BASE COMPONENTS

Injection Lasers

Semiconductor injection lasers serve as the primary source of light in an OEIC. Although vertical-cavity surface-emitting lasers (VCSEL) have been developed that emit their light perpendicular to the semiconductor layers (1), for OEICs it is more useful to have the lasers emit in the plane of the wafer (2). Such waveguide lasers have been widely developed as stand-alone devices, but they can also be incorporated into OEICs.

These lasers contain a gain medium, which is made from a semiconductor p–n junction. When forward biased, it emits light. Gain regions can be made from a bulk direct bandgap semiconductor (like GaAs or InGaAsP) with a bandgap *slightly longer* than the desired wavelength of operation, or from multiple quantum wells (MQWs). A quantum well is a thin layer of a low-bandgap material (such as GaAs) sandwiched between two layers of a higher-bandgap material (such as AlGaAs). Lasers made from MQWs typically require lower current for operation and may operate over wider temperature range than lasers with bulk gain media.

A gain medium is made into a laser by including optical feedback elements. The simplest structure (called a Fabry–Perot laser) uses two cleaved ends of a semiconductor chip for feedback mirrors. While simple to fabricate, the two cleaved mirrors do not lend themselves well to integration. Also, Fabry–Perot lasers typically emit an optical spectrum several nanometers wide, which is a drawback for many applications. More sophisticated distributed-feedback (DFB) structures use a grating fabricated along with the gain medium, which eliminates the need for at least one of the cleaved mirrors and typically produces an optical output with a very narrow frequency spectrum.

To achieve low optical losses, it is important for the light to be confined with an optical waveguide. The optical waveguide can typically be formed in the vertical direction (i.e., the direction perpendicular to the plane of the devices) by growing layers with varying refractive indices. However, forming lateral waveguides (i.e., in the plane of the devices) is more difficult and typically involves etching away some of the active (high refractive index) material and regrowing with a lower refractive index material. Figure 1 shows a view of an etched-mesa buried-heterostructure regrown laser structure.

Modulators

A continuously running laser is a good source of light, but there is no information carried on this light. Data can be impressed on the light beam simply by varying the current supplied to the gain section of the laser. This direct modulation of the laser output is a simple and widely used technique, but it has several disadvantages. The chief one is that directly modulating the laser introduces spurious broadening in the laser's optical spectrum (called "chirp"), which, when carried over a long optical fiber, reduces the maximum possible information rate.

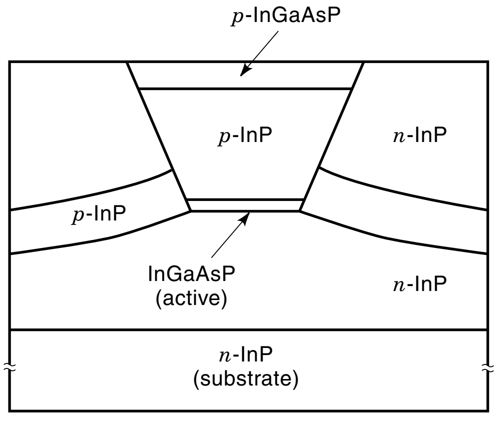

Figure 1. Schematic cross-section of an etched-mesa buried heterostructure laser structure (2). The active region (where the optical gain is generated) is the thin InGaAsP region. Current flows from the upper *p*-doped regions, through the active region, and into the *n*-doped InP regions. The *n*- and *p*-doped InP regions on the sides of the active region serve to block current from flowing through any path other than through the active region and also provide for optical waveguiding, since the InP material has lower optical index of refraction than the InGaAsP material.

This chirp can be minimized by running the laser continuously and imposing the information signal on the device with a separate optical element, an optical modulator. The most commonly used semiconductor modulators are electroabsorption (EA) modulators (3), although switches (discussed later) can be used. An EA modulator is made from a semiconductor region with a bandgap *slightly shorter* than the desired wavelength of operation. Unlike lasers, which are diodes operated in forward bias, EA modulators are *p–n* junctions operated in reverse bias, so that a variable electric field can be applied to the semiconductor layers. By varying this electric field, the optical absorption coefficient of the modulator material is varied, and thus the output optical intensity is determined by the applied voltage.

Like lasers, EA modulators can be made either from bulk direct bandgap semiconductors or quantum wells. The quantum well devices have demonstrated higher speed and lower drive voltages than their bulk counterparts.

Passive Waveguides

In addition to the active elements described previously, passive waveguides are often needed to interconnect the actives and serve as the optical "wiring" of the OEIC. The properties of these passive waveguides are important, since they may consume a sizable fraction of the OEIC and can introduce significant loss.

Passive waveguides are usually made from semiconductor layers with bandgap *much shorter* than the desired wavelength of operation. By increasing the detuning between the bandgap of the semiconductor and the wavelength of operation, the loss is reduced. To reduce loss further, it is important to keep highly doped layers, such as those found in contacts, far from the waveguide layers.

Detectors

The photodetector is an essential component in any optoelectronic system. The role of the photodetector is to convert the incident optical signal into an electrical signal efficiently and with minimal distortion. There are several types of detectors (4), including the Schottky barrier (SB) diode, metal–semiconductor–metal (MSM) photodiode, *p–i–n* photodiode, avalanche photodiode (APD), and phototransistor (PT).

The SB diode and the MSM detector (4) are probably the simplest from the point of implementation. The device uses a Schottky barrier between the metal and the semiconductor to form a depletion region where photoabsorption takes place. The MSM is a planar device with interdigitated electrode and therefore has very low capacitance. Both these detectors have high-speed performance capability, with the MSM detector trading some quantum efficiency for bandwidth. The devices do exhibit certain nonlinear characteristics with respect to bias voltages and optical intensity, and these are related to charge trapping effects at the Schottky barrier as well as zero-electric-field regions under the electrodes.

The *p–i–n* photodetector is the most commonly used device because of the high quantum efficiency, low-voltage operation, high speed, and ease of fabrication. Both planar (diffused junction) and mesa-structure *p–i–n* photodiodes have been used.

The avalanche photodetector provides gain through the avalanche multiplication mechanism but requires very critical control of the multiplication region and also operates at high bias voltages. It is an excellent candidate for very sensitive photodetection as well as in single photon counting and is a solid-state replacement for the photomultiplier tube.

The phototransistor (5) is basically a bipolar transistor with the base-collector *p–n* junction acting as a *p–i–n* photodetector and the photogenerated carriers amplified by the transistor gain mechanism. The device is attractive for its high-speed characteristics as well as the gain without the high voltage as needed for the APD, and it has found applications in optoisolators and remote control uses.

A relatively newer field of OEICs has seen the monolithic integration of photodetectors with electronic preamplifiers on a single chip. Such OEIC photoreceivers become attractive at very high speeds when the electrical connection between the photodetector and an electronic amplifier could cause noise, distortion, and signal degradation.

Transistors

The electronic functions following the photodetection are generally performed by a low-noise preamplifier. Here again, several types of transistor devices are available and can be broadly categorized as field-effect and bipolar devices.

In the field-effect transistors (FET), the transport of majority carriers from the source to drain is controlled by a field applied to a gate terminal. There have been different versions of the FETs demonstrated, the most prolific being the silicon metal-oxide-semiconductor (MOS) transistor. The others have names related to the type of semiconductor used to make them, such as the junction FETs (JFET), metal–semiconductor FET (MESFET), and the modulation-doped FETs (MODFET). The FETs are fundamentally planar devices with the speed of the device dependent on the gate length, which is defined lithographically. For a 1 μm gate length, these devices can have unity current-gain cutoff frequencies f_T in the range 15 GHz to 35 GHz. The small gate-to-source capacitance results in a large $f_{\max}$ (the maximum frequency of oscillation), typically in the range of 20 GHz to 50 GHz. Submicron electron beam lithography is used to achieve higher speeds in the 100 GHz range.

The bipolar device is the silicon bipolar junction transistor (BJT) and the heterojunction bipolar transistor (HBT) (6) in compound semiconductors. As the name indicates, transport in these devices is achieved via the combined interaction of the majority and minority carriers. The current flowing in the collector terminal is controlled by charge injected into the base terminal and is magnified by the current gain of the device. The HBTs incorporate an emitter with a bandgap wider than that of the base, and so can have higher base doping and lower emitter doping than their homojunction counterparts (BJT). This results in low base resistance and low base-emitter junction capacitance, together with high injection efficiency. In addition, when compared with the compound semiconductor FETs, the HBTs have higher transconductance and drive capability. More important, HBTs with high-speed performance can be fabricated with modest optical lithographic design rules. For emitter feature size between 2 μm and 5 μm, f_T values in the range of 50 GHz to 200 GHz and $f_{\max}$ values in the range of 35 GHz to 150 GHz have been reported. These frequencies are possible because the transport of the carriers is in the direction perpendicular to the device surface (in the direction of the epitaxial layers), and these distances can easily be controlled to hundreds of angstroms by epitaxy.

INTEGRATION TECHNOLOGY

Lateral Bandgap Engineering

As indicated previously, an OEIC typically consists of a series of devices that must be made with semiconductors of different bandgaps. Lasers, modulators, detectors, and passive waveguides all have different requirements on their bandgaps. Although some of this variation can be accomplished by growing different semiconductor layers on top of each other, in most cases, some variation of the bandgap must be achieved across the area of the OEIC. This is an important technological problem.

After growing material over the entire wafer, the simplest way to change the bandgap in one region is to etch away the material in that portion of the wafer and regrow with a different semiconductor composition. This etch and regrow technique is simple in principle and can produce high-quality devices. Figure 2 shows a schematic of a transition between a laser and a modulator formed in this way (7).

The etch and regrow technique has a number of problems in practice. First, the etch depth must be precisely controlled, or the regrown layer will not line up well with the original layer and light will not couple efficiently between the layers. A variety of defects can be formed in the processing, due to impurities in the etch or regrowth on the nonplanar transition region. Thus, this technique can be difficult to use.

An alternative approach that grows varying bandgaps across the wafer in a single step is shown in Fig. 3 (8). A silicon dioxide (SiO_2) mask is deposited on the wafer and photolithographically patterned into stripes, as shown. When a subsequent semiconductor layer is grown by metal-organic chemical vapor deposition (MOCVD), the reactants that impinge on the SiO_2 region diffuse onto the semiconductor surface before reacting. This means that the semiconductor layers that form in the gap between the stripes grow faster than layers grown far from the stripes. If quantum wells are grown

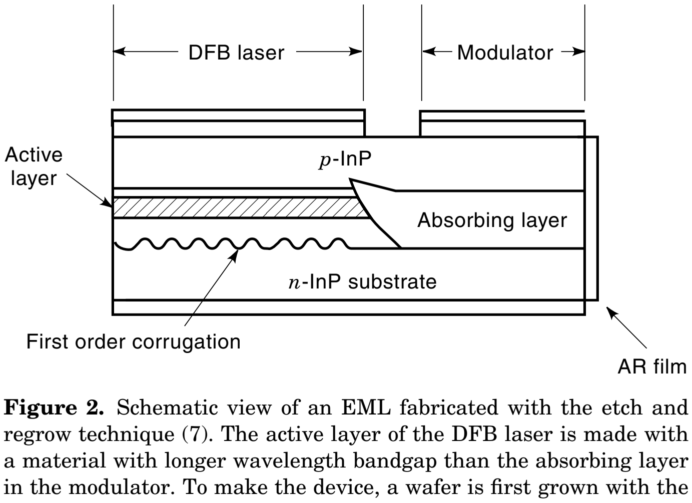

Figure 2. Schematic view of an EML fabricated with the etch and regrow technique (7). The active layer of the DFB laser is made with a material with longer wavelength bandgap than the absorbing layer in the modulator. To make the device, a wafer is first grown with the active layer everywhere. This is etched away in the portion of the modulator of the wafer where modulators are to be placed, and the absorbing layer is regrown in the holes.

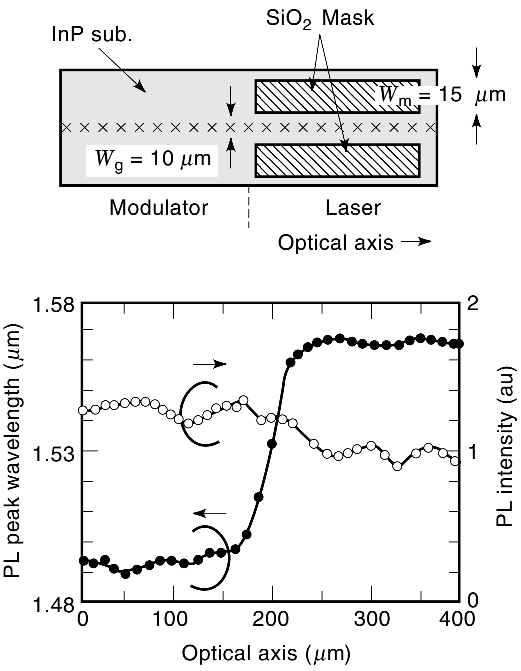

Figure 3. Schematic view of how bandgap variation is achieved using the selective area epitaxy technique (8). The upper portion of the figure shows the silicon dioxide (SiO$_2$) stripes patterned on the InP substrate. In this case, the stripes are 15 μm wide and the gap between them is 10 μm. When layers of InGaAsP are grown on this patterned substrate by MOCVD, the growth is faster in the modulator region between the SiO$_2$ stripes than it is in the laser region far from the stripes. The lower portion shows a graph of the photoluminescence (PL) peak wavelength as a function of position along the wafer. The PL peak wavelength is a measure of the semiconductor bandgap wavelength. The PL wavelength shifts from ~1.49 μm in the modulator region to ~1.57 μm in the laser section.

in this process, the wells in the gap will be thicker than those grown elsewhere, and thicker quantum wells have a longer bandgap wavelength than thin quantum wells. Thus, using this "selective area epitaxy" technique, the bandgap can be smoothly varied across the wafer without etching and regrowth.

EXAMPLE DEVICES

Transmitters—Electroabsorption-Modulated Lasers

The first and most successful commercially available PIC is the electroabsorption-modulated laser (EML), which is a single-frequency DFB laser integrated with an electroabsorption modulator. This device provides a low-chirp source of modulated light that can be placed in the same package as a directly modulated laser, which greatly simplifies transmitter design.

Figure 4 shows a schematic of one EML, which is based on using the selective area epitaxy described previously (8,9). A set of quantum wells are grown by MOCVD on an InP wafer with an SiO$_2$ mask. The quantum wells grow thicker in the laser region and thinner in the transition and modulator regions. This change in well thickness allows the modulator region to be of shorter wavelength bandgap (and thus more transparent) than the laser region at the operating wavelength.

There are a number of issues that influence the design of the integrated laser/modulator and that are not of concern in designing discrete lasers or modulators. First, the process must be designed to support both laser and modulator. An example of this is the lateral bandgap engineering discussed previously. Second, the yield on both devices must be sufficient so that acceptable yield is achieved in the larger integrated device.

However, the most significant integration issues center on making sure the integrated laser and modulator do not interact with each other. Any crosstalk between the two devices will cause the laser frequency to chirp as the modulator's voltage is changed, which will reduce the performance of the integrated device. To reduce this effect, both electrical and optical crosstalk must be reduced. Electrical crosstalk causes the current in the laser to change as the modulator's voltage is varied, which generates laser chirp. The current leakage between the two top contact pads can be reduced by engineering the doping of the top layers, by etching a portion of them away, or by using ion implantation to increase the electrical resistivity of this region.

Optical crosstalk is caused by light that is reflected off the modulator's exit facet and returns through the modulator to the laser. As the voltage on the modulator varies, the intensity of this light varies, which can change the laser's wavelength, inducing chirp. To reduce this effect, a high-quality antireflection (AR) coating needs to be applied to the exit facet of the modulator. In addition, this effect can be reduced by including a window structure in the modulator's waveguide that expands the optical beam slightly before it reaches the facet, or by designing the modulator so that it has some absorption even in its maximum transmission state.

By careful design, EMLs can have performance equal to or exceeding discrete lasers and modulators. They are being widely deployed on long-haul telephony and data transmission systems.

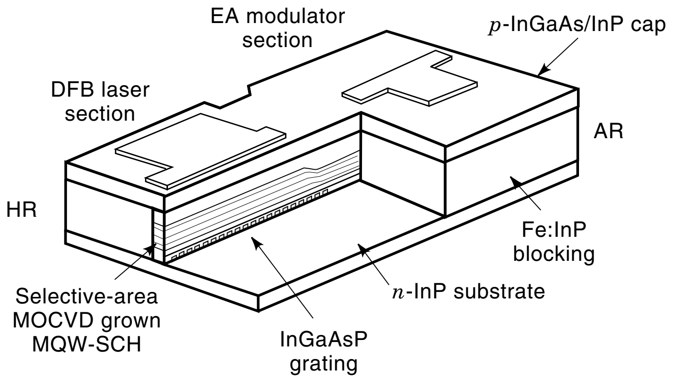

Figure 4. Schematic view of EML fabricated by the selective area epitaxy technique (9). In the DFB laser section, an InGaAsP grating is formed below the active region to make the laser emit in a single wavelength. The relatively thick quantum wells are seen in the laser section and transition smoothly to thinner quantum wells in the EA modulator section. A high-reflectivity (HR) coating is applied to the rear facet of the EML, while an antireflection (AR) coating is applied to the front.

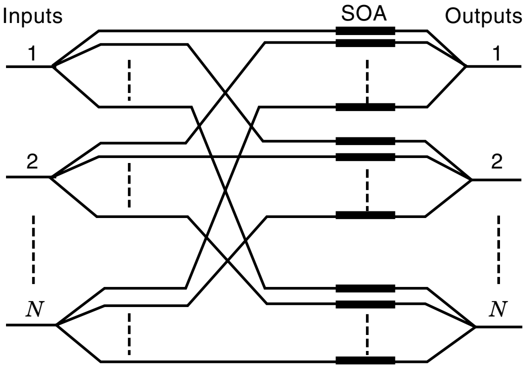

Figure 5. Schematic view of a broadcast and select optical cross connect based on the use of semiconductor optical amplifiers (SOA) (10). With this device, optical signals from any of N inputs can be routed to any of N outputs. Waveguide branches at the inputs broadcast the optical signals to the combiners at the outputs. SOAs along the way can be turned off (to eliminate unwanted signals) or turned on (to provide amplification of desired signals).

Cross Connects, Routers

Although useful, the modulators in EMLs can only act like a shutter, turning light on and off. For many applications, it would be useful to be able to route light from N input waveguides to any of N output waveguides. This optical cross connect has application in optical switching systems, where light from a number of sources must be switched to a number of different destinations.

One class of cross connects is based on arrays of Mach–Zehnder interferometric space switches (10). These switches equally split light from two inputs into two separate waveguides. The phase of the light in these two waveguides can be controlled by the application of an external voltage. After propagating through the waveguides, the light is mixed in a second two-output splitter. By varying the external voltage, the light can be switched from either input port to either output port. Up to 4 × 4 optical switches have been fabricated made from arrays of these 2 × 2 switch elements.

A diagram of a second class of cross connects, based on a "broadcast and select" architecture, is shown in Fig. 5 (10). Light from each of the N inputs is split N ways and recombined in each of N outputs. A semiconductor optical amplifier (SOA) is included in each intermediate waveguide. This SOA has two functions. When no current is applied, the SOA strongly absorbs the light that is not desired at a given output port. When current is applied, the SOA not only passes desired light, but it provides gain for this light, making up for the splitting and other losses in the cross connect. Up to 4 × 4 optical switches using this technology have been reported.

There are tradeoffs in the choice of technology for these cross connects. The Mach–Zehnder switches often have significant insertion loss, and the extinction ratio of the unwanted light is often limited to ~15 dB, which is not adequate for some applications. The SOA-based switches have a better extinction ratio, and the amplifiers can often provide enough gain to more than compensate for the loss through the switch. However, the amplifier introduces noise into the optical signal, and a degradation of on/off ratio for modulated optical signals as well as crosstalk can occur due to changes in the SOA's carrier concentration. Heat dissipation of the SOA may also limit the size of such a switch.

No significant commercial deployment of these cross connects has occurred to date, but as the need for rearrangeable optical networks increases and the switch technology matures, it seems likely that these devices will be put into service.

Wavelength Converters

Wavelength division multiplexing (WDM) techniques offer an increased utilization of the fiber bandwidth by transporting information on several wavelengths. The number of wavelengths in WDM networks determines the number of independent wavelength addresses or paths. Sometimes it will not be large enough to support a large number of nodes, causing blocking due to wavelength contention (for example, when two channels of the same wavelengths are to be routed at the same output). One method of overcoming this limitation is to

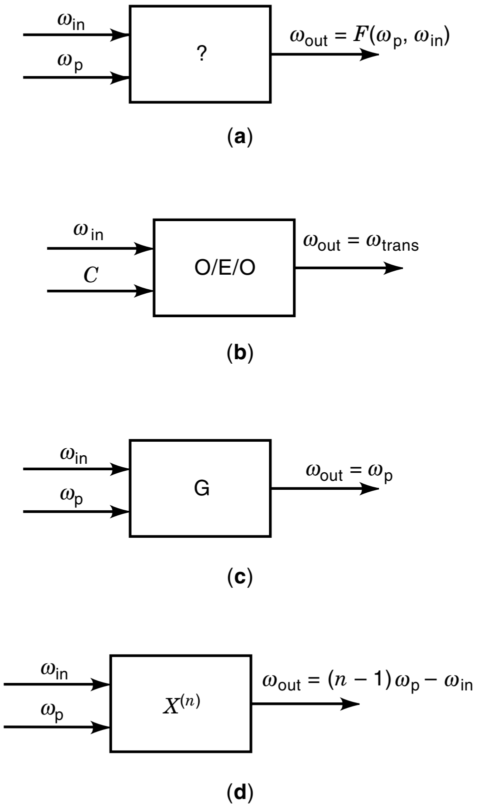

Figure 6. (a) Functional block diagram of a general wavelength converter (11). The rectangular box could be one of three possible functions depicted in (b), (c), and (d). An incoming optical signal at a frequency ω_{in} interacts with another optical signal ω_p in the box to generate the desired output frequency ω_{out}, which is related to the input signals via a transfer function F. (b) An optoelectronic wavelength converter. In this case the box contains a photoreceiver and an output laser. The input signal corresponding to ω_p is not needed and is therefore depicted as a control signal C, which essentially sets the output laser wavelength to the desired value. (c) An optical gating-based wavelength converter. In this case the box could have an optical gain medium such as a semiconductor amplifier and the control signal is an optical probe at a frequency ω_p. The two interact via either the cross-gain modulation or the cross phase modulation to generate the output frequency ω_{out}. (d) An optical wave-mixing-based wavelength converter. In this case the box contains a nonlinear optical medium where optical mixing of the input signals generates the output wavelength.

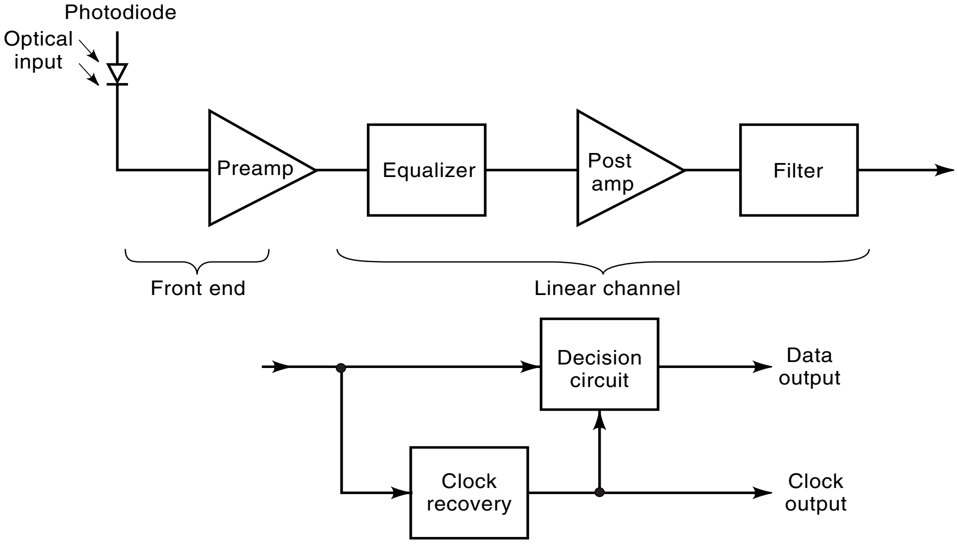

Figure 7. Block diagram of a basic lightwave receiver (13). The photodetector detects light and is amplified by the preamplifier. The two constitute the front end. The signal is then amplified further in a linear channel and then split two ways. On one path, the clock is recovered in a clock recovery module, which is then used to gate a decision circuit whose input is the other path. A conditioned signal is recovered, along with the clock.

convert signals from one wavelength to another (11). In another situation, there could be a network user whose wavelength is not compliant with the rest of the network. In such cases, the noncompliant wavelength is converted to a compliant one and then launched into the network. Wavelength conversion is beginning to appear both necessary and important in WDM networks.

Figure 6(a) shows a functional block diagram of a general wavelength converter, which can be broadly classified into three categories: optoelectronic, optical gating, and wavelength mixing, as shown in Figs. 6(b) to 6(d).

Optoelectronic Wavelength Conversion. The most straightforward of all wavelength conversion techniques is a detection of the optical signal in a photoreceiver and retransmission of the signal using a transmitter, such that the output wavelength is a compliant wavelength. This method is a variable-input-fixed-output wavelength converter with very little freedom to change the data rate.

Optical Gating Wavelength Conversion. This type of wavelength converter changes its characteristics depending on the intensity of the input signal. This change is monitored by a continuous wave (cw) signal called probe, and this probe signal will contain the information in the input signal. There are several wavelength conversion methods that fall into this category, which includes semiconductor optical amplifier cross-gain modulation, semiconductor optical amplifier cross-phase modulation, semiconductor lasers with saturable absorption, and nonlinear optical loop mirrors. This is a classical example of variable-input-fixed-output wavelength converters using all-optical techniques. The method could, in principle, be data-rate independent.

Wave-Mixing Wavelength Converters. The least explored but offering the highest degree of transparency is wavelength conversion based on nonlinear optical wave mixing. The nonlinear interactions among the waves present in the nonlinear optical material give rise to the mixing products. This mechanism is sensitive to both amplitude and phase information. To date, four-wave mixing based on third-order optical nonlinearity and difference frequency generation based on second-

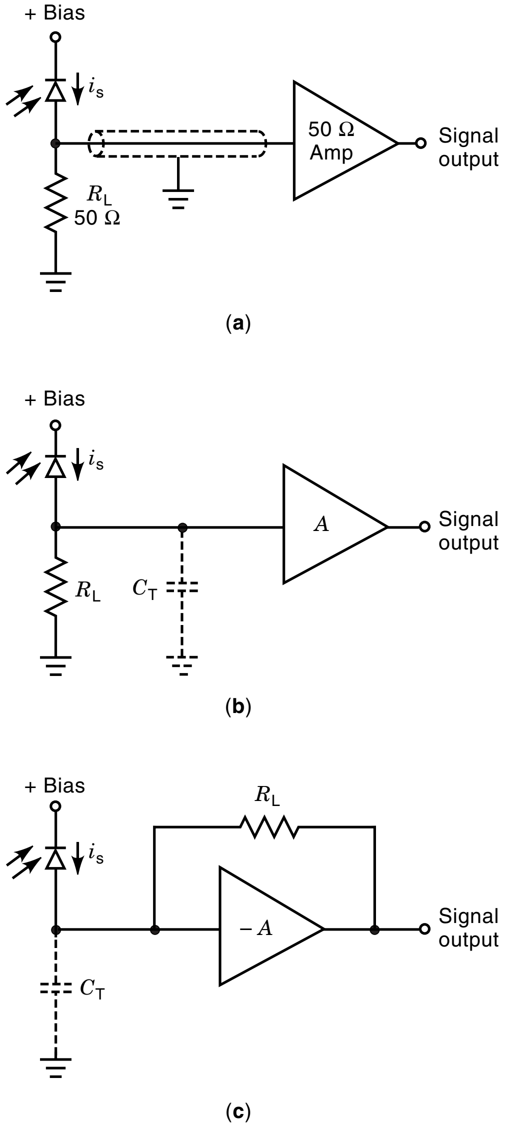

Figure 8. Three different types of front end receivers (13). (a) Low-impedance front end. (b) High-impedance front end. (c) Transimpedance front end.

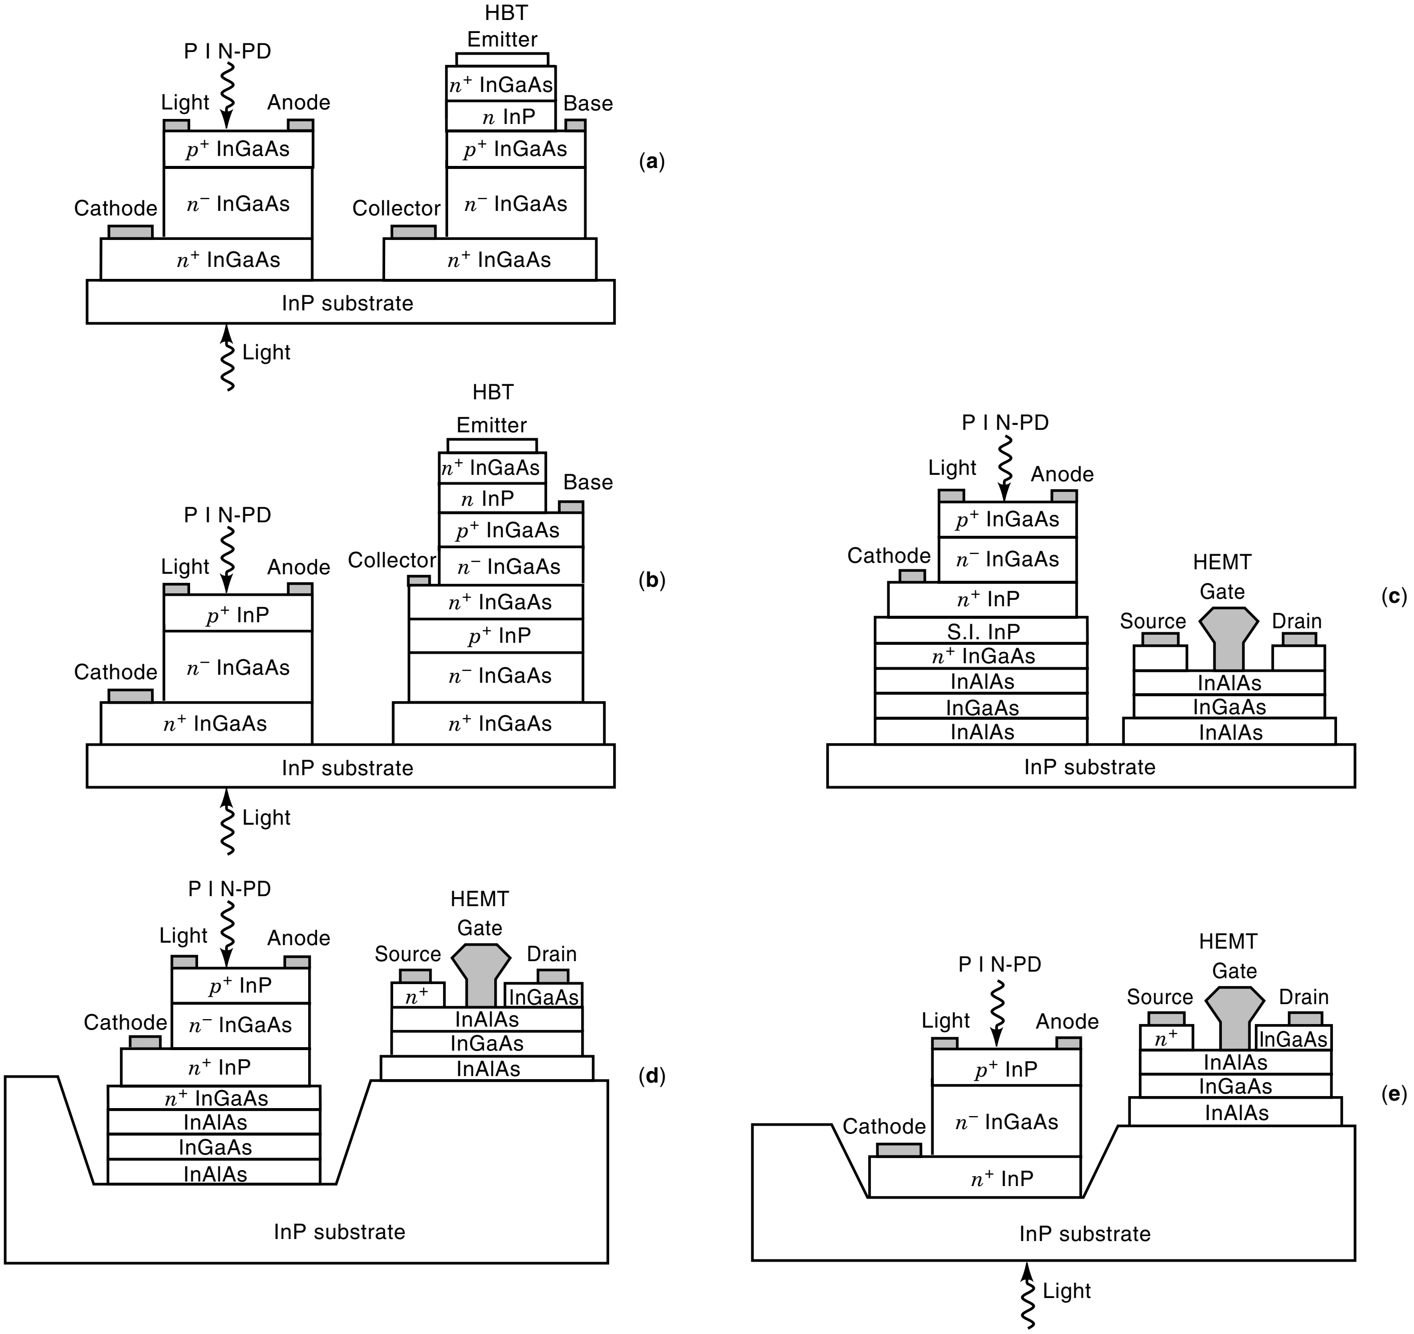

Figure 9. Schematic representation of different possible combinations of photodetector and transistor to realize a monolithic photoreceiver (OEIC) (14). In (a), the *p–i–n* and HBT share a group of epitaxial layers. In (b), the *p–i–n* and HBT have independent layers and the *p–i–n* layers are below the HBT layers. In (c), the *p–i–n* layers are above the HEMT layers on a planar substrate. In (d), the structure is the same as in (c), except that the substrate has been trenched to result in a planar surface. Finally, in (e), the *p–i–n* layers are grown in a trench in the substrate, while the HEMT layers are grown on the plane surface, so that in the end the devices are almost planar.

order optical nonlinearity have been demonstrated. The utilized mixing functions are parametric, and the mapping functions allow one-to-one mapping of an input wavelength to an output wavelength. Thus, the conversion process is variable-input-variable-output. One unique feature of this category of wavelength converters is that they allow simultaneous conversions of multiple input wavelengths to multiple output wavelengths.

Several types of wavelength converter devices have been demonstrated, such as SOAs, semiconductor lasers, and harmonic generators. In addition, there have been several PICs which combine smaller building blocks to achieve a functional wavelength converter. As an example, Ref. 12 describes an interferometric wavelength converter which includes a continuous-wave source, four SOAs, and four Y-branching waveguides, all on one InP substrate.

Receivers

In its simplest form, an optical receiver incorporates a photodetector to convert the optical signal into an electrical signal, followed by a low-noise electronic amplifier to raise the electrical signal level to a value that can be used for further signal processing (13). This constitutes the front end of the optical receiver. The back end typically consists of an automatic gain control (AGC) amplifier, followed by a decision circuit that is clocked by a clock signal recovered from the incoming signal via a clock recovery circuitry. The block diagram of a typical optical receiver is shown in Fig. 7. The front end of the receiver constitutes an important part as the quality of the signal detected and amplified at this stage determines the performance of the overall receiver.

The front end is constructed either as a hybrid integration of any one of the photodetectors described previously with an electronic preamplifier integrated circuit chip, or a monolithic integration of the two elements in a compatible semiconductor technology. Different transistor technologies have been used to implement the preamplifier circuit, with each technology having its own strengths and weaknesses. The preamplifier itself could be of three types—namely, a low-impedance, a high-impedance, and a transimpedance design, as shown in Fig. 8.

Low-impedance front ends typically consist of a photodiode connected to a low-impedance (e.g., 50 Ω) amplifier. This configuration generally does not provide high sensitivity because only a small signal voltage can be developed across the amplifier input impedance. However, because of the ready availability of 50 Ω RF (radiofrequency) and microwave amplifiers, such front ends are often used in situations where sensitivity is not a consideration.

High-impedance front ends, as the name indicates, consist of a high-input impedance amplifier that amplifies the photocurrent from the photodiode to a large signal. Such preamplifiers integrate the input signal and therefore have higher sensitivities. However, because of this nature, the preamplifier has to be followed up with an equalizer in order to compensate for the reduced bandwidth. One disadvantage of the high-impedance design is the limited dynamic range.

The transimpedance amplifier is a feedback amplifier with the load resistor connected as a feedback resistor around an inverting amplifier. This design has wider bandwidth for a given photodiode capacitance and load resistor than the high-impedance design, lower noise than the low-impedance design, and larger dynamic range among the different designs. It is the preferred design in most practical implementations.

Monolithically integrated photoreceiver front ends (also called optoelectronic integrated circuits) (14) have been demonstrated with different combinations of photodetectors and transistors. Heterojunction bipolar transistors and field effect transistors have been integrated with p–i–n photodetectors, M–S–M detectors, as well as heterojunction phototransistors to realize photoreceivers covering a wide spectrum of applications and operating speeds from 100 Mbit/s to 20 Gbit/s. A few combinations of p–i–n photodetector and transistor integration methodologies are shown in Fig. 9. The performance of monolithic photoreceivers has been shown to be comparable to and in some cases even better than hybrid approaches. One expects that as the speed of lightwave communication systems increases beyond 10 Gbit/s, monolithic photoreceivers will begin to play an important role.

Moreover, monolithic integration allowed demonstration of multichannel array receivers for potential application in WDM systems. Arrays from 2 to 16 channels have been demonstrated with a total throughput on a single chip of up to 160 Gbit/s. An issue common to monolithic arrays is crosstalk due to simultaneous operation of several channels. The crosstalk is usually electrical in origin and occurs predominantly from channels adjacent to the signal channel, and its magnitude increases with increasing signal level in the interfering channels. Nevertheless, acceptable levels of crosstalk have been demonstrated under carefully designed integration and packaging technology.

CONCLUSION

We have shown that a variety of complex functions can be performed using integrated optoelectronic circuits. Although these OEICs are made on a much smaller scale than the corresponding electronic-only integrated circuits, they perform key functions in several systems. As the fabrication technology for OEICs evolves, we expect more sophisticated and less expensive devices to become available.

BIBLIOGRAPHY

1. T. E. Sale, *Vertical Cavity Surface Emitting Lasers,* New York: Wiley, 1995; Kent D. Choquette, Vertical-cavity surface emitting lasers: moving from research to manufacturing, *Proc. IEEE,* **85**: 1730–1739, 1997.
2. G. P. Agrawal and N. K. Dutta, *Long-Wavelength Semiconductor Lasers,* New York: Van Nostrand Reinhold, 1986; N. Holonyak, Jr., The semiconductor laser: A thirty-five-year perspective, *Proc. IEEE,* **85**: 1678–1693, 1997.
3. T. H. Wood, MQW waveguide modulators, *J. Lightwave Technol.,* **6**: 743–757, 1988.
4. S. M. Sze (ed.), *Physics of Semiconductor Devices,* New York: Wiley, 1981.
5. J. C. Campbell, in *Phototransistors for Lightwave Communications,* vol. 22, part D, *Semiconductor and Semimetals, Lightwave Communications Technology,* Orlando: Academic Press, 1985.
6. H. Kroemer, Heterostructure bipolar transistors and integrated circuits, *Proc. IEEE,* **70**: 13–25, 1982.
7. H. Soda, K. Nakai, and H. Ishikawa, Frequency response of an optical intensity modulator monolithically integrated with a DFB laser, *14th Eur. Conf. Opt. Commun. (ECOC),* Brighton, UK, 1988, pp. 227–230.
8. M. Aoki et al., InGaAs/InGaAsP MQW electroabsorption modulator integrated with a DFB laser fabricated by band-gap energy control selective area MOCVD, *IEEE J. Quant. Electron.,* **29**: 2088–2096, 1993.
9. J. E. Johnson et al., High-speed integrated electroabsorption modulators, *High-Speed Semicond. Lasers Commun., Proc. SPIE, Vol.* **3038**, 1997, pp. 30–38.
10. M. Renaud, M. Bachman, and M. Erman, Semiconductor optical space switches, *IEEE J. Select. Topics Quantum Electron.,* **2**: 277–288, 1996.

11. S. J. B. Yoo, Wavelength conversion technologies for WDM network applications, *J. Lightwave Technol.,* **14** (6): 955–966, 1996.
12. L. H. Spiekman et al., All-optical Mach-Zehnder wavelength converter with monolithically integrated DFB probe source, *IEEE Photon. Technol. Lett.,* **9**: 1349–1351, 1997.
13. B. L. Kasper, Receiver design, in S. E. Miller and I. P. Kaminow (eds.), *Optical Fiber Telecommunications II,* San Diego, CA: Academic Press, 1988, pp. 689–722.
14. Special Issue on Integrated Optoelectronics, T. P. Lee Ed., *Int. J. High Speed Electron. Syst.,* **5** (1): 91–109, 1994.

THOMAS H. WOOD
S. CHANDRASEKHAR
Bell Labs—Lucent Technologies

INTEGRATED OPTOELECTRONICS. See ELECTRO-OPTICAL FILTERS.