Pulses from modelocked lasers have found applications in a wide variety of scientific and technical applications including: imaging techniques such as two-photon laser-scanning fluorescence imaging, ballistic imaging, and optical coherence tomography; ultrafast chemistry, including coherent control of chemical reactions; pump-probe measurements of solid-state materials; and short-pulse laser machining and processing of materials.

4.2 General ideas

We recall that the longitudinal modes of a laser cavity of length L are spaced in angular frequency by an amount,

$$\Delta\omega_{p,p-1} = 2\pi\frac{c}{2L} = \frac{2\pi}{T_c}, \tag{4.1}$$

where T_c is the time taken for light to travel once round the cavity. In order to avoid clutter, it will be convenient to let $\Delta\omega \equiv \Delta\omega_{p,p-1}$ throughout this lecture.

The right- and left-going waves of a cavity mode may be written in the simplified form,

$$E(z,t) = a_p \exp\left(\mathrm{i}[k_p z \mp \omega_p t + \phi_p]\right), \tag{4.2}$$

where a_p is the amplitude of the wave, ϕ_p is the phase of the wave, the upper (lower) sign refers to waves propagating towards positive (negative) z, and we have ignored the details of the transverse profile of the waves.

Suppose that the laser oscillates simultaneously on N such modes. For simplicity, we will take there to be an odd number of oscillating modes labelled by $p = -P/2 \ldots P/2$. The amplitude of the laser radiation emerging from the output coupler would have the form,

$$E(z,t) = \sum_{p=p_0-P/2}^{p=p_0+P/2} a_p \exp\left(\mathrm{i}[k_p z - \omega_p t + \phi_p]\right), \tag{4.3}$$

where p_0 is the number of the central longitudinal mode.

Figure 4.1(a) shows a calculation of the output intensity of a laser oscillating simultaneously on 61 modes with equal amplitudes, but with random — though constant — relative phases ϕ_p. The output has two key features,

1. During an interval equal to T_c the output contains sharp fluctuations in intensity which are randomly distributed;

2. The sharpest of these peaks has a duration of order $2\pi/\Delta\omega_{\rm osc}$, where $\Delta\omega_{\rm osc} = (N-1)\Delta\omega$ is total bandwidth occupied by the oscillating modes;

3. The pattern is periodic with a period equal to the cavity round-trip time T_c.

In contrast, Figure 4.1(b) shows the intensity calculated for the case when the modes all have the same phase ($\phi_p = 0$). Now the output consists of a train of short pulses separated by the cavity round-trip time. This situation is known as **modelocking**.

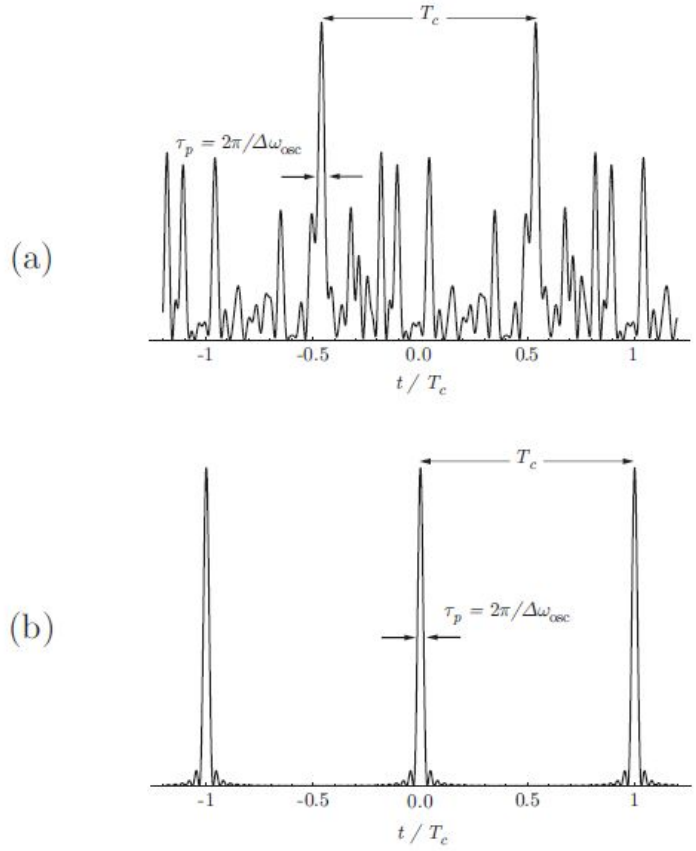

Figure 4.1: Calculated intensity as a function of time emerging from a laser with $N = 61$ modes oscillating simultaneously with equal amplitudes and (a) random (but fixed) phases; (b) identical phases. Notice that for both plots the pattern repeats every cavity round-trip time T_c, and that the sharpest features have the same width. It should also be emphasized that the vertical scales of the two plots are quite different: the peak intensity of the modelocked pulses are approximately N times greater than the mean intensity in (a).

4.2.1 Simple analysis

To gain further insight, let us calculate the output for the case where all the modes have the same amplitude, E_0. The frequencies of the oscillating modes may be written,

$$\omega_p = \omega_{p_0} + p\Delta\omega, \tag{4.4}$$

where ω_{p_0} is the frequency of the centre mode. Similarly, we consider the case in which the phases of adjacent modes differ by a constant amount, $\Delta\phi$, so that we have,

$$\phi_p = \phi_{p_0} + p\Delta\phi. \tag{4.5}$$

From eqn (4.3) the amplitude of the radiation emitted by the laser will then be:

$$E(\boldsymbol{r}, t) = E_0 \sum_{p=-P/2}^{p=+P/2} \exp\left[\mathrm{i}(k_p z - \omega_{p_0} t - p\Delta\omega t + \phi_{p_0} + p\Delta\phi)\right]. \tag{4.6}$$

For convenience we consider the amplitude at $z = 0$ and also take the special case[1] in which $\Delta\phi = 0$. Then,

For convenience we consider the amplitude at $z = 0$ and also take the special case[1] in which $\Delta\phi = 0$. Then,

$$E(0,t) = E_0 \exp\left[-\mathrm{i}(\omega_{p_0} t - \phi_{p_0})\right] \sum_{p=-P/2}^{p=+P/2} \exp\left(-\mathrm{i}p[\Delta\omega t]\right). \tag{4.7}$$

The sum appearing in the above is a geometric progression, which may be summed in the usual way to give,

$$E(0,t) = E_0 \exp\left[-\mathrm{i}(\omega_{p_0} t - \phi_{p_0})\right] \frac{\sin\left(\frac{N}{2}\Delta\omega t\right)}{\sin\left(\frac{1}{2}\Delta\omega t\right)}. \tag{4.8}$$

Hence the output takes the form of a (travelling) wave of mean frequency ω_{p_0} multiplied by an envelope function $\frac{\sin\left(\frac{N}{2}\Delta\omega t\right)}{\sin\left(\frac{1}{2}\Delta\omega t\right)}$.

Taking the intensity to be given by $|E(0,t)|^2$ (i.e. ignoring constants of proportionality):

$$I(t) = E_0^2 \frac{\sin^2\left(\frac{N}{2}\Delta\omega t\right)}{\sin^2\left(\frac{1}{2}\Delta\omega t\right)}. \tag{4.9}$$

This is just the function plotted in Fig. 4.1(b).

The temporal behaviour calculated above corresponds to a train of intense modelocked pulses. Notice that the peak intensity of each pulse is equal to $N^2 E_0^2$, and hence the peak intensity is a factor of N greater than the mean intensity of N modes oscillating with random phases[2].

$$\tau_p = \frac{2\pi}{N\Delta\omega} = \frac{T_c}{N} \approx \frac{2\pi}{\Delta\omega_{\rm osc}}, \tag{4.10}$$

where $\Delta\omega_{\rm osc} = (N-1)\Delta\omega$ is the (angular) frequency bandwidth covered by the oscillating modes.

It should be clear from the above that if we were able to lock the phase of N longitudinal modes of a laser we could generate a stream of pulses of high peak intensity and short duration. In practice the number of modes that can be locked ranges from several hundred to several thousand, leading to a huge increase in peak power. The duration of the pulses is determined by the frequency width of the laser transition; transitions with broad linewidths can (in principle) generate shorter pulses. In practice pulses with durations from a few picoseconds to a few femtoseconds can be generated; a very active area of current research is the development of techniques related to modelocking to generate pulses in the attosecond range (1 as $\equiv 10^{-18}$ s).

4.3 Active modelocking techniques

Active modelocking techniques fall into two classes: **amplitude modulation (AM)** and **frequency modulation (FM)**. We note that we usually require a single modelocked pulse to circulate within the laser cavity. For a linear laser cavity this is achieved by locating the active component near one of the end mirrors, and driving the component at a frequency equal to the longitudinal mode spacing $\Delta\omega$. However, it is possible to generate N such modelocked pulses within the laser cavity by locating the modulator suitably and driving it at $N\Delta\omega$. **Harmonic modelocking** of this type is occasionally used to increase the pulse repetition frequency. In principle it can also be used to generate shorter pulses in homogeneously broadened lasers (see eqn (4.19)).

4.3.1 AM modelocking

Once modelocking is achieved, the intra-cavity radiation takes the form of a short pulse which circulates round the laser cavity. An obvious way, therefore, to achieve modelocking is to introduce some sort of 'shutter' into the cavity which lets through short pulses, but which blocks — or increases the loss of — longer pulses. Radiation which happens to reach the shutter when it is open will pass through with low loss, will experience net round-trip gain, and hence will grow. In contrast, the radiation arriving at the shutter outside this interval will experience a round-trip loss which is greater than the round-trip gain, and so will be absorbed.

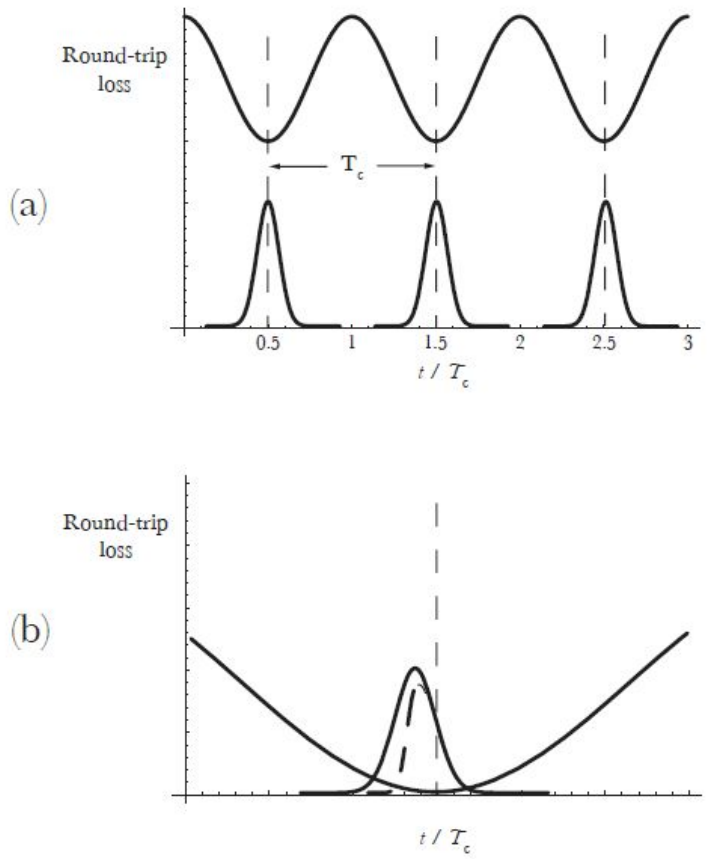

Figure 4.2: Schematic diagram of AM modelocking. In (a) the pulses arrive perfectly synchronized with the shutter and hence experiences minimum loss. The pulse shown in (b) arrives slightly too early, and hence the leading edge is reduced in intensity. The effect is to move the peak of the pulse to later times, closer to synchronization with the shutter.

After many round-trips a pulse will form which propagates round the cavity in synchronization with the shutter, as illustrated schematically in Figure 4.2. The duration of the modelocked pulse so formed is determined by an equilibrium between two competing effects. The shutter tends to decrease the duration of the pulse (and hence increase the bandwidth) since the losses of the leading and trailing edges of the pulse are greater than those experienced by the centre of the pulse. However, the gain medium can only amplify a finite frequency bandwidth which limits the minimum possible duration of the modelocked pulse.