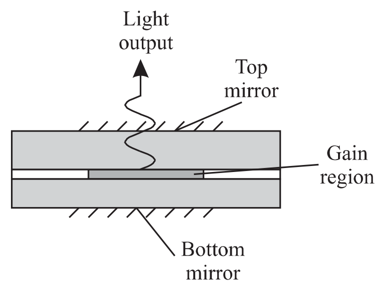

Figure 3.47 The structure of a VCSEL.

spacing between the modes of an MLM laser is $c/2nl$, where l is the length of the cavity and n is its refractive index. If we were to make the length of the cavity sufficiently small, the mode spacing increases such that only one longitudinal mode occurs within the gain bandwidth of the laser. It turns out that making a very thin active layer is much easier if the active layer is deposited on a semiconductor substrate, as illustrated in Figure 3.47. This leads to a vertical cavity with the mirrors being formed on the top and bottom surfaces of the semiconductor wafer. The laser output is also taken from one of these (usually top) surfaces. For these reasons, such lasers are called *vertical cavity surface-emitting lasers* (VCSELs). The other lasers that we have been discussing hitherto can thus be referred to as *edge-emitting lasers*.

Since the gain region has a very short length, very high mirror reflectivities are required in order for laser oscillation to occur. Such high mirror reflectivities are difficult to obtain with metallic surfaces. A stack of alternating low- and high-index dielectrics serves as a highly reflective, though wavelength-selective, mirror. The reflectivity of such a mirror is discussed in Problem 3.13. Such dielectric mirrors can be deposited at the time of fabrication of the laser.

One problem with VCSELs is the large ohmic resistance encountered by the injected current. This leads to considerable heating of the device and the need for efficient thermal cooling. Many of the dielectric materials used to make the mirrors have low thermal conductivity. So the use of such dielectric mirrors makes room temperature operation of VCSELs difficult to achieve since the heat generated by the device cannot be dissipated easily. For this reason, for several years after they were first demonstrated in 1979, VCSELs were not capable of operating at room temperature. However, significant research effort has been expended on new materials and techniques, VCSELs operating at 1.3 μm at room temperature have been demonstrated [Har00].

The advantages of VCSELs, compared to edge-emitting lasers, include simpler and more efficient fiber coupling, easier packaging and testing, and their ability

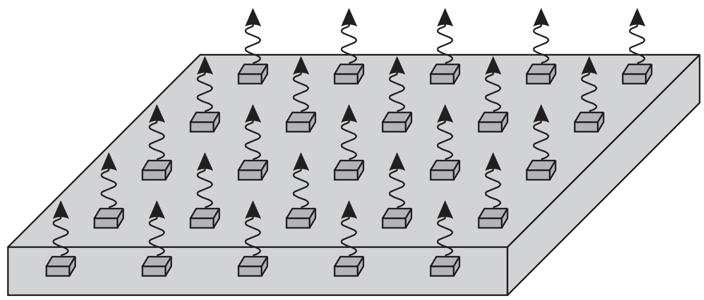

Figure 3.48 A two-dimensional array of vertical cavity surface-emitting lasers.

to be integrated into multiwavelength arrays. VCSELs operating at 0.85 μm are commercially available and used for low-cost, short-distance multimode fiber interconnections. In addition, 1.3 μm VCSELs have been commercially available.

In a WDM system, many wavelengths are transmitted simultaneously over each link. Usually, this requires a separate laser for each wavelength. The cost of the transmitters can be significantly reduced if all the lasers can be integrated on a single substrate. This is the main motivation for the development of arrayed lasers such as the DFB laser arrays that we discussed earlier. Moreover, an arrayed laser can be used as a tunable laser simply by turning on only the one required laser in the array. The use of surface-emitting lasers enables us to fabricate a two-dimensional array of lasers, as shown in Figure 3.48. Much higher array packing densities can be achieved using surface-emitting lasers than edge-emitting ones because of this added dimension. However, it is harder to couple light from the lasers in this array onto optical fiber since multiplexers that work conveniently with this two-dimensional geometry are not readily available. These arrayed lasers have the same yield problem as other arrayed laser structures; if one of the lasers does not meet specifications, the entire array will have to be discarded.

Mode-Locked Lasers

Mode-locked lasers are used to generate narrow optical pulses that are needed for the high-speed TDM systems that we will study in Chapter 12. Consider a Fabry-Perot laser that oscillates in N longitudinal modes, which are adjacent to each other. This means that if the wavelengths of the modes are $\lambda_0, \lambda_1, \ldots, \lambda_{N-1}$, the cavity length l satisfies $l = (k+i)\lambda_i/2$, $i = 0, 1, \ldots, N-1$, for some integer k. From this condition, it can be shown (see Problem 3.7) that the corresponding frequencies $f_0, f_1, \ldots, f_{N-1}$ of these modes must satisfy $f_i = f_0 + i\Delta f$, $i = 0, 1, \ldots, N-1$. The oscillation at

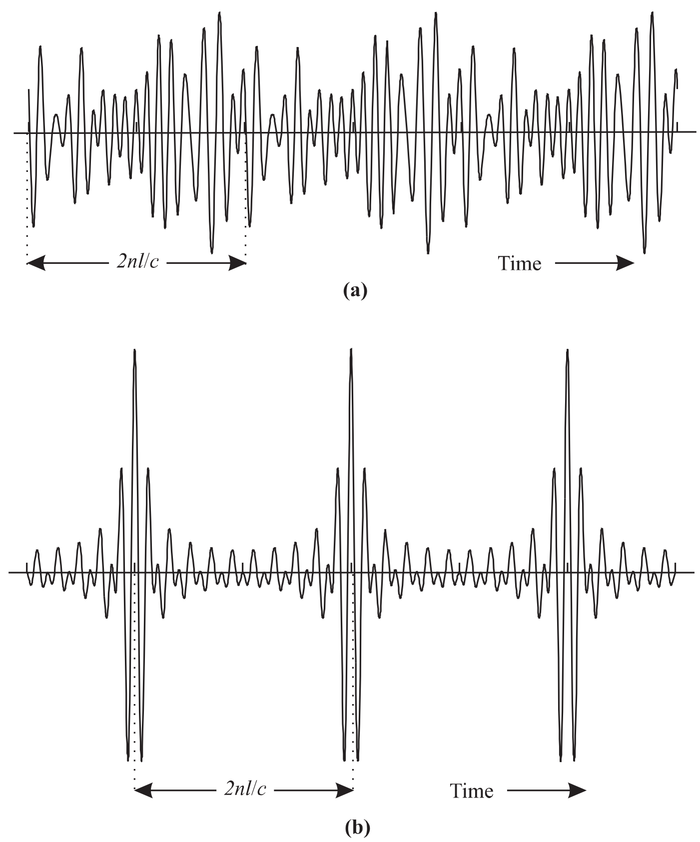

Figure 3.49 Output oscillation of a laser oscillating simultaneously in 10 longitudinal modes. (a) The phases of the modes are chosen at random. (b) All the phases are equal to each other; such a laser is said to be mode locked.

frequency f_i is of the form $a_i \cos(2\pi f_i t + \phi_i)$, where a_i is the amplitude and ϕ_i the phase of mode i. (Strictly speaking, this is the distribution in time of the electric field associated with the longitudinal mode.) Thus the total laser output oscillation takes the form

$$\sum_{i=0}^{N-1} a_i \cos(2\pi f_i t + \phi_i).$$

This expression is plotted in Figure 3.49 for $N = 10$, for different sets of values of the ϕ_i. In Figure 3.49(a), the ϕ_i are chosen at random, and in Figure 3.49(b), they are chosen to be equal to each other. All the a_i are chosen to be equal in both cases, and the frequency f_0 has been diminished from its typical value for the purpose of illustration.

From Figure 3.49(a), we observe that the output amplitude of an MLM laser varies rapidly with time when it is not mode locked. We have also seen in Figure 3.43(a) that the frequency spacing between adjacent longitudinal modes is $c/2nl$. If $n = 3$ and $l = 200\ \mu$m, which are typical values for semiconductor lasers, this frequency spacing is 250 GHz. Thus these amplitude fluctuations occur extremely rapidly (at a time scale on the order of a few picoseconds) and pose no problems for on-off modulation even at bit rates of a few tens of gigabits per second.

We see from Figure 3.49(b) that when the ϕ_i are chosen to be equal to each other, the output oscillation of the laser takes the form of a periodic train of narrow pulses. A laser operating in this manner is called a *mode-locked laser* and is the most common means of generating narrow optical pulses.

The time interval between two pulses of a mode-locked laser is $2nl/c$, as indicated in Figure 3.49(b). For a typical semiconductor laser, as we have seen earlier, this corresponds to a few picoseconds. For modulation in the 1–10 GHz range, the interpulse interval should be in the 0.1–1 ns range. Cavity lengths, l, of the order of 1–10 cm (assuming $n = 1.5$) are required in order to realize mode-locked lasers with interpulse intervals in this range. These large cavity lengths are easily obtained using fiber lasers, which require the length anyway to obtain sufficient gain to induce lasing.

The most common means of achieving mode lock is by modulating the gain of the laser cavity. Either amplitude or frequency modulation can be used. Mode locking using amplitude modulation is illustrated in Figure 3.50. The gain of the cavity is modulated with a period equal to the interpulse interval, namely, $2nl/c$. The amplitude of this modulation is chosen such that the average gain is insufficient for any single mode to oscillate. However, if a large number of modes are in phase, there can be a sufficient buildup in the energy inside the cavity for laser oscillation to occur at the instants of high gain, as illustrated in Figure 3.50.

Gain modulation of the fiber laser can be achieved by introducing an external modulator inside the cavity.

3.5.2 Light-Emitting Diodes

Lasers are expensive devices and are not affordable for many applications where the data rates are low and distances are short. This is the case in many data communications applications (see Chapter 6) and in some access networks (Chapter 11). In such cases, *light-emitting diodes* (LEDs) provide a cheaper alternative.

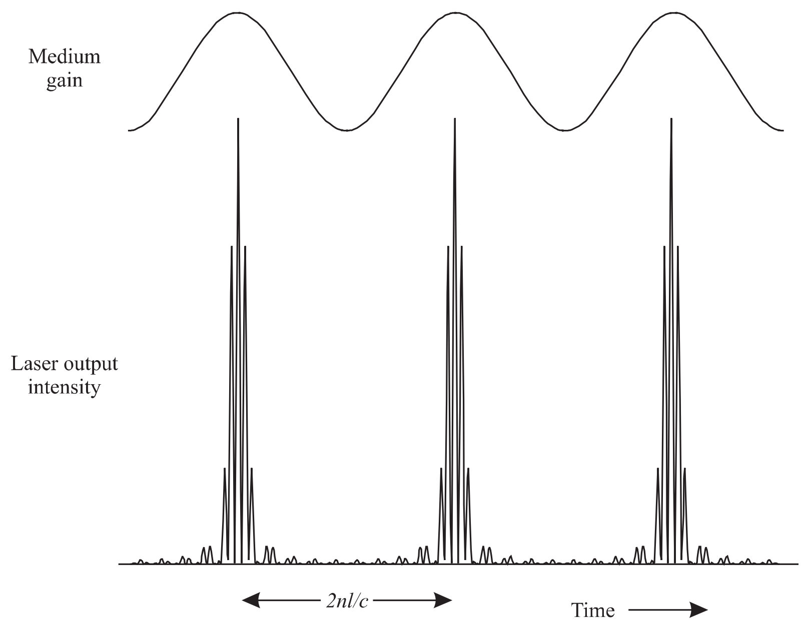

Figure 3.50 Illustration of mode locking by amplitude modulation of the cavity gain.

An LED is a forward-biased *pn*-junction in which the recombination of the injected minority carriers (electrons in the *p*-type region and holes in the *n*-type region) by the spontaneous emission process produces light. (Unwanted nonradiative recombination is also possible and is an important factor affecting the performance of LEDs.) Because spontaneous emission occurs within the entire bandwidth of the gain medium (corresponding to all energy differences between the valence and conduction bands for an LED), the light output of an LED has a broad spectrum, unlike that of a laser. We can crudely think of an LED as a laser with facets that are not very reflective. Increasing the pump current simply increases the spontaneous emission, and there is no chance to build up stimulated emission due to the poor reflectivity of the facets. For this reason, LEDs are also not capable of producing high-output powers like lasers, and typical output powers are on the order of −20 dBm. They cannot be directly modulated (see Section 3.5.4) at data rates higher than a few hundred megabits per second.

In some low-speed, low-budget applications, there is a requirement for a source with a narrow spectral width. DFB lasers provide narrow spectral widths but may be too expensive for these applications. In such cases, *LED slicing* provides a cheaper alternative. An LED slice is the output of a narrow passband optical filter placed in front of the LED. The optical filter selects a portion of the LED's output. Different filters can be used to select (almost) nonoverlapping spectral slices of the LED output. Thus one LED can be shared by a number of users. We will see an application for this technique in Chapter 11.

3.5.3 Tunable Lasers

Tunable lasers are highly desirable components for WDM networks for several reasons. Fixed-wavelength DFB lasers work very well for today's applications. However, each wavelength requires a different, unique laser. This implies that in order to supply a 100-channel WDM system, we need to stock 100 different laser types. The inventory and sparing issues associated with this are expensive and affect everybody from laser manufacturers to network operators. Laser manufacturers need to set up multiple production and test lines for each laser wavelength (or time-share the same production and test line but change the settings each time a different laser is made). Equipment suppliers need to stock these different lasers and keep inventories and spares for each wavelength. Finally, network operators need to stockpile spare wavelengths in the event transmitters fail in the field and need to be replaced. Having a tunable laser alleviates this problem dramatically.

Tunable lasers are also one of the key enablers of reconfigurable optical networks. They provide the flexibility to choose the transmit wavelength at the source of a lightpath. For instance, if we wanted to have a total of, say, four lightpaths starting at a node, we would equip that node with four tunable lasers. This would allow us to choose the four transmit wavelengths in an arbitrary manner. In contrast, if we were to use fixed-wavelength lasers, either we would have to preequip the node with a large number of lasers to cover all the possible wavelengths, or we would have to manually equip the appropriate wavelength as needed. We will see more of this application in Chapter 7. The tuning time required for such applications is on the order of milliseconds because the wavelength selection happens only at the times where the lightpath is set up, or when it needs to be rerouted in the event of a failure.

Another application for tunable lasers is in optical packet-switched networks, where data needs to be transmitted on different wavelengths on a packet-by-packet basis. These networks are primarily in their early stages of research today, but supporting such an application would require tuning times on the order of nanoseconds to microseconds, depending on the bit rate and packet size used.

Finally, tunable lasers are a staple in most WDM laboratories and test environments, where they are widely used for characterizing and testing various types of optical equipment. These lasers are typically tabletop-type devices and are not suitable for use in telecom applications, which call for compact, low-cost semiconductor lasers.

The InGaAsP/InP material used for most long-wavelength lasers is enhanced by the use of *quantum well* structures and has an overall gain bandwidth of about 250 nm at 1.55 μm, large enough for the needs of current WDM systems. However, the tuning mechanisms available potentially limit the tuning range to a small fraction of this number. The following tuning mechanisms are typically used:

- Injecting current into a semiconductor laser causes a change in the refractive index of the material, which in turn changes the lasing wavelength. This effect is fairly small—about a 0.5–2% change in the refractive index (and the wavelength) is possible. This effect can be used to effect a tuning range of approximately 10–15 nm in the 1.55 μm wavelength window.
- Temperature tuning is another possibility. The wavelength sensitivity of a semiconductor laser to temperature is approximately 0.1 nm/°C. In practice, the allowed range for temperature tuning is about 1 nm, corresponding to a 10°C temperature variation. Operating the laser at significantly higher temperatures than room temperature causes it to age rapidly, degrading its lifetime.
- Mechanical tuning can be used to provide a wide tunable range in lasers that use a separate external cavity mechanism. Many of these lasers tend to be bulky. We will look at one laser structure of this type using a micro-electro-mechanical tuning mechanism, which is quite compact.

As we will see, the tuning mechanisms are complex and, in many cases, interact with the modulation mechanisms, making it difficult to directly modulate most of the tunable lasers that we will study here.

The ideal tunable laser is a device that can tune rapidly over a wide continuous tuning range of over 100 nm. It should be stable over its lifetime and easily controllable and manufacturable. Many of the tunable laser technologies described here have been around for many years, but we are only now beginning to see commercially available devices due to the complexity of manufacturing and controlling these devices and solving the reliability challenges. The strong market demand for these devices has stimulated a renewed effort to solve these problems.

External Cavity Lasers

External cavity lasers can be tuned if the center wavelength of the grating or other wavelength-selective mirror used can be changed. Consider the grating external cavity laser shown in Figure 3.46. The wavelength selected by the grating for reflection

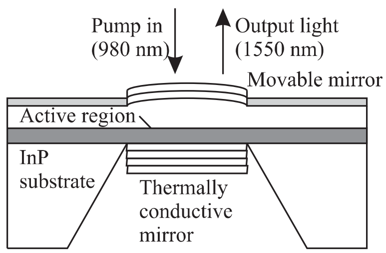

Figure 3.51 Structure of a tunable micro-electro-mechanical vertical cavity surface-emitting laser (MEM-VCSEL) (from [Vak99]).

to the gain cavity is determined by the pitch of the diffraction grating, its tilt angle with respect to the gain cavity, and its distance from the gain cavity (see Section 3.3.1, specifically, (3.9)). Thus by varying the tilt angle and the distance of the diffraction grating from the gain cavity (shown by the dotted arrows in Figure 3.46), the laser wavelength can be changed. This is a slow method of tuning since the tilt and position of the diffraction grating have to be changed by mechanical means. However, a very wide tuning range of about 100 nm can be obtained for semiconductor lasers by this method. This method of tuning is appropriate for test instruments but not for a compact light source for communication systems.

Tunable VCSELs

We studied VCSELs in Section 3.5.1. There we saw that the main challenges in realizing long-wavelength 1.55 μm VCSELs were in obtaining sufficient cavity gain, obtaining highly reflective mirror surfaces, dealing with the heat dissipation, and making the laser operate in a single-longitudinal mode. Figure 3.51 shows a VCSEL design [Vak99] that attempts to solve these problems, while also making the laser itself tunable. The tunability is achieved by having the upper mirror be a movable micro-electro-mechanical (MEM) membrane. The cavity spacing can be adjusted by moving the upper mirror by applying a voltage across the upper and lower mirrors. The upper mirror is curved to prevent beam walk-off in the cavity, leading to better stability of the lasing mode.

To conduct the heat away from the bottom mirror, a hole is etched in the InP substrate. The design uses a 980 nm pump laser to pump the VCSEL cavity. Any pump wavelength lower than the desired lasing wavelength can be used to excite the semiconductor electrons to the conduction band. For example, the 980 nm semiconductor pumps used to pump erbium-doped fiber amplifiers can be used here as well. By designing the pump spot size to match the size of the fundamental lasing mode, the laser can be made single mode while suppressing the higher-order Fabry-Perot cavity modes. Using gain to perform this function is better than trying to design the cavity to provide higher loss at the higher-order modes. The high gain also allows the output coupling reflectivity to be reduced, while still maintaining sufficient inversion inside the cavity to prevent excessive recombination.

The laser described in [Vak99] was able to put out about 0 dBm of power in continuous-wave (CW) mode over a tuning range of 50 nm.

Two- and Three-Section DBR Lasers

We saw earlier that we can change the refractive index of a semiconductor laser by injecting current into it. This can result in an overall tuning range of about 10 nm. The DFB laser shown in Figure 3.44 can be tuned by varying the forward-bias current, which changes the refractive index, which in turn changes the effective pitch of the grating inside the laser cavity. However, changing the forward-bias current also changes the output power of the device, making this technique unsuitable for use in a DFB laser.

A conventional DBR laser also has a single gain region, which is controlled by injecting a forward-bias current I_g, as shown in Figure 3.44(b). Varying this current only changes the output power and does not affect the wavelength. This structure can be modified by adding another electrode to inject a separate current I_b into the Bragg region that is decoupled from the gain region, as shown in Figure 3.52(a). This allows the wavelength to be controlled independently of the output power.

As in a conventional DBR laser, the laser has multiple closely spaced cavity modes corresponding to the cavity length, of which the one that lases corresponds to the wavelength peak of the Bragg grating. As the wavelength peak of the grating is varied by varying I_b, the laser hops from one cavity mode to another. This effect is shown in Figure 3.52(a). As the current I_b is varied, the Bragg wavelength changes. At the same time, there is also a small change in the cavity mode spacing due to the change in refractive index in the grating portion of the overall cavity. The two changes do not track each other, however. As a result, as I_b is varied and the Bragg wavelength changes, the laser wavelength changes, with the laser remaining on the same cavity mode for some time. As the current is varied further, the laser hops to the next cavity mode. By careful control over the cavity length, we can make the wavelength spacing between the cavity modes equal to the WDM channel spacing.

In order to obtain continuous tuning over the entire wavelength range, an additional third *phase* section can be added to the DBR, as shown in Figure 3.52(b). Injecting a third current I_p into this section allows us to obtain control of the cavity mode spacing, independent of the other effects that are present in the laser. Recall from Section 3.3.5 that it is sufficient to vary the effective cavity length by half a

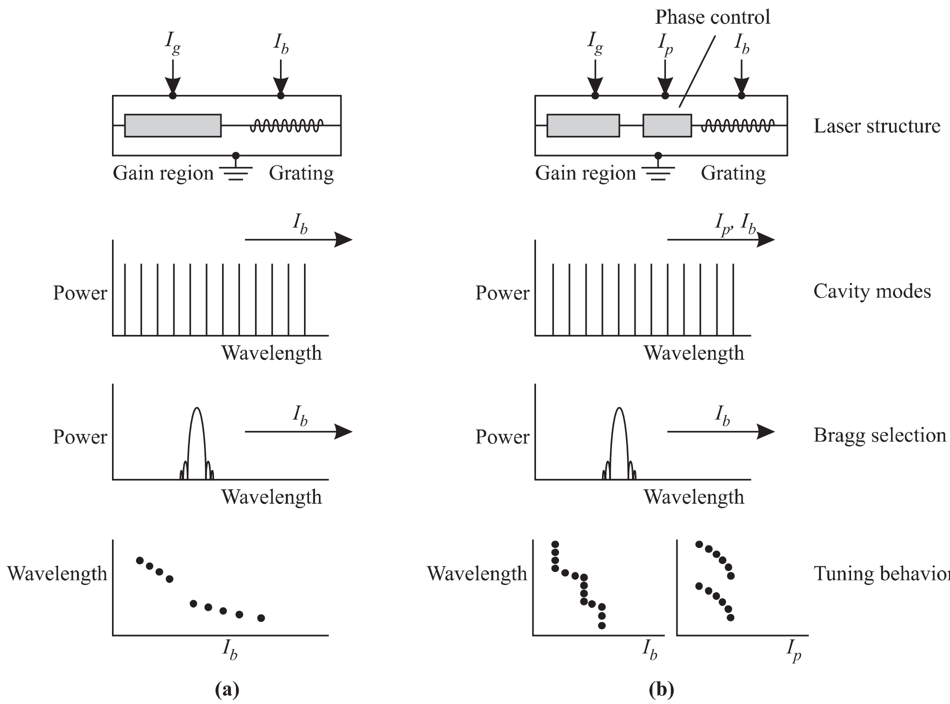

Figure 3.52 Two- and three-section DBR lasers and their principle of wavelength selection. (a) Two-section DBR showing separate control of the gain and Bragg sections. (c) Three-section DBR, which adds an additional control for the cavity phase.

wavelength (or equivalently, the phase by π) in order to obtain tuning across an entire free spectral range. This is a small fraction of the overall cavity length and is easily achieved by current injection into the phase section. By carefully controlling I_p to line up a cavity mode to correspond to the wavelength peak of the Bragg grating determined by I_b, the wavelength can be tuned continuously over the tunable range.

Two- and three-section DBRs capable of tuning over 32 channels in 50 GHz increments were demonstrated several years ago [KK90, Kam96] and are nearing commercial availability.

Clearly, a major problem that needs to be solved is in the control of these lasers, which can be quite complicated. As the laser ages, or temperature changes, the control currents may need to be recalibrated; otherwise the laser could end up hopping to another wavelength. The hopping could happen back and forth rapidly, and could manifest itself as relative intensity noise (RIN) at the laser output. In a sense, we are eliminating the very fact that made DFB lasers so wavelength stable—a fixed grating. These problems are only compounded further in the more complex laser structures that we will discuss next.

The DBRs that we have looked at so far are all limited to about a 10–15 nm tuning range by the 0.5–2% change in refractive index possible. Increasing the tuning range beyond this value requires a new bag of tricks. One trick makes the laser wavelength dependent on the *difference* between the refractive indices of two different regions. The overall variation possible is much higher than the variation of each of the individual regions. The so-called vertical grating-assisted coupler filter (VGF) lasers [AKB+92, AI93] make use of this principle. The second trick is to make use of the Vernier effect, where we have two combs of wavelengths, each with slightly different wavelength spacing. The combination of the two combs yields another periodic comb with a much higher wavelength spacing between its peaks. Problem 3.28 explains this effect in more detail. Even if each comb can be tuned only to a small extent, the combination of the two combs yields a much higher tuning range. The *sampled grating* (SG) DBRs and the *super-structure grating* (SSG) DBRs [JCC93, Toh93] use this approach. Finally, the *grating-coupled sampled reflector* (GCSR) laser [WMB92, Rig95] is a combination of both approaches.

VGF Lasers

Figure 3.53 shows the schematic of a VGF laser. It consists of two waveguides, with a coupling region between them. Its operation is similar to that of the acousto-optic tunable filter of Section 3.3.9. Using (3.17), wavelength λ is coupled from one waveguide of refractive index n_1 to the other of refractive index n_2 if

$$\lambda = \Lambda_B(n_1 - n_2)$$

where Λ_B is the period of the Bragg grating. Changing the refractive index of one region, say, n_1 by Δn_1, therefore results in a wavelength tuning of $\Delta\lambda$ where

$$\frac{\Delta\lambda}{\lambda} \approx \frac{\Delta n_1}{n_1 - n_2}.$$

This is significantly larger than the $\Delta n_1/n_1$ ratio that is achievable in the two- and three-section DBRs that we studied earlier.

In Figure 3.53, current I_c controls the index n_1, and current I_g provides the current to the gain region in the other waveguide. Just as with the two- and three-section DBRs, in order to obtain continuous tuning, the cavity mode spacing needs to be controlled by a third current I_p. Lasers with tuning ranges over 70 nm have been demonstrated using this approach.

One major problem with this approach is that the cavity length needs to be fairly long (typically 800–1000 μm) to get good coupling between the waveguides. This

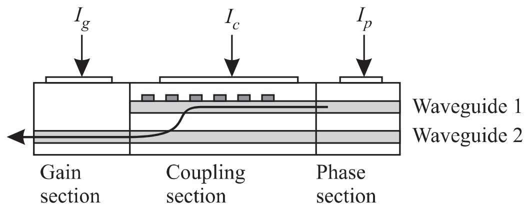

Figure 3.53 A vertical grating-assisted coupler filter tunable laser.

causes the cavity modes to be spaced very closely together. The laser therefore tends to hop fairly easily from one cavity mode to another, even though all the control currents are held steady. This effectively results in a poor side-mode suppression, making the laser not as suitable for high-bit-rate long-distance transmission.

Sampled Grating and Super-Structure Grating DBR Lasers

A sampled grating DBR laser is shown in Figure 3.54. It has two gratings, one in the front and one in the back. The Bragg grating in front is interrupted periodically (or *sampled*) with a period Λ_1. This results in a periodic set of Bragg reflector peaks, spaced apart in wavelength by $\lambda^2/2n_{\text{eff}}\Lambda_1$, as shown in Figure 3.54, where λ is the nominal center wavelength. The peaks gradually taper off in reflectivity, with the highest reflection occurring at the Bragg wavelength $2n_{\text{eff}}\Lambda$, where Λ is the period of the grating. The grating in the back is sampled with a different period Λ_2, which results in another set of reflection peaks spaced apart in wavelength by $\lambda^2/2n_{\text{eff}}\Lambda_2$. In order for lasing to occur, we need to have an overlap between the two reflection peaks of the Bragg gratings and a cavity mode. Even though the tuning range of each reflection peak is limited to 10–15 nm, combining the two sets of reflection peaks results in a large tuning range. Just as with the two- and three-section DBR lasers, a separate phase section controls the cavity mode spacing to ensure continuous tuning. An additional complication with this approach is that because the reflection peaks taper off, the current in the gain region needs to be increased to compensate for the poorer reflectivity as the laser is tuned away from the primary Bragg reflection peak.

Another way of getting the same effect is to use periodically chirped gratings instead of the gratings shown in Figure 3.54. This structure is called a super-structure grating DBR laser. The advantage of this structure is that the chirped gratings provide a highly reflective set of peaks over a wider wavelength range than the sampled grating structure.

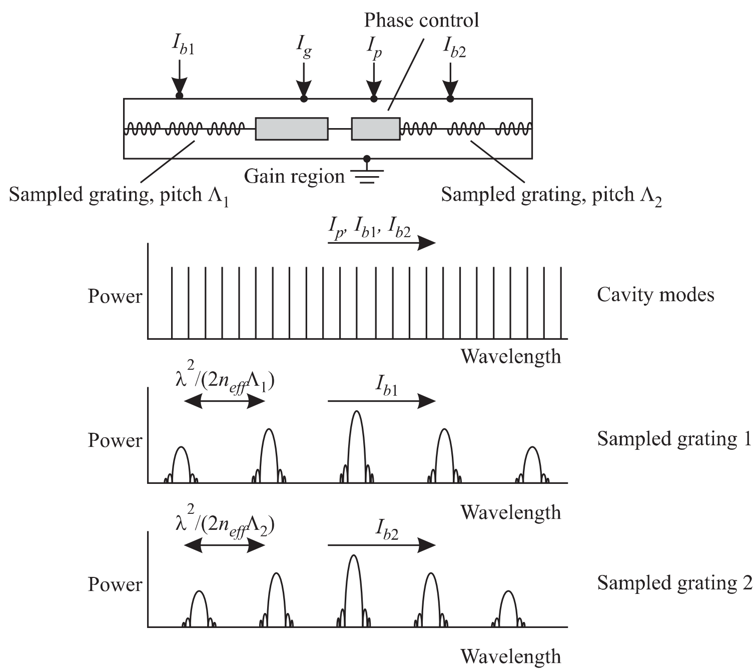

Figure 3.54 A sampled grating DBR laser and its principle of wavelength selection.

Grating-Coupled Sampled Reflector Laser

The GCSR laser is a combination of a VGF and a sampled or super-structure grating, as shown in Figure 3.55. The VGF provides a wide tuning range, and the SSG grating provides high selectivity to eliminate side modes. In a sense, the VGF provides coarse tuning to select a wavelength band with multiple cavity modes in the band, and the SSG grating provides the wavelength selection within the band. Just as in the two- and three-section DBR lasers, an additional phase section provides the fine control over the cavity modes to provide continuous tuning within the band to suppress side modes.

Laser Arrays

Another way to obtain a tunable laser source is to use an array of wavelength-differentiated lasers and turn one of them on at any time. Arrays could also be used to replace individual light sources.

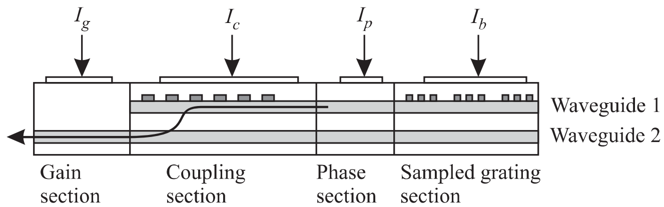

Figure 3.55 A grating coupled sampled reflector laser.

One approach is to fabricate an array of DFB lasers, each of them at a different wavelength. Combined with temperature tuning, we can use this method to obtain fairly continous tuning. A major problem with this approach is in the wavelength accuracy of the individual lasers in the array, making it difficult to obtain a comb of accurately spaced wavelengths out of the array. However, if only one laser is to be used at any given time, we can use temperature tuning to make up for this inaccuracy. Lasers using this approach have been demonstrated and used in system experiments [Zah92, You95].

Another approach is to use Fabry-Perot–type laser arrays and use an external mechanism for selecting the lasing wavelength. Several structures have been proposed [Soo92, ZJ94], one using an external waveguide grating and the other using an external arrayed waveguide grating. With these structures, the wavelength accuracy is determined by the external grating. The long cavity length results in potentially a large number of cavity modes within the grating wavelength selection window, which could cause the laser to hop between cavity modes during operation.

3.5.4 Direct and External Modulation

The process of imposing data on the light stream is called *modulation*. The simplest and most widely used modulation scheme is called *on-off keying* (OOK), where the light stream is turned on or off, depending on whether the data bit is a 1 or 0. We will study this in more detail in Chapter 4.

OOK modulated signals are usually realized in one of two ways: (1) by *direct modulation* of a semiconductor laser or an LED, or (2) by using an *external modulator*. The direct modulation scheme is illustrated in Figure 3.56. The drive current into the semiconductor laser is set well above threshold for a 1 bit and below (or slightly above) threshold for a 0 bit. The ratio of the output powers for the 1 and 0 bits is called the *extinction ratio*. Direct modulation is simple and inexpensive since no other components are required for modulation other than the light source

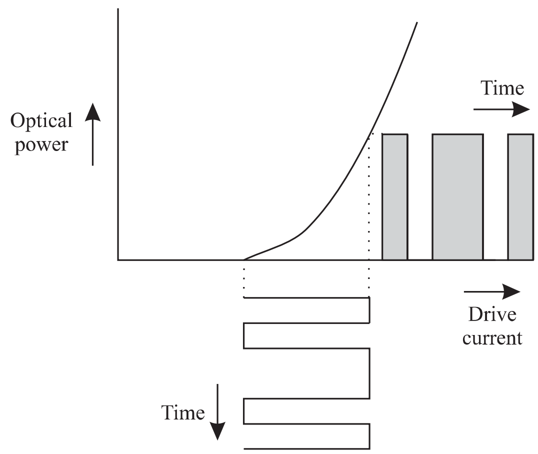

Figure 3.56 Direct modulation of a semiconductor laser.

(laser/LED) itself. In fact, a major advantage of semiconductor lasers is that they can be directly modulated. In contrast, many other lasers are continuous wave sources and cannot be modulated directly at all. These lasers require an external modulator. For example, because of the long lifetime of the erbium atoms at the E_2 level in Figure 3.35, erbium lasers cannot be directly modulated even at speeds of a few kilobits per second.

The disadvantage of direct modulation is that the resulting pulses are considerably *chirped*. Chirp is a phenomenon wherein the carrier frequency of the transmitted pulse varies with time, and it causes a broadening of the transmitted spectrum. As we saw in Section 2.4, chirped pulses have much poorer dispersion limits than unchirped pulses. The amount of chirping can be reduced by increasing the power of a 0 bit so that the laser is always kept well above its threshold; the disadvantage is that this reduces the extinction ratio, which in turn, degrades the system performance, as we will see in Section 5.3. In practice, we can realize an extinction ratio of around 7 dB while maintaining reasonable chirp performance. This enhanced pulse broadening of chirped pulses is significant enough to warrant the use of *external modulators* in high-speed, dispersion-limited communication systems.

An OOK external modulator is placed in front of a light source and turns the light signal on or off based on the data to be transmitted. The light source itself is continuously operated. This has the advantage of minimizing undesirable effects, particularly chirp. Several types of external modulators are commercially available and are increasingly being integrated with the laser itself inside a single package to reduce the packaging cost. In fact, transmitter packages that include a laser, external modulator, and wavelength stabilization circuits are becoming commercially available for use in WDM systems.

External modulators become essential in transmitters for communication systems using solitons or return-to-zero (RZ) modulation (see Section 2.6). As shown in Figure 3.57(a), to obtain a modulated train of RZ pulses, we can use a laser generating a train of periodic pulses, such as a mode-locked laser (see Section 3.5.1) followed by an external modulator. The modulator blocks the pulses corresponding to a 0 bit. (Usually we cannot directly modulate a pulsed laser emitting periodic pulses.) Unfortunately, cost-effective and compact solid-state lasers for generating periodic pulses are not yet commercially available. More commonly, as shown in Figure 3.57(b), practical RZ systems today use a continuous-wave DFB laser followed by a two-stage external modulator. The first stage creates a periodic train of short (RZ) pulses, and the second stage imposes the modulation by blocking out the 0 bits. Dispersion-managed soliton systems (see Section 2.6.1) require the generation of RZ pulses with a carefully controlled amount and sign of chirp. This can be accomplished by using another phase modulation stage.

Two types of external modulators are widely used today: lithium niobate modulators and semiconductor electro-absorption (EA) modulators. The lithium niobate modulator makes use of the electro-optic effect, where an applied voltage induces a change in refractive index of the material. The device itself is configured either as a directional coupler or as a Mach-Zehnder interferometer (MZI). Figure 3.58 shows the directional coupler configuration. Applying a voltage to the coupling region changes its refractive index, which in turn determines how much power is coupled from the input waveguide 1 to the output waveguide 1 in the figure.

Figure 3.59 shows the MZI configuration, which operates on the principles that we studied in Section 3.3.7. Compared to a directional coupler, the MZI offers a higher modulation speed for a given drive voltage and provides a higher extinction ratio. For these reasons, it is the more popular configuration. In one state, the signals in the two arms of the MZI are in phase and interfere constructively and appear at the output. In the other state, applying a voltage causes a π phase shift between the two arms of the MZI, leading to destructive interference and no output signal. These modulators have very good extinction ratios ranging from 15 to 20 dB, and we can control the chirp very precisely. Due to the high polarization dependence of the device, a polarization maintaining fiber is used between the laser and the modulator.

The EA modulator is an attractive alternative to lithium niobate modulators because it can be fabricated using the same material and techniques used to fabricate semiconductor lasers. This allows an EA modulator to be integrated along with a DFB laser in the same package and results in a very compact, lower-cost solution, compared to using an external lithium niobate modulator. In simple terms, the EA

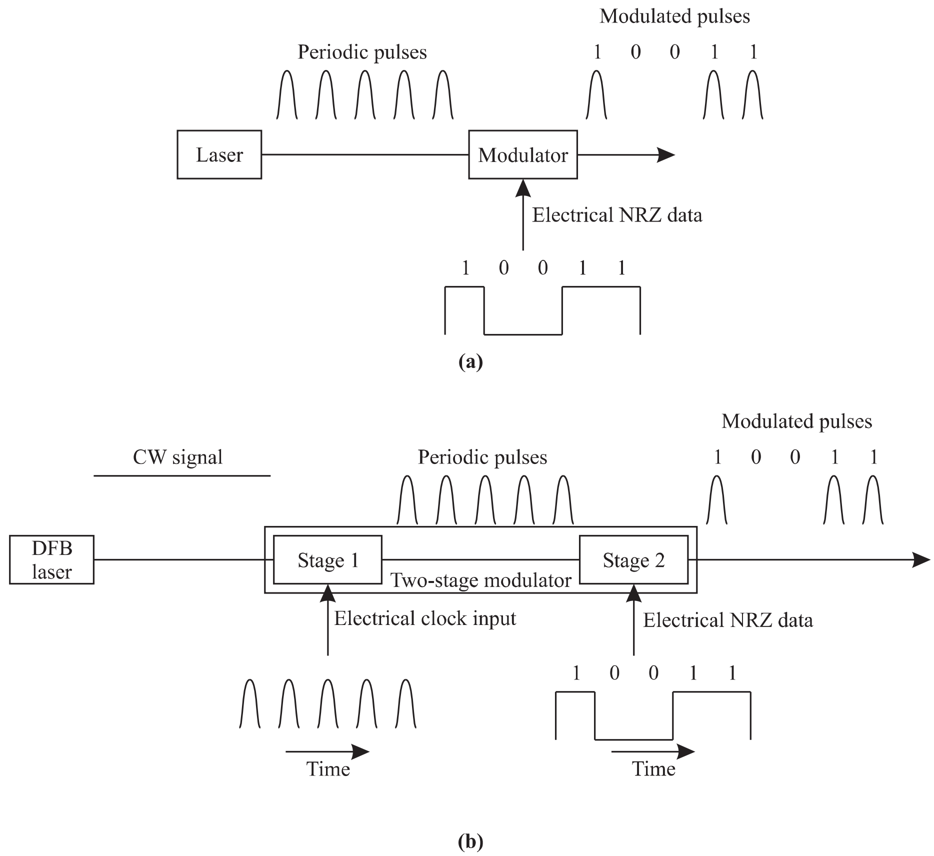

Figure 3.57 Using external modulators to realize transmitters for systems using RZ or soliton pulses. (a) A laser emitting a periodic pulse train, with the external modulator used to block the 0 bits and pass through the 1 bits. (b) A more common approach using a continuous-wave (CW) DFB laser followed by a two-stage modulator.

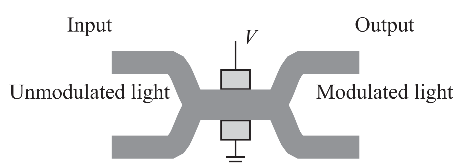

Figure 3.58 A lithium niobate external modulator using a directional coupler configuration.

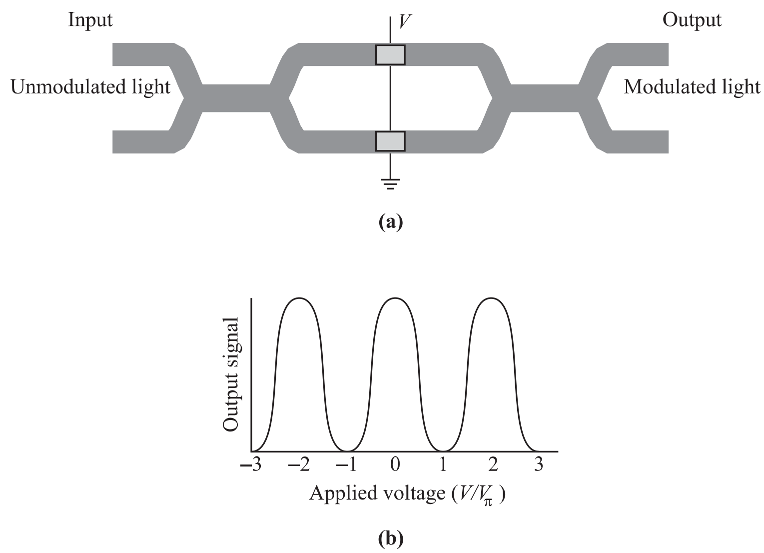

Figure 3.59 A lithium niobate external modulator using a Mach-Zehnder interferometer (MZI) configuration. (a) Device configuration. (b) Theoretical switching response as a function of applied voltage, V. V_π denotes the voltage required to achieve a π phase shift between the two arms. Note that the MZI has a periodic response.

modulator uses a material such that under normal conditions, its bandgap is higher than the photon energy of the incident light signal. This allows the light signal to propagate through. Applying an electric field to the modulator results in shrinking the bandgap of the material, causing the incident photons to be absorbed by the material. This effect is called the *Franz-Keldysh effect* or the *Stark effect*. The response time of this effect is sufficiently fast to enable us to realize 2.5 Gb/s and 10 Gb/s modulators. The chirp performance of EA modulators, though much better than directly modulated lasers, is not as good as that of lithium niobate MZI modulators. (While ideally there is no chirp in an external modulator, in practice, some chirp is induced in EA modulators because of residual phase modulation effects. This chirp can be controlled precisely in lithium niobate modulators.)

3.5.5 Pump Sources for Raman Amplifiers

One of the biggest challenges in realizing the Raman amplifiers that we discussed in Section 3.4.4 is a practical high-power pump source at the right wavelength. Since

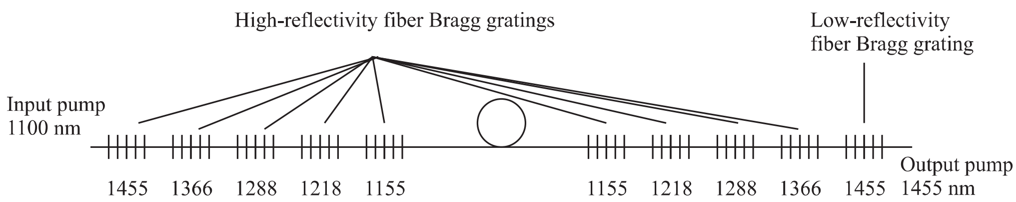

Figure 3.60 A high-power pump laser obtained by cascading resonators (after [Gru95]).

the Raman effect is only seen with very high powers in the fiber, pump powers on the order of several watts are required to provide effective amplification.

Several approaches have been proposed to realize high-power pump sources. One method is to combine a number of high-power semiconductor pump lasers. The power that can be extracted from a single semiconductor pump laser diode is limited to a few hundred milliwatts. Multiple semiconductor pump lasers can be combined using a combination of wavelength and/or polarization multiplexing to obtain a composite pump with sufficiently high power.

The other challenge lies in realizing the laser at the desired pump wavelength. One interesting approach is the cascaded Raman laser, shown in Figure 3.60.

Starting with a high-power pump laser at a conveniently available wavelength, we can generate pump sources at higher wavelengths using the Raman effect itself in fiber, by successively cascading a series of resonator structures. The individual resonators can be realized conveniently using fiber Bragg gratings or other filter structures. In Figure 3.60, a pump input at 1100 nm provides Raman gain into a fiber. A Fabry-Perot resonator is created in the fiber between by using a pair of matched fiber Bragg gratings that serve as wavelength-selective mirrors (see Section 3.3.5 for how the resonator works). The innermost resonator converts the initial pump signal into another pump signal at 1155 nm. It passes through signals at other wavelengths. The next resonator converts the 1155 nm pump into a 1218 nm pump. In principle, we can obtain any desired pump wavelength by cascading the appropriate series of resonators. The figure shows a series of resonators cascaded to obtain a 1455 nm pump output. The fiber Bragg grating at the end is designed to have lower reflectivity, allowing the 1455 nm pump signal to be output. This pump signal can then be used to provide Raman gain around 1550 nm. Due to the low fiber loss and high reflectivity of the fiber Bragg gratings, 80% of the input light is converted to the output.