Chapter 14
Optical Frequency Oscillators: Lasers

14.1 Fundamentals

14.1.1 Introduction

The first international conference at which papers were presented on the subject of "optical masers," as they were then called, was at Ann Arbor, Michigan in June 1959. The principal topic of the conference was not lasers, but the optical pumping method of observing magnetic resonance in free atoms, a technique that had recently been introduced by Kastler in Paris. The session devoted to lasers was a "miscellaneous session" in which papers on theoretical aspects of laser oscillation in gas discharges and ruby crystals were presented by Gould, Javan, and Schawlow, among others.

The workings of a laser do not involve any physical theory or even practical technique that was not already known for some time. The quantum theory of the absorption and emission of light by atoms and molecules was well established, there was abundant spectroscopic data such as wavelengths and line intensities, and the theory of light wave optics and the techniques of optical interferometry were well advanced. The study of electrical discharges through rarefied gases and crystal optics had been pursued since the 19^{th} century. This undoubtedly explains the veritable explosion that occurred in the number of reports of laser action once the first appeared.

The special properties of lasers as quantum oscillators in the infrared to ultraviolet regions of the spectrum will now be treated in the broader context of quantum oscillators in general. Their impact on the design and performance of the microwave standards (except the H-maser), which has literally transformed these standards, and their development as frequency standards in the infrared and optical regions of the spectrum will be treated in succeeding chapters.

F.G. Major, The Quantum Beat

14.1.2 The Resonance Frequency Width of Optical Cavities

We have seen in the case of atomic and molecular beam masers that the spectral width of the resonant response of the atoms or molecules is far sharper than the resonant modes of the cavity in which the particles interact with the radiation field. Since these cavity modes are well separated, this means that oscillation occurs in a unique mode of oscillation of the field in the cavity: the one tuned to resonance with the particle frequency.

In contrast, at optical frequencies it is the "cavity", which may in fact consist of only two small parallel mirrors some distance apart, whose resonance modes have narrow spectral widths compared to those of the "active" atoms or molecules. The arrangement of two highly reflecting, precisely parallel surfaces to form an optical resonator has as precursor a high-resolution spectroscopic device called a *Fabry–Pérot* interferometer, whose introduction in 1899 far predates the era of modern optics. Its capability as a high resolution spectroscopic device derives from the sharp resonant peaks in the intensity of light transmitted through it. The degree to which it is able to resolve close resonances due to neighboring wavelengths is specified by a quantity called the *finesse*. A useful physical insight into the significance of this quantity is obtained by the following approximate but instructive argument: Suppose the light wave to be analyzed is reflected back and forth traversing the distance between the mirrors an average of $2N$ times before decaying, and that the distance between the mirrors is such that the light resonates in the n^{th} longitudinal mode, so that the length of the cavity is $n\lambda/2$. It would follow that the light wave is coherent over a time of $Nn\lambda/c$, and therefore the relative width of the Fourier spectrum $\Delta\nu = \nu/(nN)$, hence the resolving power $\lambda/\Delta\lambda$ is given by:

$$\frac{\lambda}{\Delta\lambda} = nN \tag{14.1}$$

It is this quantity N that is a measure of the fineness of resolution, given the name *finesse* (F), defined in terms of the (power) reflectivity R of the mirrors as follows:

$$F = \frac{\pi R^{1/2}}{(1-R)} \tag{14.2}$$

Currently extraordinary experimental values of R reaching up to an incredible 99.998% have been reported, corresponding to $F = 157{,}000$. Another quantity important in the design of a Fabry-Pérot cavity is the *free spectral range* (FSR), which specifies the range of values of resonant frequencies/wavelengths for which there is no ambiguity in the assignment of the order n. In terms of frequencies it is $c/2n_rL$, where n_r is here the refractive index. In what follows we shall assume $n_r = 1$ unless the dispersion of the medium is relevant. For a given cavity length, the larger the order number n, the closer will be the wavelengths of the resonances. In fact it is not difficult to show that:

$$\lambda_n - \lambda_{n+1} = \frac{\lambda_n\lambda_{n+1}}{2L} \approx \frac{\lambda^2}{2L}. \tag{14.3}$$

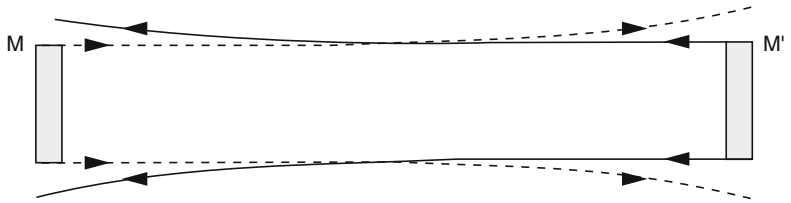

Figure 14.1 Diffraction loss at the mirrors of an optical cavity

Since the Fabry–Pérot interferometer was originally applied to conventional light sources with limited light wave coherence, the spacing of the mirrors was typically on the order of millimeters, rather than centimeters or tens of centimeters as is typical of gas laser cavities. Strictly speaking, such an open arrangement of mirrors does *not* have discrete resonance frequencies, unlike a completely closed cavity with reflecting walls. However, detailed computations on such a cavity have shown that there exist more or less discrete sets of modes with the optical field localized along the axis between the two mirrors. For these modes the loss of optical energy from the cavity due to diffraction, the inevitable spreading out of a wave, is small. We can readily show this is plausible if we accept the result from wave theory that a plane wave reflected by a circular plane mirror of diameter D (large compared to the wavelength) will be diffracted at an angle of about λ/D (radians), as shown in Figure 14.1. Such a wave traveling to another similar mirror a distance L away will partly fall outside the rim of that mirror and suffer a fractional loss of $4\lambda L/D^2$, provided that we make the crude assumption that the energy of the beam is spread uniformly over the expanded area. Clearly, even if the mirrors were perfectly reflecting, the cavity field would still decay in energy due to diffraction, with a consequent broadening of the resonance spectrum. The fractional loss of $4\lambda L/D^2$ occurs at each mirror and repeats at intervals equal to the transit time of the light wave between the two mirrors, namely L/c, where as usual, c is the velocity of light. It follows that the average fractional loss of energy per unit time is $4\lambda c/D^2$. Now we can apply this result to the two counter-traveling waves of equal amplitude whose sum is one of the stationary axial modes, belonging to the quasi-discrete set supported by the cavity. These longitudinal modes are analogous to vibrations on a string at frequencies such that the phase of the wave after traveling $2L$ is a whole number of cycles: Symbolically, this means that $(2L/c)\nu_n = n$, where n is a whole number, and ν_n is the frequency of oscillation of the optical field in the nth mode. Then the diffraction limited Q-factor for that mode is as follows:

$$Q_n = \frac{2\pi\nu_n}{\left(\frac{1}{E}\frac{dE}{dt}\right)} = \frac{\pi D^2}{2\lambda_n^2}, \tag{14.4}$$

where $\lambda_n = c/\nu_n$ is the wavelength of the nth mode. For example, if we take for the mirror diameter $D = 2.5$ cm and $\lambda = 632.8$ nm (the common wavelength of the He–Ne laser output), we find for that pure axial mode $Q \approx 10^{10}$. This is

significantly larger than is typical of microwave cavities and is even higher than most optical atomic transitions. Actually, the situation is far more complicated, in that the intensity profile of the beam is far from uniform over the cross section of the beam; rather, it has a radial distribution that can be analyzed in terms of certain radial modes. These are designated as TEM modes (transverse electromagnetic) with indices specifying the order and hence the number of zeros in the intensity distribution. For example, if we take the mirror axis as the z-axis of a coordinate system, then in the TEM_{21} mode, the field intensity has two zeros along, say, the x-direction and one zero along the y-direction. Beam profiles for some of the lower-order radial modes are shown in Figure 14.2. The least complicated mode (TEM_{00}) has only one maximum, which occurs on the axis and is described by the form $\exp(-r^2/{r_0}^2)$, called a Gaussian function. The output beam of a laser oscillating in this mode is called a Gaussian beam, and the theory of the action of optical elements such as lenses and mirrors on such beams is called Gaussian optics. It differs from ordinary ray optics, which takes no account of the wave nature of light; it is characterized by the absence of sharply defined focal points and beam profiles. The usual lens and mirror formulae of ray optics are not valid.

Optical wave theory shows that the diffraction loss is radically smaller if instead of plane mirrors, concave mirrors are used in a *confocal* arrangement, in which the focal points of the two mirrors coincide at the midpoint between them, as shown in Figure 14.3.

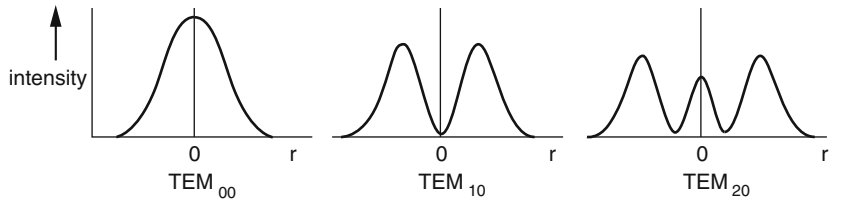

Figure 14.2 The intensity distribution for some low order-radial modes in an optical cavity

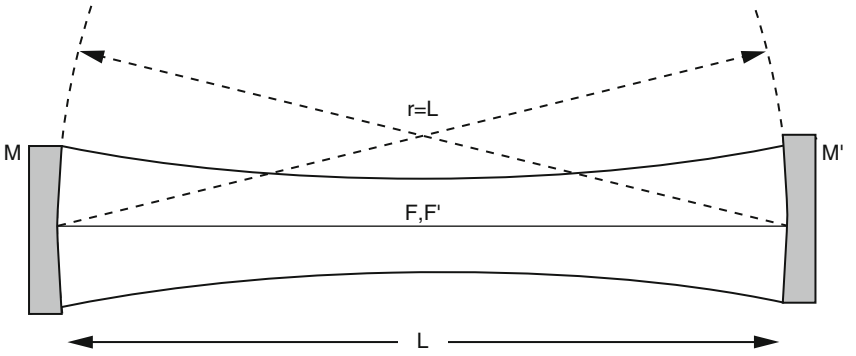

Figure 14.3 A confocal optical cavity

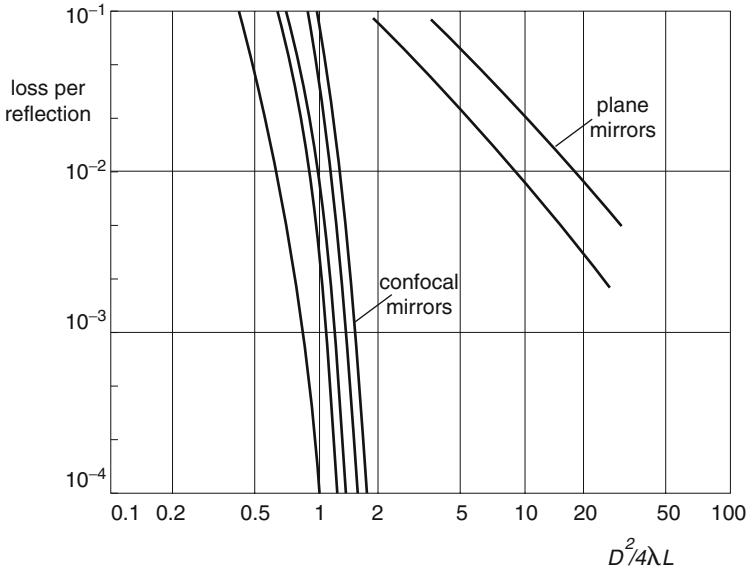

Figure 14.4 A plot of the optical loss as a function of $D^2/4\lambda L$ (Boyd, 1961)

The results of detailed computations on the diffraction loss for various oscillation modes in an optical two-mirror cavity are reproduced in Figure 14.4, in which the fractional loss is plotted against the parameter $D^2/4\lambda L$ (Boyd, 1961).

Note that according to the approximate theory outlined above, the graph for the plane mirrors should be linear with a slope of -1. We see from these results that even for the unfavorable case of plane mirrors, fractional losses as low as 10^{-4} are attainable (assuming the mirrors are perfectly aligned) for $D^2/4\lambda L \approx 100$, a practical figure. This points to the limit on the Q of these cavity modes not really being set by diffraction, but rather by the imperfect reflectance of the mirrors, which in practice rarely exceeds 99.99%, corresponding to a fractional loss of 10^{-4}. In any event, to be useful, a laser oscillator must provide an output beam, and therefore at least one of the mirrors must be partially transparent, with a consequent power loss from the cavity.

Since an optical cavity commonly has dimensions very large compared to the wavelength, the various resonance modes, with their characteristic stationary field distributions, have frequencies that differ fractionally very little from each other. The mode spacing, of course, need not be small for microwave cavities, since in that case the cavity dimensions can be of the same order of magnitude as the wavelength, and the lower mode frequencies are relatively far apart. If the optical cavity consisted of a truly closed cavity with reflecting walls, analogous to a microwave cavity, then if its dimensions were large compared to the wavelength, there would be a very large number of modes lying within any given

frequency band, all having an appreciable Q-value. In fact, that number in the general 3-dimensional case appears in the theory of blackbody radiation and is given by the following:

$$\Delta N = \frac{8\pi V}{\lambda^3}\frac{\Delta\nu}{\nu}, \tag{14.5}$$

where ΔN is the number of modes having a frequency in the interval $\Delta\nu$ centered on the frequency ν, and V is the volume of the cavity, which classical theory has shown can have any shape, provided that its dimensions are very much larger than the wavelength. To illustrate just how large ΔN can be, let us assume $\lambda = 500$ nm, $V = 100\,\text{cm}^3$, and $\Delta\nu/\nu = 10^{-7}$; substituting into the formula yields $\Delta N \approx 2 \times 10^9$! It is indeed fortunate that quasi-discrete Gaussian modes do exist in a wide-open Fabry–Pérot resonator, with only a few radial modes having a high Q-value, to restrict the number of modes into which the stimulated emission from the atoms occurs. Not only does sustained laser action in these modes become possible, but it yields the extraordinary directionality of the laser output beam, and with proper selection of axial modes, great spectral purity.

14.1.3 Conditions for Sustained Oscillation

As with beam masers, sustained optical frequency oscillation of the field in a resonator is possible if atoms or molecules are present that through stimulated emission yield a net gain in the field energy sufficient to make up for all losses, including that represented by the output beam. However, while the general principles are identical for microwave and optical frequency oscillators, there are many important practical differences that give them quite different physical aspects. In addition to the obvious differences attendant upon the very different wavelengths, the kind of coupling of the atoms or molecules with the optical field is also different: The atomic beam masers involve magnetic dipole transitions, while lasers involve the much more strongly induced electric dipole transitions. Furthermore, while in the magnetic dipole transitions the field acts on a permanent atomic magnetic moment, the existence of a permanent electric dipole moment is excluded on the basis of a fundamental symmetry property of the internal forces holding an atom together. The symmetry is broken when an external electric field is introduced; in that case, oppositely directed forces are exerted on the positive nucleus and negative electrons, resulting in a dipole moment. Transitions result from the action of the electric component of the optical field, which induces an oscillating electric dipole moment in the atom. The amplitude of that dipole moment depends on the dynamical response of the particular atom; we recall the classical model of an atom having elastically bound electrons used by Lorentz in his theory of optical dispersion to predict that response.

To achieve a net gain of power from an atomic or molecular system, it must be prepared with a preponderance of population in the upper of two quantum energy levels, between which transitions are to be stimulated. This, it will be recalled, is

simply because the probability per atom per unit time for stimulated emission is exactly the same as for absorption. The achievement of this preferential population of the upper energy state in the case of optical transitions is complicated by two circumstances: First, unlike the microwave case, an optical quantum (photon) has considerably greater energy than the mean thermal energy of particles in equilibrium at ordinary temperatures. This, it will be recalled from Boltzmann's theory, implies that in thermal equilibrium the number of atoms of a gas in the lower of the two states will be far greater than the number in the upper state. Thus according to the theory, if the number in the upper quantum state is N_1 and the number in the lower state N_2, then we have in equilibrium at absolute temperature T,

$$\frac{N_1}{N_2} = e^{-\frac{h\nu}{kT}}. \tag{14.6}$$

The temperature T is, of course, positive; hence in equilibrium we must always have $N_1 < N_2$. For example, if $T = 300°\text{K}$ and $\nu = 6 \times 10^{14}\,\text{Hz}$, we find that, on an average, only one atom in 10^{41} is in the upper state! Clearly, then, for laser action we require very nonequilibrium conditions; in fact, we require what is called *population inversion*, or a "negative (absolute) temperature."

The second essential difference caused by the greater energy of the optical photon is that the probability (per atom) per unit time for *spontaneous* emission is far from being negligibly small, as it was in the microwave case. We can see this from the Einstein expression for the ratio of his A- and B-coefficients

$$\frac{A_{nm}}{B_{nm}} = \frac{8\pi h\nu^3}{c^3}, \tag{14.7}$$

if we recall that the probability per unit time for stimulated emission is $B_{12}\rho_\nu$, where ρ_ν is the spectral energy density of the optical field causing the transitions, given by $\rho_\nu = I_\nu/c$, for a parallel light beam of spectral intensity I_ν. At ordinary light levels such as might exist in a conventional lamp, where I_ν is perhaps on the order of $10^{-8}\,\text{watt/m}^2 \cdot \text{Hz}$, we find that the probability of spontaneous emission is about 2000 times greater than that of stimulated emission. This shows why stimulated emission plays an insignificant role in the operation of a conventional lamp; in fact, all ordinary sources of light, from the common tungsten lamp to the sun, are examples of spontaneous emission. However, in a lasing medium the energy of the optical field is concentrated in a narrower spectral width, which means a far larger spectral energy density and the emergence of stimulated emission as an important process. It is important to recall that spontaneous emission, in contrast with stimulated emission, is indifferent to whether an optical field is present and is induced with *random phase* by "zero-point" quantum fluctuations in the optical field. Stimulated emission/absorption, on the other hand, results from induced electric dipole oscillation in the individual atomic systems and is correlated in phase to the common stimulating optical field. In quantum theory, stimulated emission and absorption are different consequences of the same process; which one is manifested is simply a matter of whether the initial state is the one with lower energy or higher

energy. From what has been said about the Boltzmann distribution of populations in a system in thermal equilibrium, it follows that a light beam passing through a medium in thermal equilibrium will always suffer absorption at those frequencies in its spectrum that are resonant with transitions in the system. In such a case, the intensity of a monochromatic light beam becomes weaker as it passes through a medium with a resonant transition. On the other hand, if by some means, such as an electrical discharge or intense optical pumping, a population inversion is sustained, then that monochromatic light beam *increases* in amplitude; it is amplified, as shown schematically in Figure 14.5.

To express these notions more quantitatively, suppose a monochromatic parallel beam of intensity I_ν watts/m^2 passes through an atomic/molecular medium with n_1 particles/m^3 in the upper quantum state and n_2 particles/m^3 in the lower; and let them have a frequency response (resonance line shape) $g(\nu)$, so that: according to the definition of B_{12}, the probability of the optical field stimulating emission (in unit volume) at the frequency ν is $n_1 B_{12}(I_\nu/c)g(\nu)$ per unit time, with an identical expression for absorption, except that n_1 is replaced by n_2. In practice, because of the various spectral broadening mechanisms such as the Doppler effect in gases, the function $g(\nu)$ will have a bell-shaped graph, generally broader than the spectral width of the light, so that the assumption of a monochromatic beam having the single frequency ν is not an unrealistic one. It follows that since each transition involves the exchange of one quantum of energy $h\nu$, the *net* rate of energy exchanged coherently with the field (per unit volume) is $(n_1 - n_2)B_{12}(I_\nu/c)h\nu g(\nu)$. If we choose the direction of the beam to be the z-axis of a coordinate system and balance the energy flow in the beam with the amount emitted (or absorbed) by the atoms, we are led to the following:

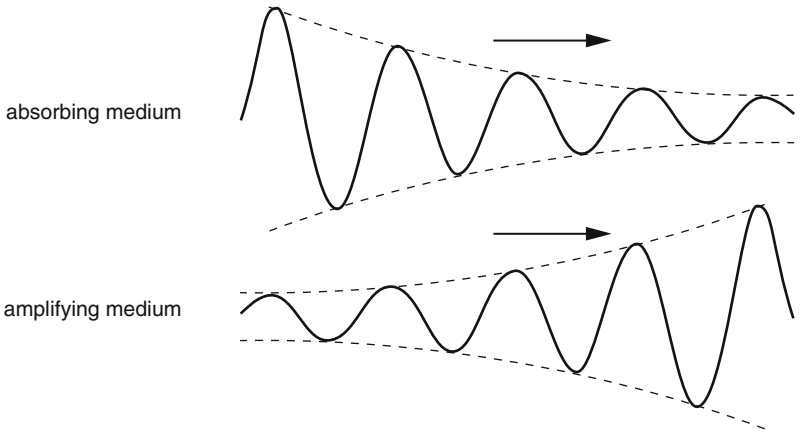

Figure 14.5 A schematic illustration of a light wave passing through absorbing and amplifying media

$$\frac{dI_\nu}{dz} = (n_1 - n_2)h\nu B_{12} g(\nu)\frac{I_\nu}{c}. \tag{14.8}$$

This has a solution of the form

$$I_\nu(z) = I_\nu(0)e^{\gamma z}, \tag{14.9}$$

where

$$\gamma = (n_1 - n_2)B_{12}\frac{h\nu}{c}g(\nu) \tag{14.10}$$

is the (exponential) gain constant. As expected, this shows that for the light beam to be amplified, we must have $(n_1 - n_2) > 0$, that is, population inversion with more atoms in the upper state than the lower. In a system in thermal equilibrium, as already emphasized, we have necessarily $(n_1 - n_2) < 0$, and the intensity falls exponentially, in agreement with the classical experimental law, sometimes called Lambert's law. The two cases are illustrated in Figure 14.5. We can usefully rewrite the gain constant γ in terms of Einstein's A-coefficient, since the latter is related to the mean lifetime of the upper state against spontaneous emission, a lifetime that can be deduced from the empirical "natural" width $\Delta\nu_n$ of the emission line. The result is as follows:

$$\gamma = \frac{1}{4}(n_1 - n_2)\lambda^2 \Delta\nu_n g(\nu). \tag{14.11}$$

The line shape factor $g(\nu)$, which gives the spectral response of the atoms to the optical field, may result from a number of different processes. For some applications it is important to draw a distinction between two different types of broadening mechanisms: homogeneous and inhomogeneous. As we saw in Chapter 7, the distinction applies to a group of atoms: If a broadening mechanism affects all atoms *identically*, such as the natural lifetime of the radiating state or collisions with other particles that interrupt the radiation process, then it is homogeneous. Lifetime broadening, for example, we know leads to a Lorentzian line shape:

$$g(\nu) = \frac{1}{\pi}\frac{\frac{\Delta\nu}{2}}{(\nu - \nu_0)^2 + \left(\frac{\Delta\nu}{2}\right)^2}, \tag{14.12}$$

where $\Delta\nu$ is the width of the $g(\nu)$ versus ν curve at half its maximum, which occurs at $\nu = \nu_0$ and has a value there of $2/\pi\Delta\nu$.

On the other hand, it can happen that each individual atom in the group has its own slightly different frequency because, for example, the atoms have different velocities, and therefore different Doppler shifts in their frequency, or perhaps because each atom sees a slightly different environment; in this case we say the broadening is inhomogeneous. For the case of Doppler broadening in a gas in thermal equilibrium, we found

$$g(\nu) \approx \exp\left(-4ln2\frac{(\nu - \nu_0)^2}{\Delta\nu^2}\right). \tag{14.13}$$

In order to set up the conditions for sustained oscillation at optical frequency, we combine the essential elements of a feedback oscillator by placing the amplifying atomic medium inside an optical resonator. And as with any other feedback oscillator, the threshold condition for oscillation to break out is that the feedback be regenerative and the loop gain $G = 1$. To obtain an expression of these two conditions explicitly in terms of a specific model, assume that we have a Fabry–Pérot resonator filled with a population inverted gas acting as a distributed amplifier. Let R_1 and R_2 represent the ratios of reflected to incident light intensity at the two mirrors, and let α be an absorption constant to account for all distributed loss of intensity due to interaction with the gas, so that we can write the condition on the loop gain as follows:

$$R_1 R_2 \exp\left[(\gamma - \alpha)\, 2L\right] = 1, \tag{14.14}$$

from which we deduce the threshold value of γ to be

$$\gamma = \alpha + \frac{1}{2L} ln \left(\frac{1}{R_1 R_2}\right). \tag{14.15}$$

The condition on the phase is a little more complicated, since the light travels through an amplifying medium that is dispersive; that is, the velocity of a light wave through it depends on the frequency of the wave. The interaction of the light with an atomic medium near a resonance can strongly affect the velocity of the light wave in a frequency-dependent way. If we define $c/n(\nu)$ as the velocity of light in the medium, then for a light wave starting from any point in the cavity, to have the same phase after making a complete round trip between the two mirrors requires the following phase condition:

$$\frac{2L}{c/n(\nu)}\nu = m, \tag{14.16}$$

where m is a whole number. In the absence of the atoms, $n(\nu) = 1$ and $\nu = \nu_m$, the cavity resonant frequency in the mth order axial mode. We will not attempt to derive the expression for $n(\nu)$ but merely state the important result that it involves the frequency dependence of absorption (or in our case the stimulated emission), and its substitution in the phase condition leads to the following approximate result for the actual frequency of oscillation:

$$\nu = \nu_m \left\{1 - \frac{(\nu_0 - \nu_m)}{\Delta\nu} \frac{\gamma}{k}\right\}, \tag{14.17}$$

where $k = 2\pi/\lambda$ is the magnitude of the wave vector. This shows that oscillation does not take place exactly at the resonant mode frequency ν_m of the cavity: an example of *frequency pulling* such as we already encountered in the hydrogen maser. It can be shown that this result can be rewritten to show the same dependence on the relative line widths of the atomic and cavity resonances; here, however, it is the cavity that has the sharper resonance, rather than the atoms.

14.1.4 The Sustained Output Power

The threshold conditions alone do not, of course, tell us anything about how and to what level the optical power builds up in the laser cavity. As with all feedback oscillators, once the threshold is passed, oscillations will start from ever-present incoherent zero-point excitations of the field or, as in this case, the spontaneously emitted light. To predict the further buildup of the optical field, we must take into account the dependence of the population difference itself on the amplified field it generates. This involves taking the theory of the interaction of the atoms with the optical field to a higher order of approximation, beyond the approximation so far implied. This was done by W.E. Lamb, who by developing the quantum theory of interaction between atom and field to the third order in the field amplitude gave explicit expressions for the coefficients α_n and β_n in the "equation of motion" for the field amplitude E_n of the nth mode:

$$\frac{dE_n}{dt} = \alpha_n E_n - \beta_n E_n^3, \tag{14.18}$$

which applies to a loop gain $G > 1$ (Lamb, 1964). This leads to a steady state when $dE_n/dt = 0$, which occurs for $E_n^2 = \alpha_n/\beta_n$. For the stationary field modes assumed in the theory, Lamb found that the linear gain factor α_n as a function of tuning has a Gaussian shape arising from Doppler broadening, while the nonlinear "saturation" factor β_n is much less Doppler broadened, exhibiting mainly lifetime broadening. This causes the overall frequency dependence of the laser output to exhibit what is now called the *Lamb dip*, a local minimum at the peak of the Doppler line shape, as shown in Figure 14.6. This sharp spectral feature, whose width reflects the homogeneous line width of the atoms, rather than the much broader Doppler width, has proved very useful in stabilizing the frequency of oscillators in the infrared and optical regions of the spectrum.

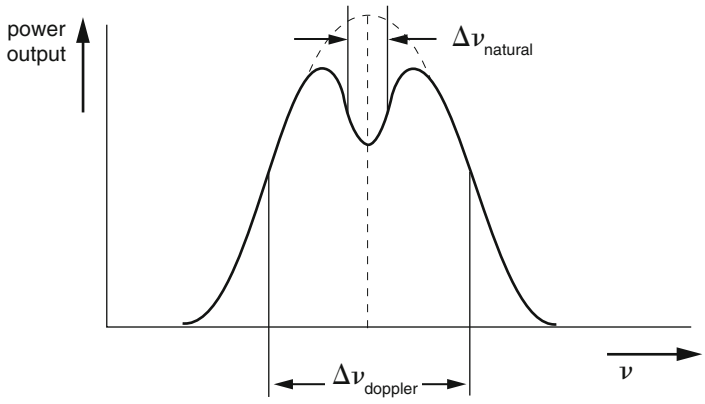

Figure 14.6 The Lamb dip in the output of a gas laser

14.1.5 Theoretical Limit to Spectral Purity of Lasers

From the point of view of an optical frequency standard, a laser oscillator serves to provide a strong, spectrally pure source of radiation, much as a klystron might do for a passive microwave Cs standard. Long-term stability and reproducibility are achieved by locking the laser frequency on resonance with a suitable reference atomic or molecular quantum transition, free of Doppler and other sources of spectral line distortion. In this role the essential attribute of the laser is the spectral purity of its output. In practice, this is broadened by the fluctuations in the optical and mechanical properties of the cavity, particularly those due to environmental conditions, such as temperature and mechanical vibrations. We must distinguish, however, between these "technical" or "artificial" sources of phase/frequency fluctuations and those that are fundamental, that is, those that arise from the quantum properties of radiation and its interaction with atoms. These residual fluctuations would remain even if we had an ideal, perfectly stable cavity.

To understand the origin of this inherent, fundamental limit on the spectral purity of the laser output, and therefore the limit on the phase stability of the laser as a frequency standard, we must go back to the fundamental processes involved in its operation. There are two light-emission processes that atoms of the laser medium undergo: spontaneous and stimulated emission. In spontaneous emission, which occurs with a probability independent of the prior presence or absence of photons, the photons emitted by different atoms bear no phase relationship to each other, nor do photons emitted by the same atom at different times. In contrast, the stimulated emission of photons occurs with a probability proportional to the number of interacting photons already present, and the phases of photons emitted by different atoms, or by the same atom at different times, have a definite relationship; that is, they are coherent. It is the inevitable presence of spontaneously emitted, incoherent photons in the otherwise coherent stream of photons constituting the laser output beam that sets the limit on spectral purity mentioned above. Quantitatively, it can be shown that the mean square deviation in phase $\langle \Delta\phi^2 \rangle$ is given by

$$\left\langle \Delta\phi^2 \right\rangle = \frac{N_{\text{spont}}}{2N_{\text{tot}}}, \tag{14.19}$$

where the average $\langle\ \rangle$ is taken over a time during which N_{spont} photons are spontaneously emitted, and N_{tot} is the total number of photons in the given field mode. This result can be made plausible by noting that the ratio of photon numbers is proportional to $(E_{\text{spont}}/E_{\text{tot}})^2$, where E_{spont} and E_{tot} are the corresponding optical field amplitudes. That is, in terms of a field picture we have an oscillating optical field vector with its phase randomly fluctuating over a narrow range because of the addition of a (small) phase incoherent field component (the spontaneous photon). The size of the phase shift we can deduce if we recall further that the phase of a harmonically oscillating quantity $E\cos(\omega t + \phi)$ can be represented as the angle of rotation of a radius vector of length E turning with constant angular velocity ω, and that

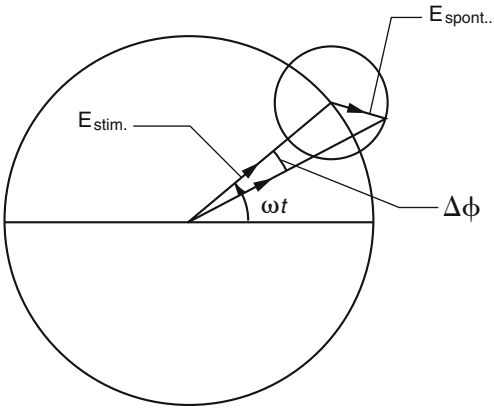

Figure 14.7 The phasor representation of the stimulated and spontaneous optical fields

therefore the phase of the resultant of two fields oscillating at the same frequency can be obtained by vector addition of the two rotating vectors representing them, as shown in Figure 14.7. Now, under steady oscillating conditions where the population inversion is sustained by pumping at a constant rate, the mean optical field amplitude E_{tot} remains constant, and in cases of practical interest $E_{\text{spont}} \ll E_{\text{tot}}$. From the figure it is clear that the maximum change $\Delta\phi$ in the phase of the optical field due to the addition of a small vector increment occurs when the phase of the latter is at 90° to the main field vector. It follows that small fluctuations in the amplitude E_{spont} can produce a maximum phase change (in radians) given by $\Delta\phi = E_{\text{spont}}/E_{\text{tot}}$. Of course, the effect of the spontaneous component varies randomly, sometimes advancing the phase of the resultant, at other times retarding it. The situation will be recognized as reminiscent of a *random walk*, of which we have already given a simple model in an earlier chapter. As we saw there, while the average of the fluctuations is zero, being equally likely to be positive as negative, the average of the *square* of the fluctuations increases linearly with their number. In this case each spontaneously emitted photon corresponds to a new fluctuation in the phase, and since the emission occurs at a constant rate, we conclude that $\langle \Delta\phi^2 \rangle$ increases linearly with time. It can be shown that this leads to a laser output with a Lorentzian spectral intensity distribution with a spectral line width $\Delta\nu$ given by

$$\Delta\nu = \frac{\pi h \nu (\Delta\nu_c)^2}{P}, \tag{14.20}$$

where $\Delta\nu_c$ is the *passive* cavity resonance linewidth (in the absence of the lasing medium) and P is the power in the cavity mode. It was on the basis of an expression of this form, first derived in 1958 by Townes and Schawlow, prior to the realization of a working laser, that the extraordinary potential spectral purity of lasers was predicted. Assume, for example, that we have a laser operating at 633 nanometers

with an ideal cavity of length $L = 1$ m and an output mirror with 1% transmission, so that 1% of the cavity intensity emerges as the output beam. The cavity resonance line width can be obtained from the average lifetime of a photon in the cavity. Thus a given photon has a 1% chance of leaving the cavity in the time required to traverse the cavity in both directions, and will therefore spend on the average 200 L/c before leaving the cavity. The corresponding (full) spectral width of the cavity resonance is then $\Delta\nu_c = c/\pi 200L$; that is, in this case $\approx$0.5 MHz. For a laser output power of 1 mW we find on substituting into the expression for the laser spectral line width $\Delta\nu = 2.5 \times 10^{-4}$ Hz! This quantum limit is so small that it was thought at the time it was first calculated that it was of no practical consequence; however, as we shall see, recent work on laser stabilization has led to claims of extraordinary spectral purity, approaching the quantum limit.

14.1.6 Laser Stabilization: The Pound-Drever-Hall Method

An experimental technique that has played a major role in stabilizing lasers with respect to ultrastable cavities, and thereby brought within reach the realization of the quantum limit to spectral purity in a laser, is known as the Pound-Drever-Hall method (Drever et al., 1983). It involves phase modulation and feed-back to lock the frequency of a laser to a resonant frequency in an ultrastable external cavity of extremely high finesse. We recall that lasers such as tunable dye lasers and solid state lasers have a relatively broad spectral width and require frequency narrowing and stabilization arrangements to achieve spectral purity. The Pound-Drever-Hall method has been applied with great success in standards laboratories to lock lasers to high finesse cavities, isolated with extraordinary care to ensure freedom from vibration and temperature fluctuations. Such stabilized lasers, destined to be local oscillators for optical standards, have been reported to reach spectral linewidths in the sub-Hertz range, corresponding to fractional linewidths in the 10^{-16} range, a truly astounding achievement.

The principle of the method dates back to 1946, a time when microwave technology was the hot field of the day. One of the most prominent experimentalists in that field and in nuclear magnetic resonance was R.V. Pound.

The "Pound stabilizer" (Pound, 1946) was designed to stabilize the output frequency of a microwave oscillator with reference to a high-Q microwave cavity. It phase modulates the microwave output of the oscillator and couples it to a resonant cavity in one arm of a "magic-T". The error signal for the feed-back loop was obtained by first detecting the sum of the two sidebands of the phase modulated signal and a sample of the carrier signal from the source, and connecting this to a phase sensitive detector having the (phase adjusted) modulating signal as reference. The output of the latter can be shown to give a linear measure of the detuning error in the neighborhood of exact resonance of the cavity.

14.1.7 Application to Stabilizing a Laser

It might be thought that one could stabilize the frequency of a laser simply by locking the laser to one of the sharp transmission peaks of a stable, high finesse Fabry-Pérot cavity, in a fashion analogous to locking a local oscillator to an atomic resonance. This would certainly work; however, the modulation frequency and therefore the bandwidth of the servo loop would be limited by the response time of the cavity. This would limit the ability of the system to reduce higher frequency fluctuations in the laser output.

The principles of operation of the Pound-Drever-Hall method can be understood with reference to the simplified diagram shown in Figure 14.8. The laser is first isolated with a Faraday isolator and then its beam is phase modulated in a Pockels cell before going through an isolator/reflector into the optical cavity. The Faraday isolator is essential to prevent any reflected wave from destabilizing the laser. The intensity of the beam reflected from the optical isolator is measured in a detector whose output is connected together with a signal from the modulating oscillator, properly phase-adjusted, to a mixer, yielding a phase-sensitive output that passes through a low pass filter to control the frequency of the laser. The distinctive feature in this servo control system is clearly the way the error signal is derived from the *reflected* beam. As already emphasized, what is particularly notable is that the modulation frequency and therefore the bandwidth of the servo control are not limited by the bandwidth of the cavity, that is, the speed with which the optical field in the cavity can readjust.

We will not reproduce the general theory of the method, but will be content with approximate results limited to conditions of practical importance, namely fast modulation in the neighborhood of resonance with a cavity mode. To begin we need an approximate expression for the reflection coefficient of a Fabry-Pérot cavity

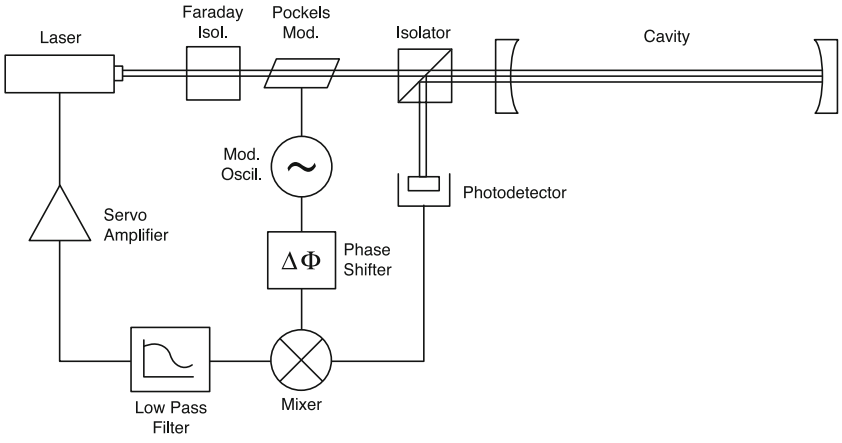

Figure 14.8 Simplified layout of Pound-Drever Hall laser stabilization method

near resonance. We give without proof the following expression for the reflection coefficient $R(\omega)$ using complex vectors:

$$R(\omega) = \frac{r[\exp(i\omega/\Delta\omega) - 1]}{1 - r^2 \exp(i\omega/\Delta\omega)} \tag{14.21}$$

where $\Delta\omega$ is the free spectral range $c/2n_r L$ and r is the *amplitude* reflection coefficient of each mirror. In the neighborhood of the N^{th} longitudinal mode we can write

$$\frac{\omega}{\Delta\omega} = 2N\pi + \frac{\delta\omega}{\Delta\omega} \tag{14.22}$$

where $\delta\omega/\Delta\omega \ll 1$ is the fractional deviation in the frequency of the laser from the exact cavity resonance. On substituting this form in (14.21) we obtain to first order in $(\delta\omega/\Delta\omega)$:

$$R(\omega) \approx \frac{ir\left(\frac{\delta\omega}{\Delta\omega}\right)}{1 - r^2}. \tag{14.23}$$

Writing $R(\omega)$ in terms of the finesse $F = \pi r/(1 - r^2)$ and using the relationship $\Delta\omega_{1/2} = \Delta\omega/F$, where $\Delta\omega_{1/2}$ is the half width of the resonance, we find:

$$R(\omega) = \frac{i}{\pi} \frac{\delta\omega}{\Delta\omega_{1/2}} \tag{14.24}$$

If we assume the modulation frequency to be so large that the sidebands are so far from resonance that we may assume they are fully reflected, then it can be shown that the reflected intensity I_r is given by:

$$I_r \approx 2I_s - \frac{4}{\pi}\sqrt{I_c I_s} \frac{\delta\omega}{\Delta\omega_{1/2}} \sin\Omega t + (2\Omega terms) \tag{14.25}$$

where I_s and I_c are the sideband and carrier intensities, respectively. On passing through the phase sensitive detector with the reference frequency Ω, the output error signal applied to control the laser, ε, is given by:

$$\varepsilon \approx \frac{4}{\pi}\sqrt{I_c I_s} \frac{\delta\omega}{\Delta\omega_{1/2}} \tag{14.26}$$

This then shows the linear dependence of the error signal on the frequency deviation; as required, going through zero at exact resonance and changing sign from negative to positive on either side of resonance.

14.2 Laser Beam Properties

14.2.1 Laser Beam Quality

The cross sectional intensity profile and spatial coherence of a laser beam are essential properties which determine how it propagates through an optical system. A standard measure of *beam quality* is clearly required in order to be able to

compare the outputs of lasers of different designs and manufacture. The theoretical definition of such a standard measure of beam quality is based on a comparison between the given beam profile and the ideal lowest order Gaussian beam. The specification of just how the comparison should be made comes under the purview of the International Standards Organization (ISO) and the U.S National Institute for Standards and Technology (NIST).

The beam quality, generally represented by M^2, can be defined as the ratio of the product of the beam's multimode diameter times the divergence angle to the same product for an ideal diffraction-limited TEM_{00} mode Gaussian beam. Symbolically the definition can be written as follows:

$$M^2 = \frac{d_m \theta_m}{d_0 \theta_0} \tag{14.27}$$

A theoretically ideal TEM_{00} beam has by definition the minimum value $M^2 = 1$; therefore the degree to which a given beam has a value approaching this value is a measure of its high quality.

Another useful parameter is the *Rayleigh range*. This is the distance along a beam from the point of minimum diameter, *the waist,* at which the diameter has increased by a factor of $\sqrt{2}$.

14.2.2 Mode-Locked Lasers

While the construction of an optical frequency standard requires that a laser be forced to oscillate in a single mode in order to attain the high spectral purity, we shall see that the "clockwork" of an optical clock requires the opposite extreme of a very wide frequency multi-mode laser, such as a dye laser or Ti:sapphire laser. In general, a laser cavity will support longitudinal modes spanning a frequency band as wide as the reflectivity of the mirrors will allow, separated by the frequency interval $\Delta\nu = c/2n_r L$. The laser may oscillate in those modes that fall within the amplification frequency band of the laser medium, unless special optical elements are present to narrow that band in order to force single mode operation. If the gain band-width reflects inhomogeneous broadening, in which the frequency dependence of the gain is due to the distribution of frequencies among *different* particles participating in the optical transitions, then clearly the laser can oscillate simultaneously in as many longitudinal modes as have sufficient gain. In this case the onset of oscillation at the frequency of maximum gain involves only a subset of the particles, leaving others with lesser gain to sustain oscillation in other longitudinal modes.

The *spectrum* of the output beam will therefore consist of an array of equally spaced lines, which may or may not be resolved, depending on the optical Q-factor of the cavity and the resolving power of the spectrum analyzer, among other things. However, the time dependence of the output intensity would appear to fluctuate randomly about a certain mean value because the longitudinal modes oscillate with

uncorrelated phase relative to each other, leading to random interference between them. On this account the *coherence* of the laser beam would be degraded, that is, two such beams would have a very limited ability to produce an interference pattern. To produce a coherent beam requires that the relative phases of the different modes be made constant, a process called *mode locking*. There are many ways in practice in which this may be achieved, and we will briefly touch on two of these in what follows.

First let us consider the consequence of mode locking on the time dependence of the output beam intensity. For simplicity assume that the optical field in the cavity is a superposition of N longitudinal modes, all having a common phase φ, then if we assume for simplicity that they all have equal amplitude, we have for the resultant:

$$E(t) = e^{i\varphi} \sum_{-(N-1)/2}^{(N-1)/2} E_n \exp[2\pi i(\nu_0 + n\Delta\nu)t] \tag{14.28}$$

The desired time dependence of the intensity is given by $E_0(t).E_0^*(t)$ averaged over a period long compared with $1/\nu_0$ but short compared with $1/\Delta\nu$, which yields the following:

$$I(t) = \frac{I_0}{N} \frac{\sin^2(\pi N \Delta\nu t)}{\sin^2(\pi \Delta\nu t)} \tag{14.29}$$

It follows from this that I(t) is zero at constant intervals of $T = 1/N\Delta\nu$ except for those times, which occur at intervals of $1/\Delta\nu$, when the denominator is also zero. For large N the result is a train of equally spaced sharp pulses of intensity NI_0 and duration on the order of $1/N\Delta\nu$ recurring every $1/\Delta\nu = 2n_r L/c$ seconds, that is the time it takes a photon to travel back and forth the length of the cavity, as shown in Figure 14.9. We note that $N\Delta\nu$ cannot exceed the total gain band width, since obviously a mode is sustained only if it falls within the gain bandwidth, but may approach its limits in practice. The striking fact to be drawn from this is that the duration of the pulses depends on the gain bandwidth of the active medium and associated optics of the laser. This gain bandwidth is as much as 3×10^{12} Hz in Nd^{3+}:glass, corresponding to extremely short pulses lasting 3.3×10^{-13} sec each, in the *femtosecond* (10^{-15} sec) range.

So far we have assumed that the active medium has an inhomogeneously broadened spectrum, in which different particles may be assumed to sustain their own modes. However in a homogeneously broadened medium, *all particles* are assumed to participate in sustaining the *same* mode and would suppress the others by reducing the population inversion for them. In fact multimode oscillation *can* be generated by intracavity phase or absorption modulation techniques to give rise to outputs similar to those from mode-locked lasers with inhomogeneous active media.

There are several approaches to mode locking a laser: perhaps the most popular are first by intracavity modulation of the optical gain or phase, and second, with

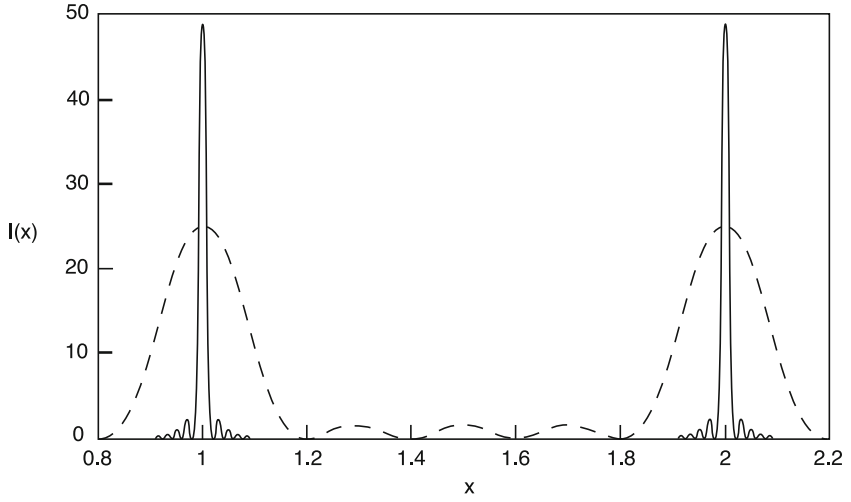

Figure 14.9 The function $\sin^2(N\pi x)/\sin^2(\pi x)$ for $N = 50$ (solid line reduced by factor 50) and $N = 5$ (dotted line)

a saturable absorber. In the first, the modulation frequency of an intracavity absorber is set to equal the frequency interval between consecutive longitudinal modes. The effect on each mode of oscillation is to generate side bands that coincide with adjacent modes and are phase coherent with each other at the value set by the modulator. If the modulation is sinusoidal only one side band is generated on each side of the carrier, but these will in turn produce more side bands at twice the modulation frequency, and so on, until all the modes within the gain bandwidth oscillate locked in phase. A detailed theory is complicated, but it helps to visualize the optical field inside the cavity in the mode-locked condition as consisting of a narrow high amplitude pulse traveling back and forth the length of the cavity. If the modulated absorber placed in the cavity is imagined to be in the shape of a thin plate placed perpendicular to the axis, then it is clear that the field configuration that is favored as having the least loss is a pulse that passes through the absorber always when it has minimum absorption.

There is another so called *passive method* of mode-locking a laser: it is through the use of an absorber that saturates, that is, becomes *less absorbent* as the intensity of light passing through it is increased beyond a certain point. This technique applies mainly to high power lasers. A simplified explanation of the physical process by which mode-locking occurs in this case begins with considering the laser as it passes through the threshold of oscillation, when the optical field with its natural random fluctuations is being established. Because of the saturation of the absorber, any strong fluctuation will "bleach" the absorber and experience a *diminished* absorption leading to stronger amplification than the average intensity.

Again this enhanced pulse will travel back and forth resulting in high peak power output pulses. It is assumed that the absorber has sufficient time to recover between pulses.

14.3 Laser Optical Elements

14.3.1 Multilayer Dielectric Mirrors and Filters

The components of a laser optical system, such as mirrors, windows, lenses, etc. are characterized by the smallness of their size compared with classical optics, and the high precision of their construction to preserve and exploit the coherence properties of the laser light. An "optically flat" surface will typically be specified as $\lambda/20$, meaning that at all points on the surface the mean departure from a geometric plane is less than one-twentieth of an optical wavelength. Furthermore, laser-grade optical surfaces have a higher degree of polish as specified by the "scratch and dig" figures, which indicate the "visibility" and number of scratches and pits in the polished surface.

In classical optics, mirrors were almost universally made by depositing a film of silver onto the desired surface either from a chemical solution (Rochelle process) or more commonly now by deposition of silver vapor in vacuum. For laser light much higher values of reflectance have been achieved using a radically different approach made possible by the fact that laser light is nearly monochromatic, and therefore the reflectance needs to be (and in some cases is preferred to be) high only for a very narrow wavelength range. The new mirrors are called *multilayer dielectric* mirrors, formed by vapor deposition onto an optically flat substrate (usually quartz or sapphire) of many thin layers of highly transparent dielectric materials, with values of refractive index alternating between high and low values. The principle underlying this type of mirror is that of superposition of light waves and the phenomenon of interference. Suppose we have a set of plane parallel films of alternating refractive indices n_1 and n_2, and let a monochromatic beam of light fall perpendicularly on them. To find the reflectance of such an arrangement we recall Fresnel's formulas for the reflection and refraction of light waves at boundary surfaces between different media. Originally derived on the basis of the "ether vibrations" theory of light, which predates Maxwell's electromagnetic theory, these formulas, with some reinterpretation, remain valid. For the particular case of a light wave in a medium with refractive index n_1 falling perpendicularly on the boundary surface with a medium of refractive index n_2, the (amplitude) reflectance for such a light wave is given by

$$r = \frac{n_1 - n_2}{n_1 + n_2}. \tag{14.30}$$

We note that if $n_1 < n_2$, the wave suffers a change of phase of 180° (a change of sign), whereas if $n_1 > n_2$, there is zero change in phase. In a multilayer dielectric mirror, the thicknesses of the films are chosen such that it takes half a period of

oscillation for the wave to traverse the thickness of any film in both directions. That is, since the velocity of light is c/n_1 and c/n_2 in the two media, we require the following:

$$\frac{2d_1n_1}{c} = \frac{2d_2n_2}{c} = \frac{\tau}{2}, \tag{14.31}$$

where d_1 and d_2 are the layer thicknesses, and $\tau = 1/\nu$ is the period of the light wave.

Now, referring to Figure 14.10 we see that light waves reflected from any boundary surface are in phase with those reflected from any other surface; that is, there is constructive interference. To derive the reflectance of a large number of dielectric layers one makes use of the boundary conditions that the total electric and magnetic components of the light wave on either side of any boundary must obey. By applying these conditions of continuity and allowing for the transit delays between boundaries, it is possible to relate the field components at one boundary to those at succeeding boundaries. The result of such an analysis we give without proof; for ideal nonabsorbing dielectrics the reflectivity is as follows:

$$r = \frac{\left(\frac{n_2}{n_1}\right)^{2N} - 1}{\left(\frac{n_2}{n_1}\right)^{2N} + 1} \tag{14.32}$$

For example, if magnesium fluoride ($n = 1.35$) and zinc sulphide ($n = 2.36$) are used in a 14-layer ($N = 7$) mirror, the intensity reflectance $r^2 = 99.8\%$. This high value, it must be emphasized, obtains only for light of a single wavelength, one for which the optical thickness of each layer is one-fourth the wavelength. Of course, in reality, the achievable reflectance is ultimately limited not only by absorption in the media, but also by scattering from irregularities in the boundary surfaces and within the media. Currently, there are commercially available mirrors fabricated using the most advanced polishing and coating technology that are claimed to have a reflectance as high as 99.99%. Of course, a few particles of dust could easily put that last decimal place in question!

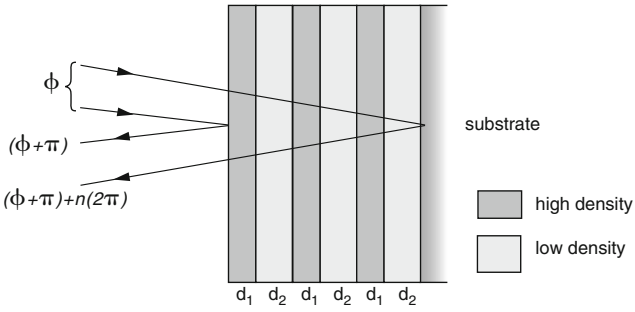

Figure 14.10 A section of a multilayer dielectric high-reflectance mirror

If a smaller number of layers is used, for example $N = 5$, we obtain $r^2 = 98.4\%$; hence, since the absorption and scattering are assumed negligible, we have a partially transparent mirror with a transmittance of $t^2 = 1-r^2$; that is, $t^2 = 1.6\%$. Such a mirror would be useful as the output mirror in an optical cavity.

Another type of multilayer mirror that has become an extremely important element in the design of optics handling fast pulses is the *chirped mirror*. In this the layers are not of equal thickness but rather vary uniformly so that a monochromatic incident beam will penetrate to the layers where the phase condition (14.20) is satisfied. This means that if two beams of slightly different wavelength are reflected by this mirror there would be a slight difference in the distance traversed, and hence in their relative phase. But a light pulse can be Fourier analyzed as a superposition of continuous waves, and varying the relative phases of these waves amounts to changing the shape of the pulse. With proper design the effect can be made one of sharpening a light pulse in time. These chirped mirrors have become common in systems designed to produce wide band frequency combs, as we shall see in Chapter 16.

14.3.2 Polarizing Optics

Another optical element that is indispensable in manipulating the spectral distribution of light is the *interference filter*. In contrast to the multilayer dielectric mirror, it is high transmittance rather than reflectance that the filter is designed to have, and this in a wavelength range as well defined as possible. The simplest bandpass filter is really a Fabry–Pérot cavity with the space between the parallel mirrors filled by a dielectric layer with an optical thickness nL equal to half the wavelength at the center of the desired band. The two mirrors are commonly in the form of the multilayer dielectric type described above, which, combined with the half-wavelength layer between them, form one integral unit.

Another optical element that has become indispensable in laser optics is the *Brewster window*, an optically flat transparent plate with parallel faces set at the Brewster, or *polarizing*, angle to the direction of light falling on it. We recall that a light wave, being transverse, can be polarized so that, for example, the electric field oscillates in one plane all along the wave. While polarization effects are generally associated with crystals such as calcite, it has been known since Malus that light can be polarized simply by reflection, a fact easily confirmed now by looking at sunlight reflected from water through polarizing sunglasses. In 1812 Brewster discovered experimentally the law that bears his name, giving quantitatively the angle at which light reflected from a dielectric is completely polarized. Brewster's law states that when a light wave is incident upon the surface of a dielectric medium, the amplitude of the component in the reflected wave whose electric field is parallel to the plane of incidence will be zero when the angle of incidence θ_B satisfies the condition

$$\tan\theta_B = \frac{n_2}{n_1}. \tag{14.33}$$

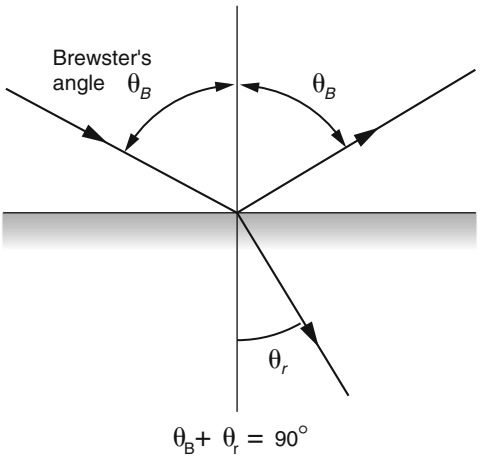

Figure 14.11 The angle relationships at the Brewster polarizing angle

For an air–glass interface $n_2/n_1 \approx 1.5$; thus $\theta_B \approx 56.3°$. If we recall Snell's law, $n_1 \sin\theta_1 = n_2 \sin\theta_2$, we find that for an angle of incidence equal to the Brewster angle, $\sin\theta_2 = \cos\theta_1$, and the refracted beam is at right angles to the reflected one, as shown in Figure 14.11. It follows from Brewster's law that if an incident beam is polarized parallel to the plane of incidence and it falls on the boundary surface at the Brewster angle, there will ideally be no reflected beam. The use of Brewster windows to avoid reflection losses is particularly important in low gain gas lasers such the helium–neon laser. In this case, if the plasma tube is to be separate from the optical cavity, its ends would be sealed with precisely oriented Brewster windows. When such a polarization-sensitive element is incorporated into a laser cavity, the result is that oscillation will take place with the field polarized in the direction having the least loss.

14.3.3 Nonlinear Crystals

The class of nonlinear devices useful beyond 30 THz ($\lambda = 10\,\mu$m) includes certain optical crystals that lack a center of symmetry. In these crystals the electric field component of the light wave induces an oscillating electric dipole moment that has a small quadratic dependence on the field amplitude. Moreover, the dipole moment in a given direction depends not only on the field component in that direction, but also on the components in other directions, reflecting the anisotropy of a crystalline medium. The nonlinear behavior is expected to be weak in general, since even in a 1 megawatt laser beam the electric field amplitude is on the order of $E \approx 10^6$ volts/m, whereas the interatomic fields in the crystal are on the order of 10,000 times that. This means that in order to gain a large cumulative effect, the light wave must interact with the crystal over large distances, containing many

wavelengths, making the crystal fall in the class of distributed devices. This introduces a requirement on the velocities of light waves of different frequencies in the crystal medium, since any sum or difference frequency wave generated at points along the path of the input waves must reinforce waves generated at subsequent points, in order that there be overall buildup of the mixed frequencies. This means ideally that all frequencies must travel at the same velocity. In practice, it is only necessary that any phase difference that develops between waves generated at the beginning and those at the end of the finite path in the crystal be less than π radians. For example, let us consider what this means in the case of second harmonic generation: Let n_ν and $n_{2\nu}$ be the refractive indices of the crystal for the fundamental and second harmonic frequencies. Now, the second harmonic component of the electric polarization of the crystal, which acts as the source of the second harmonic wave, travels at the velocity of the fundamental wave, and hence the condition on the phase difference that can be allowed to develop over a distance L in the crystal, $\Delta\phi$, is the following:

$$\Delta\phi = 2\pi(2\nu)\left[\frac{n_\nu L}{c} - \frac{n_{2\nu}L}{c}\right] \leq \pi. \tag{14.34}$$

This is called the *phase matching condition*. If it is violated, then the contributions to the second harmonic wave originating from different points along the path of the fundamental wave will combine in opposing phases, resulting in destructive interference.

Regarded from the photon point of view, the process amounts to the conversion of two photons at the fundamental frequency into a single photon of twice the frequency. Looked at this way, the phase matching condition becomes a statement of the conservation of photon linear momentum h/λ, which in quantum theory allows some discrepancy, provided that it is within the Heisenberg uncertainty in momentum. We recall that this uncertainty is related to the uncertainty in distance in the same way that the uncertainty in frequency involves the time measurement. If L is in effect the uncertainty in the positions of the photons, then the Heisenberg uncertainty in momentum is on the order of $h/2L$; it follows that for momentum conservation we require

$$2\frac{hn_\nu\nu}{c} - \frac{hn_{2\nu}2\nu}{c} \leq \frac{h}{2L}, \tag{14.35}$$

which is clearly the same as the phase matching condition.

Unfortunately, in practice, all transparent media are to some extent dispersive; that is, the wave velocity, and therefore the refractive index, varies with the frequency. It can happen that a specific material may be found whose refractive index takes the same value at widely different frequencies, allowing phase matching at those frequencies. But this is rarely the case. A far more practical approach is to exploit the birefringence (double refraction) of certain crystals. Without getting too deeply involved in crystal optics, we will assert that there are crystals of certain classes of symmetry whose optical properties have symmetry about a single axis, the *optic axis*, and are called *uniaxial.* These crystals have a 3-fold, 4-fold,

or 6-fold axis of symmetry, and their optical behavior is symmetric about these axes. Other crystals of a lower degree of symmetry have *two* preferred axes and are *biaxial.* In a uniaxial crystal, a light wave propagates along two wavefronts with the velocity of one, the *extraordinary* wave, depending on the direction relative to the optic axis. The other, the *ordinary* wavefront, advances with equal velocity in all directions. Along the optic axis both wavefronts advance with the same velocity. The directions of the electric (and magnetic) components of the two types of waves, that is, their polarization vectors, are always at right angles to each other. The velocity of the extraordinary wavefront may increase or decrease as a function of the angle with respect to the optic axis, depending on the specific material of the crystal, as shown in Figure 14.12. In such a crystal, phase matching is achieved when the phase velocity of the ordinary wavefront at one frequency matches the phase velocity of the extraordinary wavefront at the other frequency. Since the difference in phase velocity between the two varies continuously with the angle the laser beam makes with the optic axis, it is possible in principle to match phases with a suitable crystal. In practice, this requires a beam with extremely small divergence, critically adjusted to the correct angle with respect to the crystal axes. An approach less critically sensitive to beam adjustment is to find a suitable crystal in which the wavefront velocities can be adjusted by varying the temperature, while fixing the beam direction at right angles to the axis, where the velocity difference does not change (to first order) with angle Just how critical the matching of phase velocities (or equivalently, refractive indices) is in practice can be estimated from the phase matching equation $(n_{2\nu} - n_\nu) \leq \lambda/4L$, where λ is the wavelength of the fundamental frequency wave. Since the degree of nonlinearity is relatively

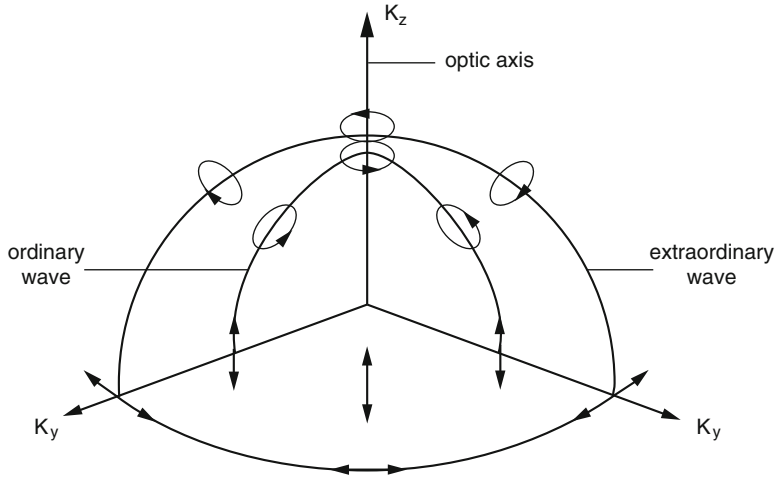

Figure 14.12 Ordinary and extraordinary wave vectors and polarizations in a uniaxial crystal, such as quartz

small, L is on the order of a centimeter; hence for a λ in the middle of the optical range, the difference in refractive index must not exceed one part in 10^5.

Of the many birefringent crystals that show nonlinearity, those few that have an adequate nonlinear coefficient, are transparent in the desired wavelength range, and are resistant to surface damage are useful as frequency mixers. Quartz has already been mentioned as the crystal in which the doubling of an optical frequency was first observed. Other crystals with greater nonlinearity that have been used include potassium dihydrogen phosphate (KDP) and ammonium dihydrogen phosphate (ADP), which are adequate for harmonic generation of 1:m infrared radiation. The widely used lithium niobate ($LiNbO_3$) crystal has a nonlinearity ten times that of KDP; potassium niobate ($KNbO_3$) has an even larger nonlinearity, providing phase matching into the blue part of the spectrum; and lithium iodate ($LiIO_3$) provides a further extension into the ultraviolet.

14.3.4 Diffraction Devices

Whenever a light wave propagates through media whose propagation characteristics vary spatially at a rate comparable to its wavelength, we observe the phenomenon of *diffraction.* The variation may take the form of high frequency acoustic vibrations in a crystal producing alternately high and low refractive index, or it may be ripples on the boundary surface between two media of differing refractive index, as in a diffraction grating. Diffraction is characterized by a redistribution in space of the light energy as determined by interference between waves scattered from different points in the source of the diffraction.

In classical optics the common wavelength-dispersive element in a monochrometer is a diffraction grating, a plane or concave mirror ruled with fine, closely spaced parallel grooves, thereby forming a periodically varying reflectivity. If the spacing of the grooves is d on a plane reflection grating, and light falls on it along the normal, then the grating equation relating the direction θ_n in which the spectral component of the light having the wavelength λ has a maximum reflection is as follows:

$$d \sin\theta_n = n\lambda \tag{14.36}$$

where n is an integer, the *order* of the diffraction. If the incident light is monochromatic then we can show that the reflected intensity as a function of θ is distributed over narrow ranges of θ centered on the values obeying the grating equation. Thus, writing

$$\delta = \frac{\pi d \sin\theta}{\lambda} \tag{14.37}$$

the sum of the partial waves reflected from the grating (applying Huygens Principle) is given by:

$$E(t) = e^{i\omega t} \sum_{n=0}^{n=N-1} E_n e^{in\delta} \tag{14.38}$$

which leads to:

$$I(\theta) = \frac{I_0}{N}\frac{\sin^2(N\delta)}{\sin^2\delta} \qquad 14.39$$

where I_0 is the incident intensity over the entire N grooves, in the absence of interference. It follows that $I(\theta)$ reaches a large maximum NI_0 whenever $\delta = n\pi$, that is, when the grating equation is satisfied. Furthermore the larger N is, the closer the zeros of $I(\theta)$ at $\delta = (n+1/N)\pi$ will be to the maximum at $\delta = n\pi$, that is the sharper the spectral resolution of the grating.

Of particular importance in the context of laser applications are acousto-optic deflectors and modulators. The diffraction of light by acoustic waves in certain crystals was predicted by Brillouin in 1921. The basis of the acousto-optic interaction is fundamentally through the change in electrical permittivity that results from a mechanical strain or deformation in the crystal. In the case of the acousto-optic effect of interest to us, the strain is a periodic function of space produced by an acoustic wave in the medium. For a plane acoustic wave propagating in a direction we will choose as the z-axis we have:

$$n(z,t) = n + \Delta n \cos[2\pi \nu_a t - k_a z] \qquad 14.40$$

where ν_a, k_a are the frequency and wave vector of the acoustic wave, and Δn is the amplitude of modulation of the refractive index due to the acoustic wave. For an incident laser beam this represents a *phase grating* traveling with the velocity of sound in the medium, producing a diffraction pattern in the far field.

In describing the diffractive behavior we must distinguish two limiting regimes: the so-called *Raman-Nath* regime and the *Bragg* regime. The Raman-Nath regime is observed at relatively low acoustic frequencies and short acousto-optic interaction lengths. It can be observed at arbitrary angles of light incidence, and the diffracted beam contains many orders symmetrically distributed about the direction of incidence as illustrated in Figure 14.13. In contrast, the Bragg diffraction regime is observed at high acoustic frequencies exceeding usually 100 MHz. The observed diffraction pattern, even at high acoustic powers, consists usually of only two diffraction maxima: the zeroth and first order. Moreover these occur only at definite angles near the so-called Bragg angle, determined by the Bragg equation familiar from X-ray diffraction theory, which for an isotropic medium is given by:

$$2\frac{k_a}{2\pi}\sin\theta_B = \frac{\lambda}{n_r} \qquad 14.41$$

where n_r here is the refractive index of the medium.

If the medium is an anisotropic crystal, then we have to take into account the birefringence of the medium: two types of diffraction are possible. An acousto-optic interaction in which there is no change in the type of wave, whether "ordinary" or "extraordinary", will obey the Bragg condition given in eq. 14.41, using the appropriate refractive index. On the other hand if the interaction leads to a transition, so that the refractive index for the incident beam differs from that

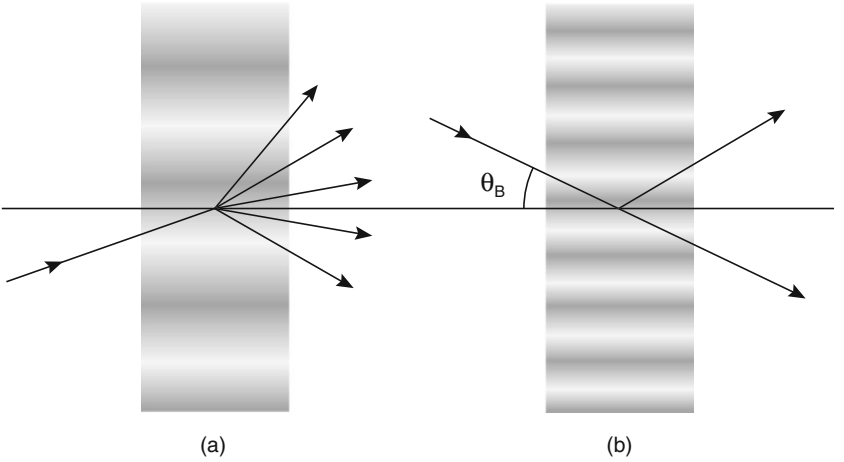

Figure 14.13 (a) Raman-Nath scattering (b) Bragg diffraction

for the diffracted beam, the theory is far more complicated; nevertheless it is of great practical importance. We will only state qualitatively that for the special case of the acousto-optic interaction plane being perpendicular to the optical axis of a uniaxial crystal, the functional dependence of the Bragg angle on the acoustic frequency differs according to the relative magnitudes of the refractive indices for the incident and diffracted beams The optimum choice of acousto-optic crystal type to perform given functions such as deflection, modulation and filtering, is dictated by the relative magnitudes of the incident and diffracted beams.