

OPTOELECTRONICS IN VLSI TECHNOLOGY

Optoelectronics in VLSI (OE-VLSI) refers to the intimate integration of optoelectronics with Very Large Scale Integration (VLSI) electronic chips (i.e., the placement of optoelectronic devices including lasers, detectors, and modulators directly on the VLSI chip) and may be performed via flip-chip bonding or with other techniques outlined later. Described here are two potential benefits of this integration. A third benefit, the production of more sophisticated optical imaging technology, is not discussed here because that requires only detectors and is covered under the concepts of focal plane arrays or VLSI cameras, whereas OE-VLSI usually refers to the integration of both detectors and transmitters.

The first benefit is the reduction in electrical parasitics associated with wire-bonding and metallic board traces in optical transceivers. The term optoelectronic integrated circuit (OEIC) has been applied to small-scale electronic integration (SSI) with optoelectronics, typically by forming electronics in the same material as the optoelectronic devices (i.e., monolithic integration). Optoelectronics are typically formed in III–V materials such as GaAs or InP with direct band gaps, and not in Si which has an indirect band gap. There are numerous examples of photoreceivers fabricated as OEICs; recent ones described in Refs. 1 and 2 have photoreceivers that contain less than 100 III–V transistors. Such monolithic techniques are usually limited to SSI because the yield of transistors fabricated in the (typically) heteroepitaxial III–V material used to form optoelectronics is low. There are VLSI GaAs technologies (3), but these are based on nonheteroepitaxial materials and not suitable for optoelectronics. Si can form detectors, and an example of a photoreceiver fabricated in Si–Ge is given in Refs. 4, but was limited to a frequency of 450 MHz. Beyond SSI, optical transceivers may perform a variety of complicated functions requiring VLSI electronics and are usually modules containing several chips including the optoelectronic chips. For example, in local area optical computing networks there are network interface cards (NICs) that perform a variety of complicated functions such as multiplexing and demultiplexing of the faster optical stream into

multiple lower-speed electrical streams, as well as functions such as data encoding and address look-up. As data rates increase, parasitic reduction may prove to be important. Additionally, if the optoelectronic integration can be performed on a batch of chips in one step as outlined later, production cost may go down. These concepts fall in the category of “system-on-a-chip.” This first area in which OE-VLSI may be important is just beginning to be explored, as data rates in optical networks exceed 10 Gbit/s, and any advantage of OE-VLSI concepts has not been shown. Indeed, it may turn out that a modular approach has a lower cost as a result of independently testing components and experiencing difficulties with mixing analog and digital signal types. These questions will be answered in the marketplace in the next few years.

A second benefit of OE-VLSI, which has actually been the mantra driving the field since its inception, derives from the needs of VLSI itself. This benefit has at various times been called optical computing, digital optics, and as a subset of computing, photonic switching and has its roots in the belief prevalent among some in the mid 1980s that more efficient computing could be performed by manipulating light rather than electrons. This belief is based on the physical fact that because light beams do not interact in free space, massive parallelism is possible in communication (5). However, it is precisely this noninteraction of photons that makes control of light by light a difficult proposition. In fact, electrons must mediate the interaction of photons, whether through a nonlinear optical process that is within a solid that does not require the actual absorption of photons but rather simply the polarization of electron clouds by the optical field or on the other end of the scale, via the detection of photons with discrete detectors, electronic processing of the signals, and retransmission of photons using discrete transmitters. The former seems so elegant and the latter so bulky that, from a physicist's view, the former must be more efficient. However, the relentless progress of Si-based electronics negates this argument. When billions of transistors can be placed upon a single chip, and it is difficult to manipulate passively thousands of light beams, the winner in this battle is clear.

However, in the last few years, the field of optical computing has been transformed so that rather than compete with electronic computing it derives strength from it. In a historically perverse fashion, now rather than hoping that electronic computing power has a limit, the field of *optics in computing* is encouraging it on. The reason for this is that optics has a well-recognized strength compared to electronics, which is its ability to communicate information. This has been made clear in long distance, and increasingly in short distance, optical fiber communication. Going forward this type of usage may become a reality in the very short distances between electronic chips or even on the chips themselves between transistors. The latter is outside the scope of this discussion, but chip-to-chip optical communication may soon become reality because the computing power of a chip requires ever-increasing communication on and off the chip, which may not be supported by electrical wiring alone.

This prediction is supported by the plot in Fig. 1. There the number of chip input/output (I/O) is plotted versus silicon VLSI Complementary Metal Oxide Semiconductor (CMOS) line rules two ways. The lower curve is the Semiconductor Industry Association's (SIA) prediction (6). This curve, which itself reaches numbers in the thousands as line rules decrease in the next decade, is based partly on what the association

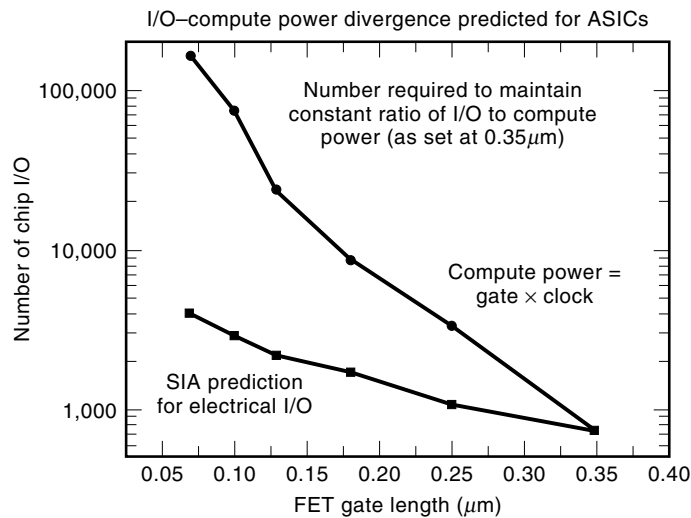


Figure 1. Plot of Semiconductor Industry Association prediction of VLSI chip I/O going into the next century (6) and that necessary for parity with computing power of chip, which predicts that for sub-tenth micrometer VLSI, >100,000 I/O required, indicating that direct optical array I/O to chips will be necessary.

thought was technically possible. However, the upper curve is given by keeping the ratio of I/O bandwidth to computational bandwidth (transistor count times clock rate) constant, assuming I/O clock rate equals chip clock rate. Keeping this ratio constant may result in the highest efficiency of the chip. The upper curve reaches numbers in the hundreds of thousands in the next decade. Below it is shown that the optoelectronic yield may support such numbers.

Given these two applications of OE-VLSI, this article focuses on the chip-to-chip communication application rather than the NIC application. There are four reasons for this:

1. Even though concepts such as system-on-a-chip are powerful and the NIC market is substantial, it pales in comparison to the possibility that all state-of-the-art computing systems will have optical communication, both in terms of how much more revolutionary this concept is scientifically and for its sheer pervasiveness.
2. It is perhaps more appropriate that the subject of OEICs expand to cover the use of more functionality in transceivers.
3. If the power of OE-VLSI in NICs is greater functionality and if the full circuitry for only one optical channel can be placed on a single chip at first, in fact it is necessary that research be spent demonstrating the benefit of the *electronic* integration (signal processing, memory, and transceiver circuitry) before as a last step a single pair of optoelectronic devices is added.
4. Following from the last statement, the true power of OE-VLSI is the ability to form large arrays of optoelectronics with electronics, sometimes referred to as smart pixels because they process an array of light beams.

Given this application choice, it is clear that the type of OE-VLSI technology used must support tens or hundreds of thousands of optoelectronic devices per chip. We therefore focus on the only optoelectronic transmitting device demon-

strated to support such yield, the surface-normal GaAs-based *p-i-n* diode optical modulator. When choosing an optical transmitting element, even though one's impulse is to choose a device that emits light for simplicity in the optical system, LEDs do not provide sufficient directed power, and lasers have not yet demonstrated sufficient yield for next generation computing chip I/O. Arrays of lasers may be formed using vertical-cavity surface-emitting lasers (VCSELs), but detailed measurements have not been made with arrays larger than 64, and no computing demonstration has been made at all. In a modulator-based system, even though it is true that the optical system has more elements because the outputs of the system are measured by applying an array of read-out beams to the modulators (6), this actually leads to some conceptual simplicity in the system because the beam array is generated by diffracting a single laser. This means that all the beams are automatically at the same optical wavelength, simplifying the design of the lenses and gratings of the system. Additionally, the single read-out laser may be clocked, providing global clock distribution to the system.

Also, because the modulator is a diode, it also functions as a detector, so that only one level of optoelectronic integration is required. The modulator works by having a material whose absorption changes with electric field (an electroabsorptive material) in the intrinsic region, so that upon reverse bias the diode's transmission (or reflection if there is an integral mirror) changes. By biasing the device in its absorbing state, it functions as a photodiode. Electroabsorption is produced by using a readout wavelength near the band gap of the intrinsic material. Via quantum mechanical effects, the band edge states are altered with electric field, leading to changes in optical absorption (7). By employing a multiple quantum well (MQW) material, which has layers of thickness the order of the wavelength of an electronic quantum mechanic wavefunction, electroabsorption is enhanced (8). Thus the modulator used in smart pixel arrays is usually called an MQW-modulator.

In arguments favoring eventual use of VCSELs in smart pixel systems, the temperature sensitivity of the modulator caused by the change of the band gap with temperature is cited. Although the emission wavelength of a VCSEL will change with temperature, if the optical system is made achromatic (i.e., the focusing and beam directing optics of the system operate the same over a wide range of wavelengths), a VCSEL-based system may operate over a wider temperature range than a modulator-based system. We show later that modulators could operate over a 60°C range. Whether this is sufficient for computing systems probably depends on the application. Here is a possible scenario: just as display technology sometimes favors emitters [cathode ray tubes (CRTs)] and sometimes modulators [liquid crystal displays (LCDs)] so OE-VLSI technology will favor one or the other or even a mixing of both on the same chip.

Mentioning LCDs reminds us that there are many types of optical modulators, but a discussion of them all would be outside the scope of this article in which we focus on the speed of semiconductor devices. There has, however, been extensive research on the integration of liquid crystal with VLSI. In fact, an LCD is itself an integration of liquid crystal with electronic circuitry, and there has been research into extending this technology into computing or as a smart display (9). Even though the speed of liquid crystal is limited to rates lower than a megahertz and so could never transmit the informa-

tion of a single electronic pin, the yield of liquid crystal modulators will always surpass any other modulator because the active element is in the electronic circuitry. Therefore, it is theoretically possible to have millions of liquid crystal modulators on a chip. However, generating such a large number of readout beams would be problematic.

In the next section, a brief overview of OE-VLSI technology based on MQW modulators is presented. This section reviews the progress made in integrating modulator devices with electronic circuits and discusses yield and manufacturing issues for future commodity processes. The section entitled "Circuit Considerations for OE-VLSI" discusses circuit considerations for the MQW modulator based smart pixel technology and reviews standard-cell designs for the optoelectronic transceivers. The section entitled "Large Switch Experiment" reviews the first major demonstration of the OE-VLSI technology, a switch with terabit capability. Next, the section entitled "Toward a Commodity OE-VLSI Technology" presents results associated with the first CMOS/SEED foundry shuttle, and reviews a future roadmap for this technology. Conclusions constitute the last section.

TECHNOLOGY OVERVIEW

The multiple quantum well modulator has driven much of the work in smart pixels because of three main principles: (i) it may function either as detector or transmitter; (ii) it has high enough yield so that large working arrays (i.e., up to thousands) have been fabricated; and (iii) it may operate at high speeds (i.e., Gbit/s). The device is a $p-i-n$ diode, usually arranged in a layered structure produced by epitaxial growth techniques such as molecular beam epitaxy. By applying a reverse bias to the diode, the electric field in the intrinsic region increases, and the light absorption in that region changes by the principle of electroabsorption. Therefore, the light transmission through the device is modulated by the applied voltage signal. Electroabsorption will occur in any semiconductor near its electronic band gap because the gap is effectively lowered by the electric field. In addition, electron-hole pairs (excitons) are produced by absorption of photons with energies near the band gap, and this absorption line will shift with the electric field. A good description of electroabsorption in bulk (uniform) semiconductors can be found in Ref. 7. The shift of an exciton line provides for strong modulation, but at room temperature in bulk semiconductors the exciton is broadened in energy, reducing its effectiveness. A MQW material, however, maintains the narrowness of the exciton at room temperature. The structure of an MQW is a repeated sequence of layers of smaller band gap material and larger band gap material, whose individual layers are so thin (i.e., 100 Å) that the excitonic quantum mechanical wave function is confined in the layers of the smaller band gap material. Hence these layers are called quantum wells. A good review of MQW electroabsorption is contained in Ref. 8. Thus the preferred modulator, usually called the MQW modulator and sometimes referred to as a $p-i(\text{MQW})-n$ diode, is simply a $p-i-n$ diode where the intrinsic layer is itself a multilayered MQW.

As a transmissive device, the contrast between the high and low transmission states is typically less than 2:1 for voltage swings below 10 V. The device may also work in reflection




Figure 2. Reflectivity spectra of MQW modulator, which has been used up to now as the optoelectronic transmitting element in all large-scale interchip optical communication system experiments.

mode if a mirror is incorporated. In reflection, two passes through the device occur, squaring the contrast. The mirror may be a multilayer quarter wave semiconductor mirror produced as part of the epitaxial structure (10), although metal reflectors can also be used if the light is brought from the substrate side. The reflectivity spectra of a representative MQW modulator is shown in Fig. 2. If both top and bottom mirrors are incorporated, Fabry-Perot cavities, which are optimally balanced, can be formed so that at a particular voltage there is essentially zero reflectivity. Thus, in principle, an infinite reflection contrast can be achieved (11,12). Such a technique is needed for systems conveying analog information in the optical beams, and OE-VLSI chips have been demonstrated with high contrast, large arrays by using post-integration etching of the modulator to tune its optical thickness (13). However, for digital information, large chip-to-chip bandwidth has been demonstrated with reflection contrasts of 2:1.

Review of MQW Modulator-Based Smart Pixels

Self-Electrooptic Effect Devices. The first smart pixels were not very "smart" but could perform a logic function. These are called self-electrooptic effect devices (SEEDs) but, more appropriately, would be considered circuits composed of a series connection of a constant voltage source, a MQW modulator, and a load. This load could be another MQW modulator, in which case it is called a symmetric SEED (S-SEED). In operation, this circuit is bistable, which allows it to perform logic functions. It is bistable because when the operating wavelength is on the exciton absorption line, the photocurrent will decrease as reverse bias is applied, causing negative photoconductivity. Therefore, as the light intensity is increased, the voltage on the modulator, and hence its transmission, will undergo a discontinuous change at some point. As the light intensity is then decreased, the voltage on the modulator will undergo a discontinuous change at a light level lower than produced in the upward part of the curve, causing a hysteresis in the transmission-versus-light-intensity curve. Thus, at a particular light intensity, or, in the case of the S-SEED,

when the light intensities on the two modulators are equal, the circuit can exist in one of two different states, depending on the history of operation. Logic can then be performed. In particular, logic level restoration can be achieved in the following manner. The S-SEED circuit could be set into a state where the light intensity level on diode 1 was higher than on diode 2. Then, in the next clock cycle, data could be read onto the circuit, and unless the input to diode 2 was of higher intensity than the input to diode 1, the circuit state would not change. Thus, a logic decision could be performed. Finally, in the next clock cycle, the state of the circuit could be read by beams of equal and *higher* intensity. Because the read beams, if equal, can be of higher intensity than the write-beams without changing the state of the circuit, signal gain and, hence logic level restoration, can be achieved. The logic decision is simple, but full logic can be performed by using more than one beam per diode. A review of SEED operation is contained in Ref. 14. SEED arrays consisting of 64×64 arrays of modulators have been fabricated (15).

In principle, the SEED circuit is the fastest smart pixel because its speed is simply limited by the charging time of the diodes, without any intervening electronics. However, that charging time is limited by the photocurrent and hence the optical power. If the capacitance of the SEED is 100 fF, for a voltage change of 10 V, it would take approximately 20 mW to switch the device in 100 ps, or 20 W for an array of 1000 elements.

To reduce this power and to make for broader functionality, electronics can be integrated with the modulators. It is a philosophical point whether, when electronics are integrated with the modulators, the resulting circuit is more appropriately considered a circuit with optical I/O, rather than a SEED. Because electronics have been added to modulators, the name of the resulting object has been appended with the word SEED, even as the intervening electronics have begun to dominate the schematic. Our viewpoint is that using the term SEED is appropriate when identical devices [$p-i(\text{MQW})-n$ diodes] are used for both the input and output, as follows from the term *self* in SEED and evidenced in the earliest SEED devices. As mentioned earlier, this quality of the $p-i(\text{MQW})-n$ diode is both the great utility of the device, because, systems can be built with just one level of optoelectronic/electronic integration, and the true inventive concept of the originators of the SEED. As multiple levels of integration are developed (either specialized devices for detection and modulation, or detection and light emission), the name SEED and the concept it implies may eventually be dropped, in favor of the term optoelectronic-VLSI. As other OE-VLSI technologies emerge, this term may further be amended to modulator-based OE-VLSI.

Field Effect Transistor SEEDs. The integration of electronics with MQW modulators can occur either by bringing the electronics to the modulators or vice versa. The first substantial attempt was the former, the field-effect transistor SEED (FET-SEED). The integration consisted essentially of forming field-effect transistors in the top-doped layer of the modulator (16). In this way, modulators and FETs are formed side-by-side in the same layer structure. In a standard MQW modulator, both doped regions have a larger band gap than the quantum well material so that these regions are transparent. In the FET-SEED, a quantum well is placed in the top (n) doped

region at a carefully controlled distance from the surface (17). Because it is so thin, it does not induce appreciable light absorption, and so does not affect the performance of the modulators. For the FETs, the quantum well serves the function of being able to fill and empty electrons, depending on the surface field, which can change via a gate metal. FETs are formed by producing ohmic contacts to either side of the gate. These transistors functioned well, and a switching network using 4×4 arrays of smart pixels, each composed of 21 FETs and 11 modulators, was demonstrated at a clock rate of 155 Mbit/s (6). However, migrating to larger array sizes and more complex circuitry put undue stress on the yield of the process. In addition, circuit design was problematic given the available modeling of the devices.

Hybrid CMOS-SEEDs. Finally, we come to our current process, which brings the modulators to the electronics, and in particular, to silicon CMOS electronics. There are several ways this may be accomplished. For instance, one may attempt to grow, by epitaxy, GaAs-based modulators on Si (18). Indeed, there are volumes written on the growth of GaAs on Si, and such growth is possible. However, integration via growth is a different matter. The growth must occur before metalizations are placed on the CMOS as a result of thermal processing constraints; the subsequent reproduction of submicron Si metalizations in a research environment would be problematic.

Herein lies another pragmatic issue concerning optical I/O to electronics, which is that if the electronics cannot be state-of-the-art, this exercise may not be profitable. One of the primary arguments for optical I/O is that electrical I/O alone will not be able to meet the needs of Si electronics as it advances. However, if the quality of the electronics must be compromised, this argument may fail. Therefore, any integration scheme must be able to use state-of-the-art electronics without degradation; hence, there is a strong preference for processes that are as noninvasive as possible to the electronics. Because heteroepitaxial growth involves temperatures of about 800°C, it is not clear that it can be integrated as a process without degrading the Si in some manner.

A second guiding principle for any integration technique derives also from the argument that electrical I/O will fail to keep pace with Si electronics. From this argument, it is apparent that smart pixels may not be useful until the number of optical channels per chip is in the thousands, because it is apparently not until this number is reached by pin-out count that there will be a problem. The ability of the optical interconnects to provide complex interconnection topologies also becomes an advantage at these large numbers. Therefore, any integration process must be capable of producing such numbers of optoelectronic devices, preferably in a single step.

The only candidate that suffices both of the preceding criteria is flip-chip bonding because it is a low-temperature process and attaches many devices at once. We have shown that thousands of solder pads can be produced on chips by photolithography and evaporation. Using our commercial bonding machine (19), aligning an optoelectronic chip laterally to a Si CMOS chip can be routinely performed to 2 μm accuracy. Thus attaching thousands of high-speed devices in a single step is routine. The temperatures involved are below 200°C and so do not affect the silicon.

We have innovated on the process of flip-chip bonding, particularly for optoelectronic devices. This is necessitated by the fact that the best MQW modulators operate at a wavelength of 850 nm (20). However, the GaAs substrate of such modulators is opaque at that wavelength. Therefore, we remove it after the bonding (21). This additional step has become fairly routine in our process, although some groups avoid it by working with reduced efficiency modulators at wavelength where GaAs is transparent (22). Our hybrid CMOS-SEED process arranges for both the n and p contacts of the diodes to be brought to the surface and to be coplanar and also ensures that each device is electrically isolated from its neighbors. Solder is deposited on one or both of the Si and GaAs chips, and the chips are bonded together. We find that gold-coated tin works well as a solder and that good bonding occurs by tacking the chips together at an elevated temperature and pressure without the need for reflow. Then epoxy is wicked between the chips by capillary action. The epoxy both provides mechanical stability and protects the front surfaces of the chips during the next step, which is the removal of the GaAs substrate. This is done by chemical etching. The etch is stopped by a stop-etch layer, which has been placed between the GaAs substrate and the device layers (Fig. 3). Each MQW modulator has two flip-chip pads, p and n , one of which is




Figure 3. Process flow for integrating MQW modulator arrays to VLSI chips by flip-chip bonding followed by substrate removal (23).




Figure 4. Forward biased emission pattern from a 64×16 array of flip-chip bonded diodes followed by substrate removal with 100% yield. Solder pads are $15 \mu\text{m} \times 15 \mu\text{m}$, and device size is $18 \mu\text{m} \times 52 \mu\text{m}$.

optically active. The active region of the device as drawn in Fig. 3 is the n -pad which also serves as a reflector. Each device is about $18 \times 52 \mu\text{m}$ in size. Each has a $15 \times 15 \mu\text{m}$ n -pad and a $15 \times 15 \mu\text{m}$ p -pad, separated laterally by $15 \mu\text{m}$. The epoxy may subsequently be removed if necessary (e.g., to expose wire-bond pads).

To test electrical connection, the devices may be forward-biased to emit. The emission pattern from a 16×64 array of devices is shown in Fig. 4. We demonstrated arrays of bonded modulators and simple circuits in (23). Our largest chips to date are $7 \text{ mm} \times 7 \text{ mm}$ Si chips with 140,000 to 450,000 CMOS gates with a 64×68 array of modulators arranged in a $5.5 \text{ mm} \times 5.5 \text{ mm}$ field (24). We have produced such chips with 100% modulator yield (corresponding to a device yield of 99.97%), but typically have about 1 to 5 nonworking diodes, with the yield reduction mostly caused by defects in the modulators. Thus, our device yield for these larger chips is about 99.9%.

A key advance in the hybrid CMOS-SEED process was demonstrated in Refs. 25 and 26. Here it was shown that the modulators may be bonded to metal interconnect layers on the Si chip directly above transistor layers, creating a three-dimensional optoelectronics/VLSI circuit. Typical submicron CMOS technologies provide at least three levels of metal, so the modulators may be bonded to the top level. This removes any constraints in the placement of the underlying circuits with respect to the modulators and allows a designer to take full advantage of sophisticated VLSI design tools. The cross-sectional structure of the resulting OE-VLSI circuit is shown in Fig. 5. The modulators form an optoelectronic overlay to the electronics, and in a sense, our mission of being able to provide optical I/O to Si CMOS without compromising in any way its ability to form circuits is complete. Using this technique we can make OE-VLSI chips with dense modulator arrays. The highest density circuit (in terms of transistors per unit area) built to date had an 8×8 array of MQW diodes




Figure 5. Structure of a three-dimensional hybrid GaAs MQW/silicon CMOS circuit.

and over 21,000 transistors in a silicon area of $1 \text{ mm} \times 1 \text{ mm}$ (26). In terms of MQW diodes per unit area, the densest OE-VLSI chips have 32×64 modulators arranged in a $2.3 \text{ mm} \times 2.3 \text{ mm}$ field with over 400,000 transistors in an silicon area of approximately $8.2 \text{ mm} \times 3.8 \text{ mm}$. We routinely have only one or zero bad diodes in these dense chips, for a yield $> 99.95\%$.

Note that there have been many other technologies developed to attach optoelectronics to chips, including variations of flip-chip bonding such as in Ref. 13, where the modulator chip is epoxied first to a glass substrate, the GaAs substrate is removed, and then flip-chip bonding is performed. This technique lends itself to easier etch tuning to achieve high contrast Fabry-Perot devices, and large arrays have been demonstrated. Other techniques include heteroepitaxy of modulators onto GaAs VLSI (27) and removal of the thin device layers from GaAs substrates and transfer onto Si VLSI chips (28). However, arrays larger than 64 have not been shown using other techniques, and until this happens, these techniques are not applicable to large-scale computing. Thus, they are not discussed here.

Yield Limits

As mentioned, VCSELs may offer some advantage because they do not require an off-chip readout light source. There may be a device advantage of modulators, however, in that they may be much less sensitive to material defects. The reason for this is that all that is required for a modulator to work is that its reverse-biased leakage current be low enough not to induce heating. Barring process or lithographic errors, nothing else can go wrong with the device as long as a gross error does not occur in the wafer growth (a large pit or particle that actually interferes mechanically with the formation of the modulator, discussed later). Typically modulators do not degrade with time because very little current flows through the device. Lasers, on the other hand, are susceptible to a number of degradation mechanisms, including most prominently dark-line defect formation, which probably leads to much higher sensitivity to microscopic defects.

To provide a framework for the discussion from the modulator side, data are given here designed to show microscopic

material defect levels *pertinent to modulators*. We do this by forming extremely large modulators with $5 \text{ mm} \times 5 \text{ mm}$ active areas. We then measure the reverse leakage current. If the leakage is below a set value, the region occupied by the modulator is “modulator defect” free, and we can quantify the number of defects per unit area.

Because we are exploring material properties, we give details of the growth conditions and substrates used in this experiment, which were from the same ingot. The growth was performed by molecular beam epitaxy on a Varian GEN II solid source system, with substrate rotation of 10 rpm, and a substrate temperature of 645°C . The V/III flux ratio was 20. The arsenic source was Furukawa lot No. 40601T, the gallium source was Rhone-Paulenc lot no. 88088, and the aluminum source was Morton Thiokol lot no. E110. The arsenic beam pressure was 1.5×10^{-5} torr, and the GaAs growth rate was $1 \mu\text{m/h}$.

The 2 in. substrates were from American Xtal Technology ingot no. E3H201, with a resistivity of 1 to $4 \times 10^{-7} \Omega\text{-cm}$ and a mobility of $7000 \text{ cm}^2/\text{Vs}$ to $9000 \text{ cm}^2/\text{Vs}$. Their orientation was $(100) \pm 0.1^\circ$ and thickness was $500 \pm 25 \mu\text{m}$. The etch pit density was $< 3000/\text{cm}^2$.

Figure 6 shows a map of the 12 diodes we fabricated on one wafer. Each diode has a $5 \text{ mm} \times 5 \text{ mm}$ active area. The vertical full scale is 1 mA. This would correspond to 4 nA for a $10 \mu\text{m} \times 10 \mu\text{m}$ device, the size we anticipate using for our flip-chip bonded technology in the future, and this leakage is an acceptable level for such a modulator. Therefore 10/12 of the diodes on this wafer contain no “modulator” defects, assuming that the nature of the leakage in those devices is distributed across many distributed small defects. Note that this makes no comment on the level of defects pertinent to other types of devices, such as VCSELs. The two diodes with large leakage seem to suffer from some other type of catastrophic defect. This defect is invisible to visual examination. For a second, identical wafer, similar behavior was observed, with 9/12 diodes show $< 1 \text{ mA}$ leakage at 10 V.

As mentioned, it is assumed in this analysis that these “soft” leakages are caused by many “small” defects. We are not attempting to state here the exact nature of these defects but simply to recount their effect. It is not known whether this leakage is caused by many defects that contribute equally or whether the defects contribute differing amounts of leakage. We can comment that all the diodes with this “soft” leakage, when left under bias for $\sim 1 \text{ h}$, actually had a reduction of leakage current while left on, so these defects do not appear to contribute to any reliability problems.

For at least one diode, though, that displayed large leakage, only soft leakage occurred until it had been biased up for a few seconds, so it appeared that something “snapped” in the device. Thus we believe that the diodes with large leakage suffer from another type of defect that is catastrophic to the modulator. It occurs in 5/24 diodes of the two wafers, and we assume here that we have actually observed five distinct point defects. Again, we do not know that the large leakage is caused by several defects in each diode, but it seems that because of the distinct effect and the fact that we see it only in a small fraction of the diodes, a single defect may be responsible for each bad diode. Then, the defect count is $5/(24 \times 5 \times 5 \text{ mm})$, or $0.83/\text{cm}^2$.

Note that the position of the fatal defects in the wafers is the same and that the wafers were cut from the same boule.

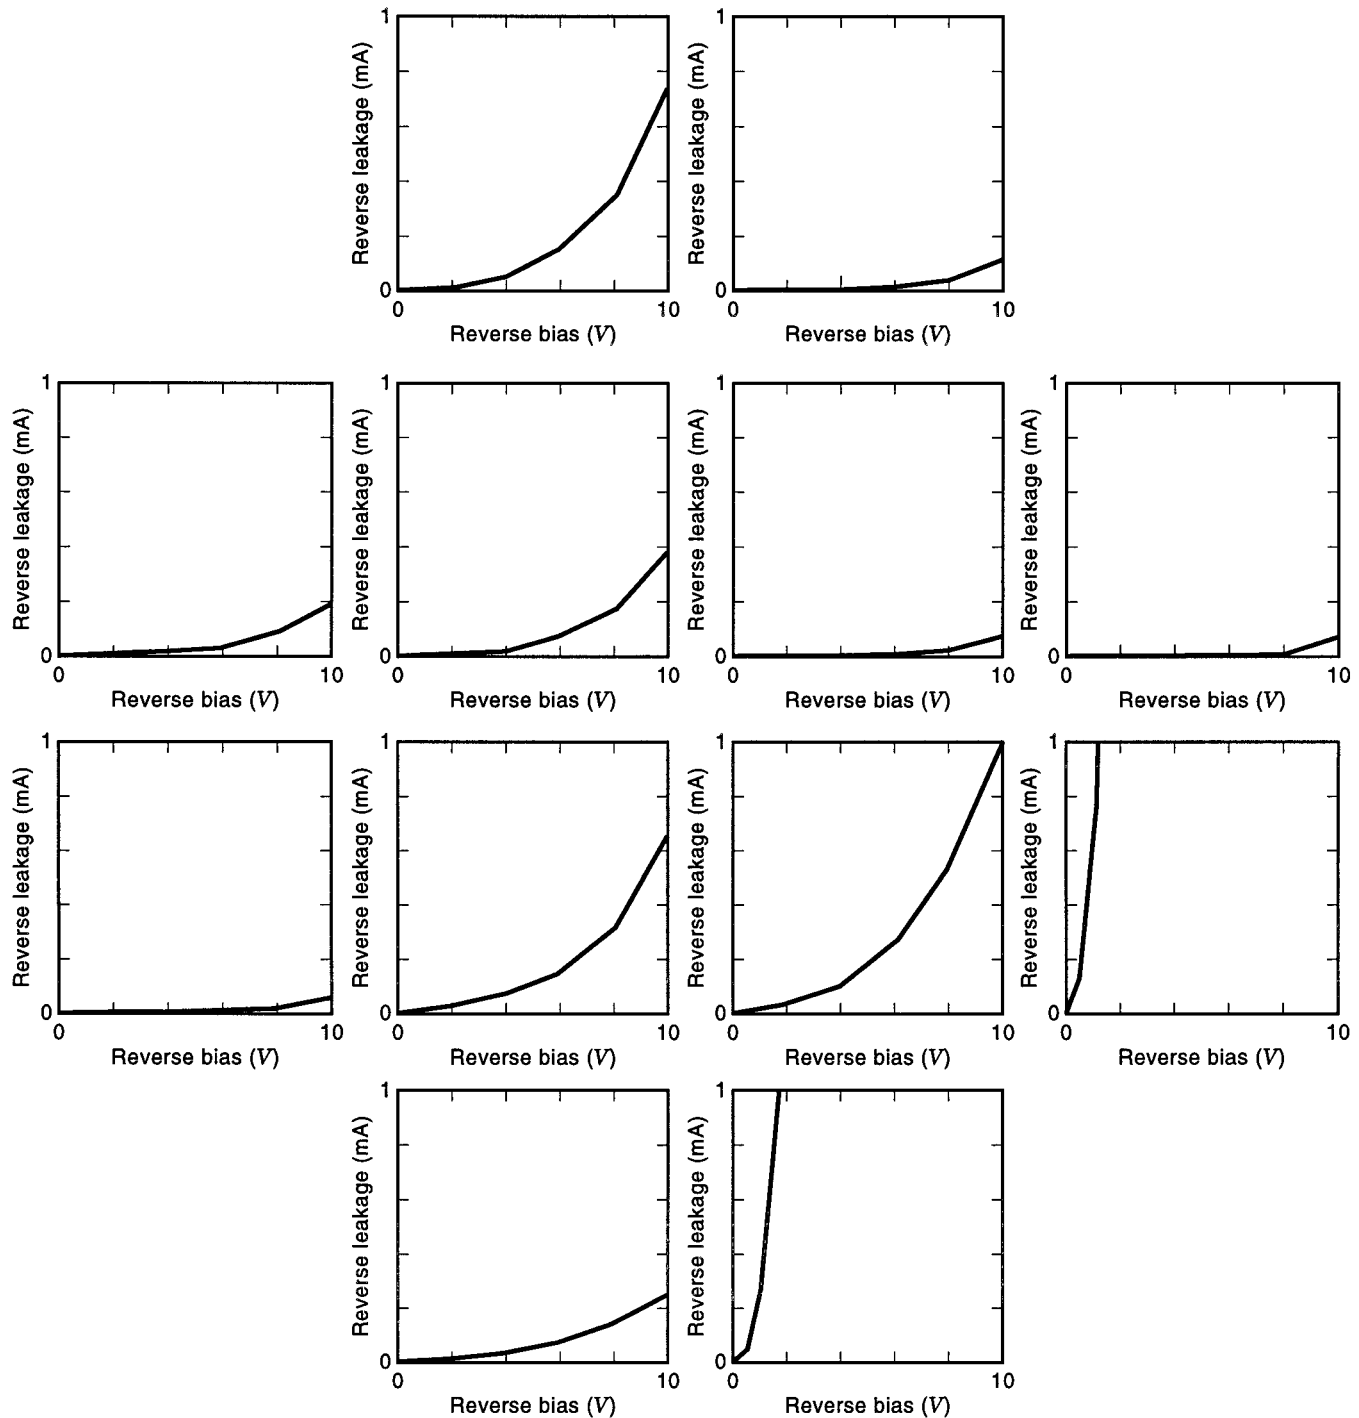


Figure 6. Reverse dark leakage current of twelve 5 mm \times 5 mm MQW modulators, with only two disabled devices, indicating that microscopic defect level for modulators is $<1/\text{cm}^2$. This indicates that arrays larger than 100,000 are possible.

We then conclude that those defects that contribute large leakage are caused by defects in the substrate that propagate during growth of the boule.

Thus we have found that microscopic material defects relevant to modulators occur once about every square centimeter. Our current design rules would allow about 350,000 diodes/ cm^2 so we conclude that GaAs/AlAs modulator arrays of at least 256×512 are intrinsically possible.

However, this is not the entire story. As mentioned, large (several microns) defects occur on the surface of MBE grown wafers. These simply interfere mechanically with the formation of devices. They can be observed by illuminating the wafer and are shown in Fig. 7 (2 in. diameter wafers shown). As can be seen, they occur in numbers of $\sim 10/\text{cm}^2$ and, therefore, limit array size to $\sim 35,000$. These large surface defects are possibly because large bits of material are ejected from the




Figure 7. Photographs of MQW modulator wafers before and after crystal growth, which has molecular beam epitaxy. The white specks are large defects on the surface occurring at densities of $\sim 10/\text{cm}^2$. For the large arrays predicted in Fig. 6, these must be eliminated, possibly by using vapor-phase growth techniques.

material sources in MBE and may possibly be avoided in other growth techniques, such as metalorganic chemical vapor deposition (MOCVD) (29).

Modulator Performance and Temperature Sensitivity

Here a comprehensive modeling and prediction of the performance of MQW modulators is presented. Performance of a modulator is characterized by contrast ratio and insertion loss. Insertion loss of these modulators can be kept low because of their simple optical coupling. It is not so much contrast ratio that is in short supply with surface-normal modulators because even without Fabry-Perot effects, which we ignore here, fairly high contrast can be achieved. However, it is the fragility of contrast ratio that reduces the performance of surface-normal modulators. The positions of the exciton, the operating wavelength, and the Fabry-Perot resonance if that is employed, all must be in very good registration to obtain high contrast. Therefore, here we optimize the tolerance of contrast to manufacturing errors or changes in operating condition. We fit the spectra of two GaAs/AlAs samples with zero-bias exciton positions λ_0 of 833.8 and 842.3 nm at each bias with Lorentzian curves, $r = 1 - (1 - r_e)/[1 + (\lambda - \lambda_e)^2/\delta^2]$, where λ_e is the wavelength of the exciton, r_e is the reflectivity at the exciton peak, and δ is its breadth. In turn, these parameters are fit with six other parameters that are a function of λ_0 . Thus we have a comprehensive model of GaAs/AlAs reflection MQW modulators with varying well width operating below the band edge (30).

Considered are modulators operating with a 3 V swing, for 0.35 μm line rule CMOS. For shorter line rules, and lower available voltage swings, we assume doubling or tripling voltage drivers as described later. Note that it is the voltage swing that is limited and not the voltage offset. The first operating mode examined allows adjustable voltage offset V_1 , and we show, for a given λ_0 , what number of wells N will result in the maximum tolerance to change in λ for a feedback circuit that changes V_1 from 0 V to 12 V. In Fig. 8, a contour plot of contrast = 2:1 is shown in the parameter space of V_1 and λ for five different cases of λ_0 . For each λ_0 , N is optimized so that the contour extends to $V_1 = 12$ V while maintaining a vertical width of 2 nm. We find that for $\lambda_0 = 850$ nm and $N = 85$, 2:1 contrast can be maintained over a change of λ of 17 nm, with a local variation of λ of ± 1 nm at any set V_1 . This would allow variations of temperature of 60°C while still allowing for variations in laser wavelength of ± 1 nm. For

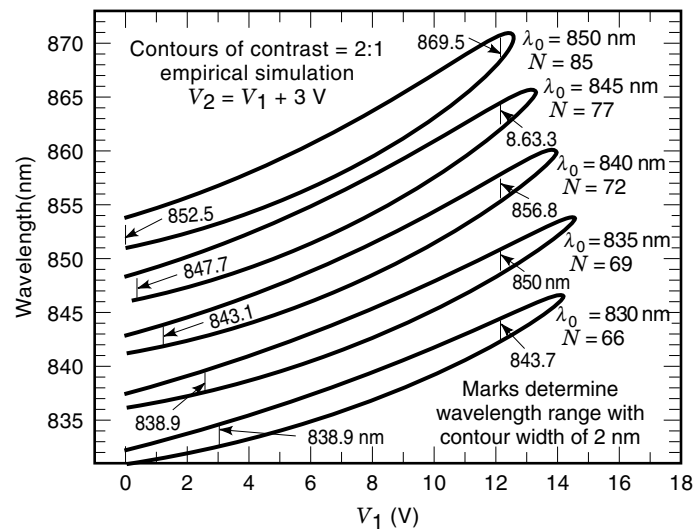


Figure 8. Contours of 2:1 optical contrast, considered necessary for intra-computer links, for five different MQW modulator designs, in the phase space of voltage offset and wavelengths (30). By allowing voltage offset to vary via feedback control, a wavelength range of 17 nm, corresponding to a temperature range of 60°C, is possible.

shorter exciton wavelengths, the range becomes smaller, showing that longer laser wavelengths are favored.

The other mode of operation studied is where the feedback circuit is not permitted. Of course, then the device is much less tolerant. Now, the offset voltage is fixed, and the sensitivity of the contrast to variations in operating condition λ and manufacture λ_0 are considered. The procedure is to plot contours of contrast ratio in λ , λ_0 space, and for an expected variation in λ_0 , measure the maximum variation in λ while maintaining contrast above a certain value. An optimized design is shown in Fig. 9, for $N = 50$ and an offset voltage of 2 V. Shown in this plot is the wafer-scale variation in exciton posi-

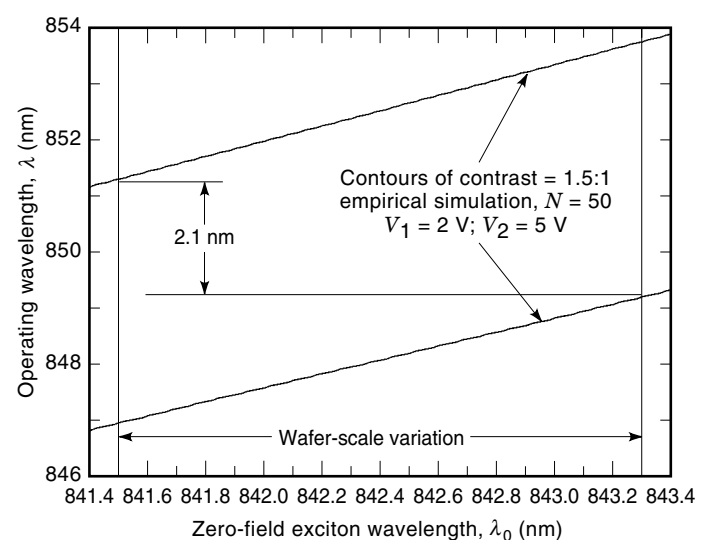


Figure 9. Contour as in Fig. 8, but in the phase space of exciton wavelength, which varies due to manufacturing control, and wavelength, keeping voltage offset fixed. The allowed wavelength variation is only 2.1 nm, too low for real systems.

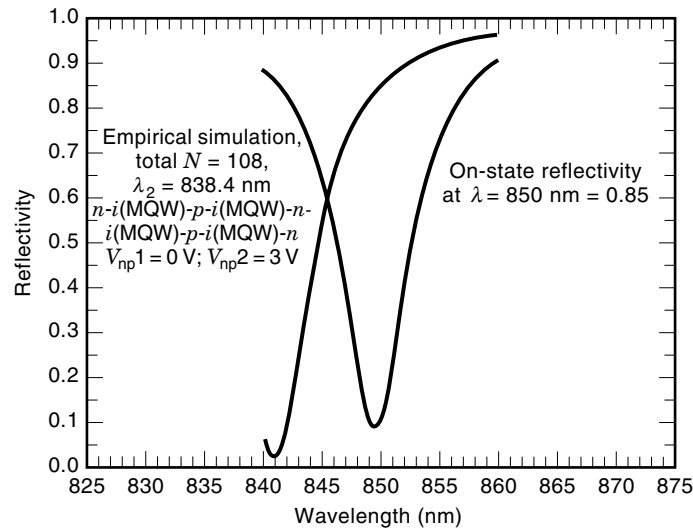


Figure 10. Simulation of spectra of advanced stacked-diode MQW modulator design, which would have sufficient temperature tolerance ($\sim 20^\circ\text{C}$) for operation without voltage offset feedback control.

tion, which, when satisfied, allows a wavelength variation of 2.1 nm while maintaining contrast above 1.5:1. We conclude that, *without feedback on the offset voltage, standard $p-i(\text{deep MQW})-n$ modulators are unusable.*

We examine other types of modulators. Shallow quantum well modulators (31) have much less wafer-scale variation of λ_0 (0.6 nm), and we find a corresponding non-feedback mode wavelength variation of 5.2 nm while maintaining 1.5:1 contrast, but with a reduced on-state reflectivity (0.33, for a unity mirror). If this is unacceptable, one may consider deep MQW modulators arranged in stacked diodes [$n-i(\text{MQW})-p-i(\text{MQW})- \dots$] (32). For four diodes, the range increases to 4.7 nm, and so is comparable to the shallow quantum well sample, whereas the on-state reflectivity for a unity mirror increases to 0.85 (Fig. 10).

Thus for a 3 V swing, with feedback circuitry to adjust the voltage offset, a wavelength range of 17 nm is possible while maintaining 2:1 contrast, corresponding to a temperature variation of 60°C . If the offset voltage is not allowed to vary, given the wafer-scale variation in exciton position, the maximum wavelength range maintaining 1.5:1 contrast of a $p-i(\text{deep MQW})-n$ modulator is 2.1 nm, deemed insufficient for most applications. A shallow quantum well sample can operate in this way over a range of 5.2 nm, but at the expense of reducing the on-state reflectivity to only 0.33. By employing stacked diode designs with deep MQWs, for a sample with four diodes the corresponding range is 4.7 nm, with an on-state reflectivity of 0.85.

CIRCUIT CONSIDERATIONS FOR OE-VLSI

SEED-Based Circuits

The original SEED and symmetric SEED displayed extremely fast optical switching. Switching times as low as 33 ps were measured in S-SEED devices with $13 \mu\text{m} \times 14 \mu\text{m}$ mesas (33). However, these devices did not provide electrical gain. As a consequence, S-SEEDs typically exhibited relatively low

sensitivity, requiring several picojoules of optical energy to be deposited on the device in order to make it switch. At these energies, absorption saturation became a concern, until it was found that this effect could be reduced by proper choice of operating wavelength (34). The end result was that the performance of switching systems based on SEED arrays was limited not by device response, but instead by the total optical power that was available in the system (35). In addition, the functionality of the arrays was typically limited, although logic operations could be performed by cascading array of these devices (36). The integration of these optoelectronic modulators with electronic FETs represented an important step in the evolution of the SEED device technology. This work was motivated by the need to reduce the optical input energy required to switch the optoelectronic device, and the desire for increased functionality in the smart pixel node. The FET-SEED technology provided a monolithic integration of metal-semiconductor FETs with MQW modulators (16,17,37). The result was a low-capacitance optoelectronic device with good sensitivity and high speed. Switching times as low as 200 ps were demonstrated in a simple FET-SEED smart pixel (38). Such devices supplied an effective input capacitance of approximately 50 fF, 60 fF to a receiver (39). Receivers were operated at about 30 fJ at 311 Mbit/s, and about 150 fJ at 1 Gbit/s, with a power consumption of approximately 3 mW (40).

The only shortcoming of the FET-SEED technology was the restricted circuit complexity resulting from the yield of the GaAs technology. In spite of the significantly improved sensitivity, it became clear that applications of the SEED device with electronic gain would ultimately be limited by the functionality and the speed of the electronic circuit. Hence, from the systems perspective, it became important to improve the *electronic* technology of the integrated device by providing the SEED device with state-of-the-art electronic FETs. Although GaAs FETs have a higher transconductance than silicon MOSFETs for a given feature size, the hope was to capitalize on the emerging deep-submicron silicon transistor technologies, the sophisticated design tools and libraries, and the rapid prototyping capability that a thriving silicon CMOS industry could provide. These were the driving forces behind the development of a silicon CMOS/SEED technology. For the reasons discussed in the previous section, the prudent course to follow was to use a hybrid integration technology. The question was whether the integration could be sufficiently "intimate" to prevent the additional parasitics associated with the flip-chip attachment from annulling the benefit of the improved transistors.

Following from the SEED work, the first CMOS-modulator chips were switching circuits. These circuits, fabricated in $0.9 \mu\text{m}$ CMOS, operated at over 250 Mbit/s (41). An input capacitance of approximately 52 fF per diode was measured for the hybrid device (42). Even with $0.8 \mu\text{m}$ CMOS technology, hybrid SEED receiver circuits provided comparable speed, sensitivity, and power dissipation to the FET-SEED receivers, with significantly reduced area. Finally, the ability to achieve a three-dimensional integrated structure of MQW diode-above circuits makes possible dense VLSI circuits and compliance with mainstream silicon design methodologies (25). More on CMOS OE-VLSI switching circuits will be discussed in the section entitled "Large Switch Experiment." This section is

completed by describing in detail the CMOS circuits unique to OE-VLSI, the optical receiver and modulator driver circuits.

Transceiver Circuits for Hybrid CMOS/SEED

A key circuit component for an OE-VLSI technology are the transceiver circuits. Because the MQW modulator may be modeled as a capacitive load element, the modulator-driver circuits to date have consisted of a simple inverter. In a CMOS technology, this circuit has zero static-dissipation (neglecting leakage currents). There is a constant component of the driver power dissipation due to the absorbed photocurrent in both “on” and “off” states. But, when operating at low optical read-powers, the receiver circuits are a greater concern in terms of electrical power dissipation than the modulator driver circuits (43). The basic reason for the receiver power dissipation is that the circuits are typically biased as small-signal amplifiers, resulting in a steady dissipation of power.

The asynchronous transimpedance receiver was chosen for initial systems experiments because it provides good sensitivity and dynamic range and because it can be operated in single-ended mode where a single diode is used to generate positive photocurrent, as well as in a differential operation mode with an additional diode at the input (Fig. 11). Small input swings in the input transimpedance stage (front end) are then amplified to logic levels at the receiver output using additional gain stages.

Single-ended CMOS/SEED receivers (Fig. 12) based on a transimpedance frontend followed by one amplification stage have been demonstrated to operate at a bit-rate of 375 Mbit/s in 0.8 μm CMOS with a dc power consumption of approximately 3.5 mW, an area of 17 $\mu\text{m} \times 18 \mu\text{m}$, a dynamic range of over 16 dB, and a sensitivity of approximately 60 fJ (25). The addition of gain stages improves the performance to below 40 fJ at 550 Mbit/s (44). Transimpedance receiver-transmitter circuits, based on two-beam differential data encoding (Fig. 13), have been built in 0.8 μm CMOS, and are capable of 1 Gbit/s transmission of digital data with a bit-error rate below 10^{-10} (45). All these circuits discussed above were provided as standard cells as part of the first workshop/foundry offered in this technology (see the section entitled

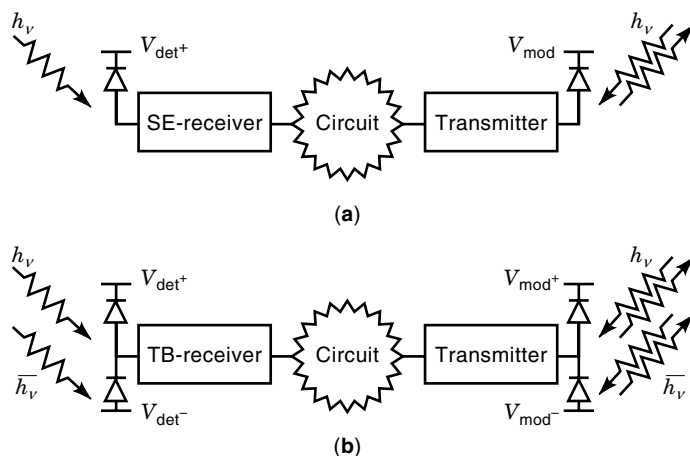


Figure 11. Schematic of (a) single-ended (SE) optical receiver; (b) two-beam (TB) receiver; V_{det} and V_{mod} are the detector and transmitter bias voltages, respectively.

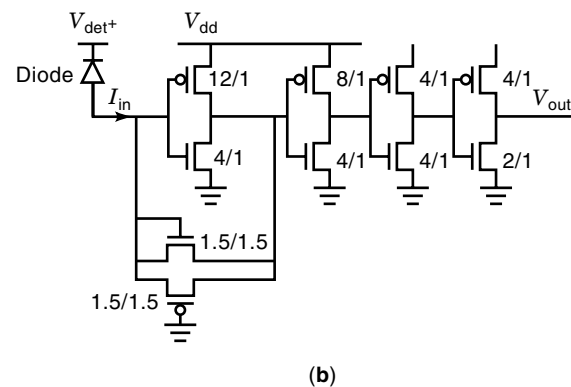



Figure 12. Single-ended transimpedance receiver standard cells: (a) 375 Mbit/s receiver with 3.5 mW power dissipation and 11 μA sensitivity (25) and (b) 550 Mbit/s receiver with 5 mW power dissipation and 6 μA sensitivity (44).

“Toward a Commodity OE-VLSI Technology”). A comparison of these circuits can be found in Ref. 46.

As supply voltages are lowered, receiver dissipation reduces and receiver bandwidth improves. However, the voltage swing provided by a normally biased inverter is reduced. This degrades the performance of a standard modulator. Equally useful for low-voltage OE-VLSI technologies are circuits that up-convert the voltage. Figure 14 presents a circuit that can provide a voltage swing that exceeds the nominal supply voltage of the CMOS technology. As shown in Fig. 14, this circuit was tested in 5 V CMOS technology and could provide swings up to twice the supply voltage in that technology. The advantage of the circuit is that it consumes no static power dissipation. Modulator-driver circuits are discussed in greater detail in Ref. 47.

LARGE SWITCH EXPERIMENT

Large switches are thought to be the first computing type to use OE-VLSI technology because of their large ratio of information channels to computing elements. In Ref. 24, a chip with 4352 optical modulators was demonstrated that switched 256 optical channels at data rates of 400 Mbit/s each, for an aggregate data throughput greater than 100 Gbit/s.

This chip was a small part of a large switching system that employed multiple electronic switching modules, each of which switched a fraction of the total channels in the system.

A switching architecture was developed that allowed smaller electronic switches to be interconnected with optical fiber via a central distribution network (48), which was the chip of Ref. 24. This chip was actually functionally a set of 16 independent crossbar switches with 16 inputs and outputs (i.e., sixteen 16×16 crossbars) on the same die. Optical fiber from each of the electronic modules was bundled and imaged onto the OE-VLSI chip. The beauty of the switching architecture was that all the buffering memory could be located on the electronic modules, whereas the OE-VLSI required no memory. Thus, it played to the strengths of the OE-VLSI technology, which emphasizes interconnect capability over memory density, which is low in CMOS compared to DRAMs.

For more on this switching architecture readers are referred to Ref. 48. Even though other technologies may be used for the distribution network, including electronic crossbars, the use of optical fiber interconnections and OE-VLSI chips represents the highest scalability in terms of total bandwidth of the system, all of which must pass through the distribution network. Interconnecting the electronic modules with optical fiber allows the physical size and clock rate of the system to expand practically without limit. A distribution module could be formed with separate optical transceiver chips and switch-

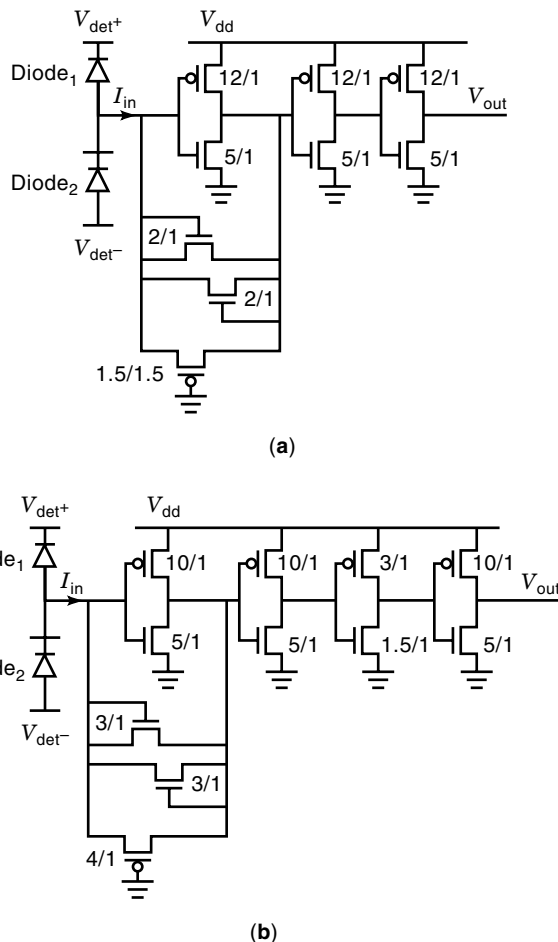


Figure 13. Two-beam transimpedance receiver standard cells: (a) 700 Mbit/s receiver with 4 mW power dissipation (41) and (b) 1 Gbit/s receiver with 8 mW power dissipation (45).




Figure 14. Voltage up-converter circuit. The circuit converts a swing of 0 V_{dd-LOW} to a swing of 0 V_{dd-HIGH}.

ing chips, but it would be much more expensive than a single OE-VLSI chip.

The distribution chip was, as stated, a collection of sixteen 16×16 crossbars. Each input was fanned out 16 times, for a total of 4096 input detectors. The fan-out could be performed electronically on the chip and was done so in later versions. For more on the exact mechanisms of the chip, the reader is referred to Ref. 24. Here, for purposes of demonstrating the capability of the OE-VLSI technology, the capabilities and requirements of the three versions of the chip as shown in Table 1.

TOWARD A COMMODITY OE-VLSI TECHNOLOGY

First CMOS/MQW OE-VLSI Foundry Shuttle

The removal of the GaAs substrate after flip-chip attachment has other positive qualities. We have performed process runs where multiple chip designs are aggregated onto a single Si chip, then a modulator array encompassing this entire large chip is processed and, after substrate removal, can be sawed into the individual chips. This is the basis for our continuing foundry runs when, in a manner like a foundry, we aggregate smaller designs into larger chips to reduce our workload. We first tested the concept of sawing a processed chip with designs from four external users. In this run, a chip designed with a perimeter of wire-bond pads and central diode field was simply sawed into four pieces. Thus it was akin to our earlier chips, with the added step of sawing the chip after bonding. This tested whether the modulators would survive the sawing procedure, which they did.

Table 1. Parameters of OE-VLSI Chips Made as Part of Large Switch Project (24)

	Chip 1	Chip 2	Chip 3
Si technology	1.0 μm	0.8 μm	0.35 μm
Functionality	256 16×1 mux nodes	64 16 I/O switches	32 16 I/O switches
Error free clock rate	450 Mbit/s	600 Mbit/s	900 Mbit/s
Total data throughput	115 Gbit/s	600 Gbit/s	404 Gbit/s
Power consumption	1 W	5 W	0.3 W

Our first foundry shuttle consisted of a 6 in. wafer of silicon that contained multiple copies of a $17 \text{ mm} \times 17 \text{ mm}$ silicon reticle. Each reticle contained 32 individual $2 \text{ mm} \times 2 \text{ mm}$ OE-VLSI chip designs as well as some larger chip designs. Designs from over 30 teams were collected into two separate $8.5 \text{ mm} \times 8.5 \text{ mm}$ subreticles. In Fig. 15 several chips are shown on the subreticle before sawing. The small rectangles are the modulators. Each subreticle contained 16 distinct $2 \text{ mm} \times 2 \text{ mm}$ CMOS circuit designs. Each of the 16 chips in a subreticle had a 10×20 array of MQW modulators bonded onto it. The horizontal and vertical pitch of the diodes were $62.5 \mu\text{m}$ and $125 \mu\text{m}$, respectively. Three columns of 80 diodes were placed between the chips, in the saw-cut lanes, for process-monitoring purposes. By forward biasing these diodes and viewing the emitted light through an infrared camera, the yield of the flip-chip attachment and substrate-removal processes could be estimated. An additional diode, placed at the reticle center was used for reflectance measurements, prior to saw-cut.

Metallization and solder deposition for the first foundry shuttle was performed at the reticle level. Flip-chip attachment and substrate removal was performed at the subreticle level. This further tested our ability to saw up the final product and also greatly extended the area that we process. The




Figure 15. Photomicrograph of subreticle from the Lucent Technologies/George Mason University coop OE-VLSI foundry, showing a close-up of several OE-VLSI chips, after integration but before final dicing. Each chip integrates 200 modulators that may be used as transmitters or detectors. Vertical and horizontal pitch of the diodes are $62.5 \mu\text{m}$ and $125 \mu\text{m}$, respectively. Each diode is approximately $25 \mu\text{m} \times 60 \mu\text{m}$ with an active area of approximately $20 \mu\text{m} \times 20 \mu\text{m}$. Any CMOS circuitry is allowed by the designer, even under the modulators, which are bonded to the top level of CMOS metallization.

diode attachment area of these chips was $7.8 \text{ mm} \times 7.8 \text{ mm}$. The substrate removal of these pieces was complicated not only by the size (linear extent) of the substrate, but also by the fact that it had to be removed entirely. In our earlier single chips, the substrate could remain at the edge of the GaAs chip because the wire bond pads were outside of the GaAs chip. Here, however, the wire-bond pads in the interior of the chip were covered by the GaAs substrate. We use a jet etcher with $\text{H}_2\text{O}_2:\text{NH}_4\text{OH}$ to remove the substrate, and it is difficult to achieve uniform removal over such a large area. This problem was solved by performing polishing of the substrate with $\text{Br}:\text{methonal}$ after flip-chip bonding to a thickness of $\sim 150 \mu\text{m}$ before the acid etching of the substrate.

Although the problems of bonding and removing the substrate of such large chips were solved, another unexpected problem specific to the foundry chip remained. Because of the specific geometry of the diode-array involved, voids were sometimes left in the epoxy. Because the foundry chips contained lanes of wire-bonding area that were separated by modulator fields, the epoxy could wick in faster or slower along certain paths, creating the opportunity for air to be trapped. We included a protective dielectric, placed on the GaAs front face between the modulators, that provided some measure of protection where voids were formed. However, damage could still be seen on some of the chips. This problem did not occur in the nonaggregated chip designs apparently because there was a single modulator array and no wire-bond pads, and the nature of the epoxy was to then produce uniform filling. We have found that the epoxy can actually be applied before flip-chip bonding because it dewets from the metallic pads. This results in fewer epoxy voids.

The ability to do batch processing for the bonding and substrate-removal steps is not only critical for foundry shuttles but also illustrates how this technology may be manufactured in a cost-effective manner. Another useful aspect of substrate removal is that after the substrate is removed, further chips may be bonded. In particular, because the modulators may be pixellated leaving spaces between them on the chip, we may bond another chip in the same area as the first chip and produce interleaved devices. Thus dissimilar devices, either separately optimized detectors and modulators or detectors and lasers, may be collocated on the Si chip.

OE-VLSI Technology Roadmap

In this section, we review a performance roadmap for this OE-VLSI technology as the silicon gate-lengths shrink. The progress of silicon CMOS technology, reviewed in Table 2, is expected to follow the projections of the SIA roadmap (6). Based on the yield considerations presented earlier, we can also estimate the expected evolution of the accompanying GaAs/Al-GaAs MQW devices. Table 3 summarizes the expected evolu-

Table 2. Projected CMOS IC Technology Parameters^a

FET Gate Length (μm)	Number of Gates (M)	Area (cm^2)	Voltage (V_{dd})	On-Chip Clock Speed (MHz)	Oxide Thickness (\AA)	Number of Electrical I/O Ports	Off-Chip Clock Speed (MHz)	Max. Power Dissipation (W)
0.7	0.15	1	5	100	130	400	75	5
0.5	0.3	2	3.3	150	100	600	110	10
0.35	0.8	4	2.8	200	80	900	150	15
0.25	2	6	2.2	350	60	1350	200	30
0.18	5	8	1.8	500	50	2000	250	40
0.12	10	10	1.5	700	40	2600	300	90
0.1	20	12	1.25	1000	35	3600	375	180

^a Generations of CMOS are expected to be spaced 3 years apart.

tion of the CMOS/MQW OE-VLSI technology as a function of CMOS feature size (gate length). The maximum number of diodes is projected to increase rapidly in the first few generations and then grow gradually as the device yield, issues of bonding large arrays, and manipulating large numbers of light beams become increasingly significant. A more detailed exposition on the scaling of the diode capacitance, as well as other device characteristics can be found in Ref. 43.

We have thus far focused on the modulator device yield, which is the main hurdle in scaling the OE-VLSI technology. In addition to the device yield, other barriers can be expected to arise because of the finite laser source power and the on-chip power-dissipation of the I/O transceiver circuits, particularly the receivers. In order to reduce the amount of laser power required for illuminating the modulators (and also to limit modulator-driver power dissipation and prevent related thermal effects in the modulators) sensitive receiver circuits must be employed at the optical-to-electrical interface of the OE-VLSI circuit. A consequence is that the receiver circuits will be a significant source of power dissipation themselves because of the need for biasing these circuits for high speed and high sensitivity.

Because the performance of the specific OE-VLSI technology is expected to be limited by the electronic circuits, we can expect that reductions in silicon feature sizes and optoelectronic device dimensions will serve to increase the aggregate optical interconnect bandwidth to a VLSI chip by improving the sensitivity and bit-rate of the receivers and by reducing their power dissipation. For a given receiver front-end, a technology figure of merit, based on the gain-bandwidth of the FET, the channel noise factor of the FET, and the detector capacitance can be defined (43). The figure of merit can be

written as ($f_T/\Gamma \cdot C_D$), where Γ is the excess channel noise factor in the FETs that constitute the receiver front-end. We note that Γ is inversely proportional to gate length and can be estimated to range between 1.8 and 3, to account for increased channel noise as the dimensions of the front-end FET are reduced. This figure of merit provides a measure of the receiver noise and hence the noise-limited receiver sensitivity of the specific OE-VLSI technology in terms of a few readily available metrics: the detector capacitance, the unity gain-bandwidth of the FETs, and the channel noise factor of the FETs. We note that for large arrays of smart-pixel receivers, the OE-VLSI technology typically cannot be exercised so as to obtain true noise-limited receiver performance. Nevertheless, for a given receiver geometry, the figure of merit is a useful means of assessing the relative receiver sensitivity performance of different OE-VLSI technologies.

To make a conservative estimate for the growth of the figure of merit, we replace the FET unity gain bandwidth with the inverse gate delay as measured from ring oscillator data. The scaling of the inverse gate delay, based on published results from experimental technologies (43), is shown in Fig. 16. Figure 17 shows the scaling of this figure of merit in terms of FET gate length. As shown in Fig. 17, the sensitivity of the receiver can be expected to improve exponentially as feature size is reduced. This sensitivity improvement is also accompanied by a reduction in power dissipation caused by reduced supply voltages and static currents in the receivers. Improved sensitivity and lower power dissipation together imply that

Table 3. Projected Evolution of Optoelectronic Technology Parameters

Feature Size (μm)	MQW Linear Dimension (μm)	Bonded Diode Capacitance (fF)	Number of Optical Diodes
0.7	18	150	1,000
0.5	14	100	3,000
0.35	12	70	6,000
0.25	10	50	12,000
0.18	8	36	24,000
0.12	7	24	40,000
0.1	5–6	20	50,000




Figure 16. Scaling trend for the gate-delay as a function of the gate length of the CMOS technology. Data correspond to measured delays from ring-oscillator data (43).

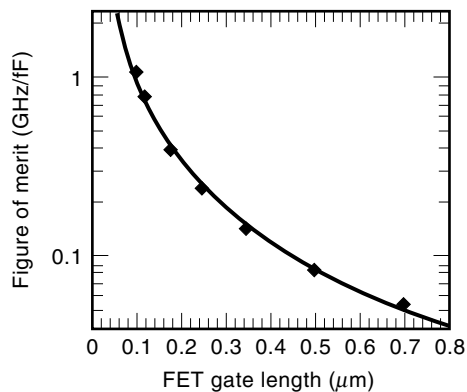


Figure 17. Scaling of the figure-of-merit for the OE-VLSI transceiver technology.

more optical transceivers (hence more optical I/O) can be accommodated on a chip for a given power budget (43). The reduction in feature size will also allow the bandwidth of each optical I/O link to grow.

As a consequence, the total optical I/O bandwidth to the VLSI chip can be expected to increase considerably. We can quantify the scaling of the total electrical and optical I/O bandwidth of the OE-VLSI chip and relate it to the computational bandwidth of the silicon chip. The latter is defined as the product of the number of gates and the on-chip clock speed of the chip. Assuming that half the projected number of pins on a chip can be used for electrical I/O at the off-chip clock speed, and assuming that the optical interconnect will function at the on-chip clock speed (a conservative assumption), we can plot these three quantities in Fig. 18. This figure depicts the expected scaling of the optical I/O bandwidth (assuming two-beam operation), the compute bandwidth (defined as the product of the number of gates and the on-chip clock frequency), and the SIA projections for high-performance electrical I/O onto and off a packaged CMOS die (6). As can be seen in Fig. 18, it appears that the growth of the I/O bandwidth of the OE-VLSI chip can be expected to match the growth of the computational bandwidth of the silicon technol-

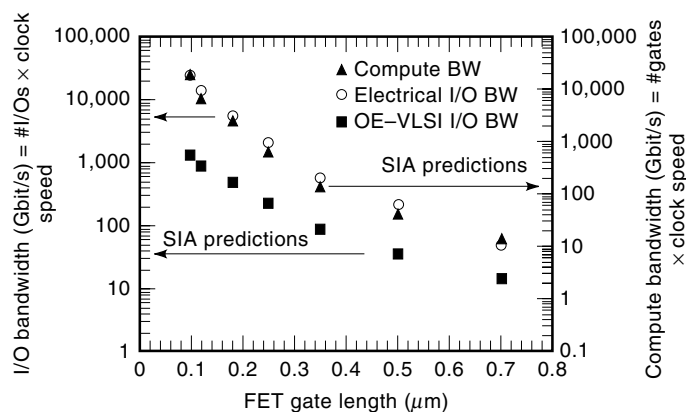


Figure 18. The I/O bandwidth and computational bandwidth of electronic and optoelectronic VLSI integrated circuits versus CMOS linewidth assuming differential optical signaling. The I/O bandwidth of an OE-VLSI chip exceeds the projected electrical I/O bandwidth of the silicon and can match the growth in the computational bandwidth of CMOS technologies (43).

ogy. This far exceeds the projected scaling of the electrical I/O bandwidth available for a VLSI circuit. As a consequence, the OE-VLSI technology should have competitive advantage for microchip applications that require high-performance communication.

CONCLUSIONS

The concept of a manufacturable technology that can provide parallel optical interconnects directly to a VLSI circuit, suggested over a decade ago in Ref. 5, now appears to be a reality. One such OE-VLSI technology, described in this article, is based on the hybrid flip-chip bonding of GaAs/AlGaAs MQW modulator devices onto silicon CMOS circuits. In this article, we have surveyed the progress in SEED-based OE-VLSI smart pixel technologies. Results from a batch-fabricated foundry shuttle incorporating multiple OE-VLSI chip designs were reviewed. This foundry represents the first successful delivery of custom-designed CMOS VLSI chips with surface-normal optical I/O technology to external customers.

Issues related to scaling of the OE-VLSI technology and compatibility with deep-submicron CMOS technologies have been discussed. In terms of the fundamental technology, the challenges are in reducing the drive voltages of the modulators to stay compatible with mainstream CMOS and to continue to improve the yield in the manufacturing and hybridizing of the MQW diodes. In terms of the circuits, the challenges will be to continue to improve receiver sensitivity while reducing power dissipation and crosstalk. A third consideration, not discussed here, is that of the systems integration, where the task will be to package systems that can efficiently transport large arrays of lightbeams to and from such chips.

Based on relatively conservative assumptions on how these components will evolve, a general conclusion is that it appears that this hybrid optical I/O technology has substantial room for continued scaling to larger numbers of higher-speed interconnects. Indeed, future OE-VLSI technologies can be expected to provide an I/O bandwidth to a chip that is commensurate with the processing power of the chip, even in the finest linewidth silicon—a task that cannot be expected from electrical interconnects. It is anticipated that the availability of such a technology will allow the production of new systems that are competitive, on a cost-per-performance basis, with all-electrical solutions.

ACKNOWLEDGMENTS

The authors express their gratitude to the Bell Labs optoelectronic-VLSI technology team members: J. A. Walker, B. Tseng, S. P. Hui, J. E. Cunningham, W. Y. Jan, L. M. F. Chirovsky, T. K. Woodward, A. L. Lentine, R. A. Novotny, L. A. D'Asaro, and R. E. Leibenguth. The authors thank R. A. Athale and K. Raj of George Mason University for their assistance in organizing the CMOS-SEED foundry and also acknowledge partial support from the ARPA/GMU CO-OP for the first CMOS-SEED foundry shuttle. The authors are also grateful to D. A. B. Miller of Stanford University for his many contributions to this work.

BIBLIOGRAPHY

1. S. Yu et al., A monolithically integrated 1×4 switchable photodiode array with preamplifier for programmable frequency filters and optical interconnects, *IEEE Phot. Tech. Lett.*, **9**: 675–677, 1997.
2. S. Chandrasekhar et al., Investigation of crosstalk performance of eight-channel p–i–n/HBT OEIC photoreceiver array modules, *IEEE Phot. Tech. Lett.*, **8**: 682–684, 1996.
3. www.vitesse.com.
4. J.-S. Rieh et al., Monolithically integrated SiGe-Si PIN-HBT front-end photoreceivers, *IEEE Phot. Tech. Lett.*, **10**: 415–417, 1998.
5. J. W. Goodman et al., Optical interconnections for VLSI systems, *Proc. IEEE*, **72**: 850–866, 1984.
6. 1993–1994 *The National Technology Roadmap for Semiconductors*, Semiconductor Industry Assoc., p. B2.
7. J. D. Dow and D. Redfield, Electroabsorption in semiconductors: The excitonic absorption edge, *Phys. Rev. B*, **1**: 3358–3371, 1970.
8. D. A. B. Miller, D. S. Chemla, and S. Schmitt-Rink, in H. Haug (ed.), *Optical Nonlinearities and Instabilities in Semiconductors*, San Diego, CA: Academic Press, 1988, p. 325; U. Efron and G. Livescu, in U. Efron (ed.), *Spatial Light Modulator Technology*, New York: Marcel Dekker, 1995, pp. 217–286.
9. K. Sayyah and U. Efron, Optically addressed spatial light modulator with a high photosensitivity and intensity adaptation range, *Opt. Lett.*, **21**: 1384–1386, 1996.
10. G. D. Boyd et al., Multiple quantum well reflection modulator, *Appl. Phys. Lett.*, **50**: 1119–1121, 1987.
11. M. Whitehead et al., Low voltage multiple quantum well reflection modulator with on:off ratio greater than 100:1, *Electron. Lett.*, **25**: 984–985, 1989.
12. R. H. Yan, R. J. Simes, and L. A. Coldren, Electroabsorptive Fabry-Perot reflection modulators with asymmetric mirrors, *IEEE Phot. Tech. Lett.*, **1**: 273–275, 1989.
13. T. L. Worchesky et al., Large arrays of spatial light modulators hybridized to silicon integrated circuits, *Appl. Opt.*, **35** (8): 1180–1186, 1996.
14. D. A. B. Miller, Quantum well self-electro-optic effect devices, *Opt. Quantum Elec.*, **22**: 561–598, 1990.
15. A. Lentine et al., A 2 kbit array of symmetric self-electrooptic effect devices, *IEEE Phot. Tech. Lett.*, **2**: 51–53, 1990.
16. D. A. B. Miller et al., Field-effect transistor self-electrooptic effect device: integrated photodiode, quantum well modulator, and transistor, *IEEE Phot. Tech. Lett.*, **1**: 62–64, 1989.
17. L. A. D'Asaro et al., Batch fabrication and operation of GaAs-AlGaAs field-effect transistor self-electro-optic effect device (FET-SEED) smart pixel arrays, *IEEE J. Quantum Electron.*, **29**: 670–675, 1993.
18. K. W. Goossen et al., Monolithic integration of GaAs/AlGaAs multiple quantum well modulators and silicon metal-oxide semiconductor transistor, *Proc. OSA Topical Meeting Photonics Switching*, Optical Soc. Amer., vol. 16, 1993, pp. 94–98.
19. Research Devices, Piscataway, NJ.
20. K. W. Goossen et al., Independence of absorption-linewidth product to material system for multiple quantum wells with excitons from 850 nm to 1064 nm, *IEEE Phot. Tech. Lett.*, **5**: 1392–1394, 1993.
21. K. W. Goossen, J. E. Cunningham, and W. Y. Jan, GaAs 850 nm modulators solder-bonded to silicon, *IEEE Phot. Tech. Lett.*, **5**: 776–778, 1993.
22. A. C. Walker et al., Development of an optoelectronic parallel data sorter based on CMOS/InGaAs smart pixel arrays, in *Optics in Computing*, vol. 8, 1997 OSA Tech. Dig. Series, Washington, DC: Optical Soc. Amer., 1997, pp. 149–151.
23. K. W. Goossen et al., GaAs MQW modulators integrated with silicon CMOS, *IEEE Phot. Tech. Lett.*, **7**: 360–362, 1995.
24. A. L. Lentine et al., High-speed optoelectronic VLSI switching chip with > 4000 optical I/O based on flip-chip bonding of MQW modulators and detectors to silicon CMOS, *IEEE J. Selected Topics in Quantum Electron.*, **2**: 77–84, 1996.
25. A. V. Krishnamoorthy et al., 3-D integration of MQW modulators over active sub-micron CMOS circuits: 375 Mb/s transimpedance receiver-transmitter circuit, *IEEE Phot. Tech. Lett.*, **7**: 1288–1290, 1995.
26. A. V. Krishnamoorthy et al., Photonic page buffer based on GaAs MQW modulators bonded directly over active silicon CMOS circuits, *Appl. Opt.*, **35**: 2439–2448, 1996.
27. H. Wang et al., Monolithic integration of SEED's and VLSI GaAs circuits by epitaxy on electronics, *IEEE Phot. Tech. Lett.*, **9**: 607–609, 1997.
28. D. L. Mathine, R. Droopad, and G. N. Maracas, A vertical-cavity surface-emitting laser applied to a 0.8-micron NMOS driver, *IEEE Phot. Tech. Lett.*, **9**: 869–871, 1997.
29. H.-C. Lee et al., Nonlinear absorption properties of AlGaAs/GaAs multiple quantum wells grown by metalorganic chemical vapor deposition, *IEEE J. Quantum Electron.*, **24**: 1581–1590, 1988.
30. K. W. Goossen et al., On the operational and manufacturing tolerances of GaAs-AlAs MQW modulators, *IEEE J. Quantum Electron.*, **34**: 431–438, 1998.
31. K. W. Goossen, J. E. Cunningham, and W. Y. Jan, Excitonic electroabsorption in extremely shallow quantum wells, *Appl. Phys. Lett.*, **57**: 2583–2585, 1990.
32. K. W. Goossen et al., Interleaved contact electroabsorption modulator using doping-selective electrodes with 25°C to 95°C operating range, *IEEE Phot. Tech. Lett.*, **5**: 181–183, 1993.
33. G. D. Boyd, L. M. F. Chirovsky, and R. A. Morgan, 33 ps optical switching of symmetric self-electro-optic effect devices, *Appl. Phys. Lett.*, **57**: 1843–1845, 1990.
34. G. D. Boyd et al., Wavelength dependence of saturation and thermal effects in multiple-quantum well modulators, *Appl. Phys. Lett.*, **63**: 1715–1717, 1993.
35. H. S. Hinton et al., Free-space digital optical systems, *Proc. IEEE*, **82**: 1632–1649, 1994.
36. A. L. Lentine and D. A. B. Miller, Evolution of the SEED technology: bistable logic gates to optoelectronic smart pixels, *IEEE J. Quantum Electron.*, **29**: 655–669, 1993.
37. T. K. Woodward et al., Operation of a fully integrated GaAs-Al_xGaAs_{1-x} FET-SEED: A basic optically addressed integrated circuit, *IEEE Phot. Tech. Lett.*, **7**: 763–765, 1995.
38. G. D. Boyd et al., Mode-locked pulse operation of GaAs/AlGaAs field effect transistor self-electro-optic effect device smart pixels and saturation considerations, *Appl. Phys. Lett.*, **65**: 3108–3110, 1994.
39. T. K. Woodward, A. L. Lentine, and L. M. F. Chirovsky, Experimental sensitivity studies of diode-clamped FET-SEED smart-pixel optical receivers, *IEEE J. Quantum Electron.*, **30**: 2319–2324, 1994.
40. T. K. Woodward, A. L. Lentine, and L. M. F. Chirovsky, 1 Gbit/s operation and bit-error rate studies of FET-SEED diode-clamped smart-pixel optical receiver, *IEEE Phot. Tech. Lett.*, **7**: 763–765, 1995.
41. A. L. Lentine et al., Arrays of optoelectronic switching nodes comprised of flip-chip bonded MQW modulators and detectors on silicon CMOS circuitry, *IEEE Phot. Tech. Lett.*, **8**: 221–223, 1996.
42. A. V. Krishnamoorthy et al., Ring oscillators with optical and electrical readout based on hybrid GaAs MQW modulators bonded to 0.8 micron CMOS, *Electron. Lett.*, **31**: 1917–1918, 1995.
43. A. V. Krishnamoorthy and D. A. B. Miller, Scaling optoelectronic-VLSI circuits into the 21st century: A technology roadmap, *IEEE J. Selected Topics in Quantum Electronics*, **2** (1): 55–76, 1996.

44. A. V. Krishnamoorthy et al., Operation of a single-ended 550Mbit/s, 41fJ, Hybrid CMOS/MQW receiver, *Electron. Lett.*, **32**: 764–765, 1996.
45. T. K. Woodward et al., 1 Gb/s two-beam transimpedance smart-pixel optical receivers made from hybrid GaAs MQW modulators bonded to 0.8 μ m silicon CMOS, *IEEE Phot. Tech. Lett.*, **7**: 763–765, 1995.
46. T. K. Woodward et al., Optical receivers for optoelectronic VLSI, *IEEE J. Selected Topics in Quantum Electronics (JSTQE)*, **2**: 106–115, 1996. [Special Issue on Smart Pixels]
47. T. K. Woodward et al., Modulator-driver circuits for optoelectronic VLSI, *IEEE Phot. Tech. Lett.*, **9**: 839–841, 1997.
48. T. J. Cloonan and G. W. Richards, *Terabit per second packet switching having distributed out-of-band control of circuit and packet switching communications*, US Patent number 5,537,403.

KEITH W. GOOSSEN
ASHOK V. KRISHNAMOORTHY
Lucent Technologies