Keynote Address

Quantum well devices for optics in digital systems

D. A. B. Miller

AT&T Bell Laboratories,
Holmdel, New Jersey 07733

ABSTRACT

Semiconductor quantum well devices are an important option for incorporating optics in digital systems. Some underlying reasons why optical devices can benefit digital systems are discussed. The potential for quantum well optical modulators and switches to help realize these benefits is summarized briefly. The importance of integration of optical and electronic components for any large scale use of optics in digital systems is emphasized.

1. INTRODUCTION

Semiconductor quantum well structures are one of the important technology options for the use of optics in digital systems. In this article, I will summarise some of the basic physical reasons for trying to use more optical and optoelectronic devices in digital switching and processing systems and briefly discuss the potential role of quantum well modulators and switches. The availability of suitable devices is crucial for realizing the potential advantages of optics. The actual discussion of quantum well devices will be very brief in this article because these devices are more extensively discussed already in the published literature.[1-4]

An important theme throughout this discussion is that the abilities of optics and electronics are largely complementary. This argues for combining the two to get better overall system performance. One clear conclusion of this argument is, however, that we need integration of optics and electronics to see the greatest benefit. This in turn has important consequences for the device technologies we should use.

2. CONTRASTS BETWEEN OPTICAL AND ELECTRONIC PHYSICS

We can start by understanding the basic physical differences and similarities of optics and electronics. First of all, it is arguably true that essentially all information between devices in either optical or electrical systems is carried by electromagnetic waves of one form or another. This is obvious in optics, but is true also for electrical lines; the velocity of propagation of information is not the electron velocity, but is set instead by overall electric or electromagnetic parameters of the lines, such as dimensions, dielectric constants and resistance. Only in certain special cases, such as electrons on their way from the filament to the screen in a cathode ray tube, is the information really carried by electrons over substantial distances. (Information flow *within* a device may well be by electrons, however.) It is also arguably true that all interaction (and hence logic) is done by electrons. (In principle, we could use other charged fundamental particles, although we are safe in neglecting muons for the time being!) This is obvious in electronics. As we look at optical effects more closely, we find that even the classical nonlinear optical processes involve the coherent motion of electrons in charge clouds; without the electrons there would be no interaction between the photons. In this sense, there is no such thing as an "all-optical" effect. Because some nonlinear optical effects only involve such transport over subatomic distances, they can in principle be very fast. Typical electronic devices may be slower internally because the transport is over longer distances (e.g., 1 μm).

If optics and electronics both use electromagnetic waves for communication and electrons for logic, what then are the differences between them? There are two clear qualitative differences, both of which result from the very high frequencies of optical waves. One of these is classical, the other is quantum mechanical. The classical effect is that the wavelength is short in optics, which has two consequences: (i) we can use dielectrics to confine or guide the waves because we can work with structures larger than the wavelength; (ii) we can use free space optics (such as lenses and mirrors) to handle many waves at once. In electrical interconnection, the wavelength is usually comparable to or larger than the entire circuit of interest. To confine the information to flow along the paths we wish, we essentially use the effectively infinite dielectric constant of conductors to confine the information to much less than its "free-space" wavelength. We are more used to thinking of wires as being conductors of current that carry electrons from place to place. In fact, as far as information flow is concerned, as discussed above it is not the electrons that carry the information but the waves, in which case the function of the wires is to confine and guide the waves. (Because optics, either in free space or in waveguides, does not require conductors to confine the waves, superconductors are of course not necessary for low loss, high speed optical interconnections.) In electrical interconnection, it is basically unthinkable to consider communication within the system in which the waves are allowed to leave the wire "waveguides" and propagate in free space; the wavelength is so large that the waves would diffract over the entire volume of the board or processor, communicating everything to everywhere. In optics, of course, this is no longer true; we can, for example, successfully image many points in one plane to corresponding points in another plane with a simple lens, even although the waves from all points overlap at the lens. It may even be that there is a psychological resistance to the use of such free space interconnection because in normal electrical interconnection we dare not even contemplate letting the information leave its controlled "guided wave" environment; in fact, a significant part of the electronic system engineer's time is devoted precisely to avoiding cross-talk. It has, however, been explicitly demonstrated that a non-trivial digital system can be made in which every interconnect is done by free-space optics.[5] The "cross-talk" of the optical waves turned out not to be an issue in this demonstration.

The quantum mechanical difference in optical and electrical interconnection is that electrical systems perform classical detection (they measure voltage) whereas optical systems use quantum detection (they count photons). The photon energy of the individual photons in an electrical interconnect is of course much less than the thermal energy, kT, because the base frequency is so low; hence we must measure what is essentially a thermodynamic average quantity, the voltage, and we must make the voltage much larger than kT/e for noise immunity. In optics, the photon energies are large, and it is probably no accident that the optical photon energy in electron volts is essentially the same as the logic voltages we use in electronics; in both cases, this gives us thermal noise immunity. Because the photon energy is so much larger than kT, we may detect instead by counting photons. Classical voltage then becomes irrelevant in the wave. This has some important consequences for interconnection energy because of "quantum impedance conversion",[6] which I will discuss briefly below.

3. FEATURES OF OPTICS

Perhaps the hardest task for optics would be to compete with electronics for logic.[7] This does not seem likely for complex systems at the present time, although the better optical logic gates that do exist have abilities complementary to electronics. A recently-demonstrated "soliton dragging" logic gate using short optical pulses in optical fibers can perform real logic operations on sub-picosecond optical pulses,[8] a time scale difficult to achieve electronically. Quantum well symmetric self-electrooptic effect devices (S-SEEDs)[4] have been made in very large parallel arrays such as $32 \times 64 = 2048$ devices,[9] all of which could be accessed in parallel with high quality optics. Such a number of chip-to-chip interconnections would be difficult to contemplate electrically. Arrays as large as $128 \times 256 = 32$ K of devices have recently been made. S-SEEDs also have switching speeds as fast as 35 *ps*.[10] One common historical criticism is that optics did not have a "three-terminal" logic gate; both the soliton gate and the S-SEED are, however, true three-terminal devices, albeit in unusual ways. As a consequence, they can be cascaded, and functional digital systems have been built with S-SEEDs.[5] Both of these devices can be run with optical signal energies ~ 1 *pJ*. Such an energy is low enough to be interesting, although it does not compete with the logic energy of the best electronic gates. It is, however, significantly lower than the communication energy between chips in electronic systems, which can readily be in the range of 100 *pJ*. One feature of optical logic devices is that there is no distinction between the logic energy and the communication energy; no line drivers are needed in optics. This is actually a consequence of the "quantum impedance conversion" arguments.[6] Optical logic devices remain experimental as far as systems use is concerned, partly because their strengths are so different from those of electronic gates; systems that took full advantage of optical gates would look very different from current electronic systems. Optical logic would seem to be useful, for example, in systems that are very communications-intensive but in which the logic complexity is low. Example systems might include telecommunications switching or image processing. Now that real, high-performance optical logic arrays are available, this may stimulate systems research to take advantage of them.[11]

As we look at the idea of using optics to help interconnect electronic systems, we can see many potentially beneficial features. First of all, it is clear that electronic system design and performance is significantly constrained by the difficulty of interconnection, especially in large, high performance systems. Issues such as clock skew, frequency dependent loss and crosstalk, pulse dispersion, impedance matching, ground "loops" (essentially the difficulty of providing a uniform ground potential throughout the electronic system) are all important. Chip power and area are becoming increasingly dominated by interconnects, both for connections on the chip itself and for driving the connections off the chip. Numbers of off-chip interconnects are limited. Electrical interconnect technology will, of course, continue to advance, but all of these difficulties result to a greater or lesser degree from the physics of electrical interconnects. Optical interconnects, however, have the potential to avoid all of these problems because the physics is different.

Optical systems do not suffer from frequency-dependent loss (there is no skin effect because there are no conductors), nor frequency-dependent crosstalk (the carrier frequency is so high that the modulation makes no difference). Pulse dispersion is also not a significant issue for optics over the distances within processors. Because of the quantum detection involved in optics, it can fundamentally address the issues of ground isolation and interconnect energy.

The fundamental issue of the effect of the quantum detection inherent in optical interconnects requires some discussion. Essentially, the quantum detection effectively gives us an unusual, unique and in some respects ideal transformer. First, it is clear that optical transmission and detection provides isolation of voltage levels between two circuits while still functioning all the way down to zero frequency (d.c.), a characteristic we would want of an ideal transformer. The ability to work down to arbitrarily low frequencies is important for digital systems since we wish to be able to handle an arbitrarily long string of "ones" or "zeros". This is the function of a simple optical isolator, usually consisting of a light emitter and a photodetector, and it enables us to avoid the issue of uniform ground potentials throughout the system. We could get the same effect by modulating any high frequency carrier and detecting the resulting received power, but the optical devices make this particularly easy. The fact that we emit (or modulate) photons avoids the need for a high-frequency classical oscillator.

A more subtle way in which optical devices can operate as a transformer is that they can effectively match impedances. This is especially important because a large fraction of the power involved in large digital systems is devoted to driving the low impedance or high capacitance of transmission lines. The optical devices, because they are quantum devices, allow us to match high impedance devices into and out of the unavoidably low impedance of wave propagation. The low impedance (or high capacitance) of electrical lines comes directly from the low impedance of free space (377 Ω). Because the impedance and capacitance only vary logarithmically with the dimensions of the lines (once the separation of the lines is somewhat larger than their dimensions), we can conclude that, as a rule of thumb, the capacitance of a line is ~ 100 pF/m, and its impedance (if it is a low loss line) is ~ 50 Ω. There is little we can do to improve these numbers, regardless of how well we engineer the lines. It requires 20 *pJ* of energy to send a 1 *V*, 1 *ns* pulse down a 50 Ω line, or to fully charge a 30 *cm* line to 1 *V*. This number is much larger than the logic energy of a good electronic logic device, a number that may be ~ 10 *fJ*. Hence, if a system requires much communication internally, the communication energy can easily dominate over the logic energy in a purely electronic system.

The reason for the effective impedance transformation inherent in optical devices is because they count photons rather than measuring a classical voltage. This is discussed in detail elsewhere, so I will not repeat the argument here.[6] The key point can be understood by example. Suppose, for example, that we had a device with a 1 $G\Omega$ input impedance connected to a good photodiode, and that we shone a 1 *nW* light beam on the diode. This would generate approximately 1 *nA* of current in the diode, giving a 1 *V* signal across the 1 $G\Omega$ load. This is remarkable, because the classical voltage in a 1 *nW* light beam is approximately 600 μV. Hence we can see that the photodiode performs an effective impedance matching from the low impedance of free space to the high impedance of the device. The example chosen here is extreme for clarity; in practice one would use some smaller impedance in order to get a suitably low time constant. Actually, the important parameter in the optical case is really the energy of communication, which is essentially the energy required to generate enough charge in the photodetector to change its internal voltage by a sufficient amount (e.g. a logic level voltage ~ 1 *V*), an energy that is independent of speed. Optics therefore requires a given energy per bit, almost independent of speed.

The same transformation is also possible at the transmitting end, either with an ideal quantum emitter (that emits one photon for every electron passed) or a suitably efficient small modulator. Light sources continue to get better, and the low threshold current necessary to take full advantage of this effective impedance transformation may yet be achieved.

Quantum well modulators can already take advantage of this impedance transformation at the transmission end because they are not threshold devices.

Instead of being concerned about line capacitance, in the optical case we are more concerned about device capacitance. To take full advantage of the optical devices in reducing communication energy requires optical devices that can be very small and efficient, and that can be integrated readily with electronics. This is a major challenge, but quantum well devices are well suited for this. One conclusion is that for all distances longer than a certain critical distance, we should use optical devices for interconnection (if we are only concerned about interconnect energy). The critical distance depends on the optical device capacitance, and occurs essentially at the point where the device capacitance becomes less than the line capacitance. For $10 \times 10\ \mu m$ devices and $1\ V$ logic levels, this distance is $200\ \mu m$. Of course, there are many considerations other that energy that determine the choice of interconnect technology, but this argument does point out one fundamental advantage of optics that we may be able to exploit in the future as integration technology improves. At the present time, we seldom can take advantage of this energy argument because the device technology is not good enough for most devices. Typical laser diodes and, even more so, light emitting diodes, have effective impedances of much less that $50\ \Omega$, and so do not give us a direct energy advantage from this mechanism. The S-SEED, in simple systems experiments, already exhibits this energy advantage however; it can be switched with input energies as low as $1\ pJ$, even with a $> 5\ V$ internal logic level, and the energy is independent of how far we wish to send the signal. There is no distinction in such a device between the logic energy and the communication energy, in contrast to electronic devices.

The more radical concept of using the two-dimensional abilities of optics in interconnection brings other potential features.[11] Clock skew can be avoided almost entirely in an imaging interconnect, since all paths from image to object are of almost equal lengths. Very large numbers of interconnections are also possible in principle with two-dimensional, free-space interconnects. Thousands of lines connected by a single lens are feasible. Although the initial cost of such an optical system might be high, the cost per line could be low; in contrast to many interconnection technologies, the cost is certainly much less that linearly proportional to the number of lines connected. Such optics is also good at global interconnect patterns, such as perfect shuffles, patterns that are more difficult in electrical interconnects because of the large numbers of crossing wires. Free space optics does have more difficulty with irregular interconnects, something that electronics finds relatively easy, at least over short distances. Irregular interconnects are not fundamentally necessary in processors, although there is a clear engineering trade-off; more logic levels or gates would probably be required if there were no irregularity possible. (An arbitrary connection between a set of inputs and a set of outputs can always be established using a switching fabric, which consists entirely of fixed, regular physical interconnections between logic gates.) Certainly, current electronic designs make great use of irregularity. As a meaningful free-space optical interconnect technology becomes possible, system design may change, however, because of the other benefits. The most likely scenario is, of course, some mix of regular and irregular optical and electronic connections.

4. OPTICAL DEVICES FOR DIGITAL SYSTEMS

It seems most likely that the introduction of optics into digital systems will be in hybrid optical/electronic configurations. This is because of both the state of the technology and the potential flexibility of such systems. It is sometimes believed, however, that such hybrid systems will have poor physical performance; historically, the conversion from optics to electronics and back again has often been an expensive and inefficient process that can actually limit the speed of the system. This, however, is not fundamental. To avoid such problems, the devices must be well integrated. The SEED provides an existence proof of this. It is simply a very well integrated optical-to-electronic-to-optical converter. As mentioned above, existing devices can run with energies of picojoules and speeds of 35 ps. As we are able to incorporate more complexity yet still keep the quality of integration, this conversion barrier will disappear. The key is, however, good integration technology.

The basic concept of the SEED is to combine a quantum well optical modulator with a photodetector, possibly with some intervening circuitry, to provide a device with both optical inputs and optical outputs. Light shining on the input photodetector will cause a voltage change across the modulator, hence changing the transmission. Both the photodetector and the modulator are usually constructed in the form of diodes. Such a device becomes practically interesting if it is integrated. As mentioned above, significant levels of integration have been achieved for a relatively simple device like the S-SEED. Many other configurations of SEEDs have also been investigated, performing various digital and analog functions.[4]. These simple SEEDs avoid the need for any electronic gain by using internal positive feedback to provide switching and gain. This makes the devices easier to fabricate. Such devices typically use the quantum well diode simultaneously as a photodetector and a modulator.

The SEEDs have relatively low operating energies for optical devices for two reasons. First, the modulation mechanism, usually the quantum-confined Stark effect (QCSE), only requires micron thicknesses of semiconductor quantum well material for significant modulation, and the QCSE is a large effect, requiring in practice drive voltages of ~ 1 - 20 V. This mechanism and related effects are discussed elsewhere. Secondly, the photogenerated charge is allowed to move macroscopic distances within the device (e.g., 1 μm), giving a large overall polarization change in the active material of the device for only a small optical energy. Specifically, carriers generated in one of the, say, 50 quantum wells in the diode move to screen the field in all 50 wells, so that we are essentially using them 50 times; this is a common feature of what are sometimes referred to as "charge transport" nonlinearities. Viewing SEEDs as optical nonlinearities, they consequently require about 100 times less optical energy density (e.g. 10 $fJ/\mu m^2$) than other semiconductor nonlinearities such as absorption saturation (e.g., 1 $pJ/\mu m^2$). These arguments are given in more detail elsewhere.[4]

If we start to introduce more electronic elements between the input photodetectors and the output modulators, we can get two additional benefits. First, from the physical point of view, we can use electronic gain to reduce the required optical input energy. For example, existing S-SEEDs typically have operating voltages of 5 - 15 V, which means we must provide enough optical energy to discharge the internal capacitance of the S-SEED over this relatively large voltage. By using electronic gain at the photodetector, we can reduce the input voltage requirement to perhaps less than 1 V, reducing the optical energy requirement proportionally. Note incidentally that we have to be careful in digital systems about trying to use too much gain for three reasons: (i) large gain, usually achieved with a multiple stage amplifier, will introduce delay; (ii) we probably must keep the system d.c.-coupled (to handle the arbitrary long strings of "ones" or of "zeros"); (iii) the input voltage swing needs to be large enough to exceed thermal noise and threshold voltage fluctuations - error correction is probably out of the question because of the complexity and delay it would introduce. This is therefore a very different problem from designing a communications receiver. Nonetheless, reduction of the optical input energy by a factor of 10 to 100 is probably feasible using electronic gain without incurring too many penalties in the digital system design.

The second reason to want to introduce more electronics is to achieve more logical complexity between inputs and outputs. For most systems that we are used to, we do not need global, longer distance interconnection between each and every device. Hence it is more efficient to have local "functional blocks" that are then connected over longer distances. The extent to which this phenomenon is a consequence of the limits of electronic technology or a natural requirement of systems is not clear, and should remain a topic of debate. Certainly the point at which one leaves a local block for more global interconnection must depend on the technology, and also on the problem to be solved. The above discussion about quantum impedance conversion suggests physically that electrical interconnects are good up to some distance, beyond which they should be replaced by optical ones, which gives a physical argument for a particular block size. Regardless of where the optimum point is for crossing from electronics to optics, it is reasonable to say that the largest range of different kinds of systems can be made if we are able to choose arbitrarily where we go from optics to electronics and *vice versa*. Hence we want the option to put arbitrary electronic logic between the optical inputs and outputs. Thus we can see that a concept such as the SEED evolves smoothly into the idea of optically interconnecting electronic blocks, provided only that we have a good integration technology. This technology is starting with quantum well modulators. Already there have been successful demonstrations of both integration of quantum well modulators with field effect transistors[12] and growth of modulators on silicon substrates.[13] Both of these approaches are targeted at integrating modulators and photodetectors with sophisticated electronic technology.

An important question in any optoelectronic integration technology is whether we should use modulators or directly driven lasers as the optical output devices. Modulators have the disadvantage that they require an external optical power source and additional optics to deliver the beams to the modulators. It may also be necessary to match the wavelength of the laser to that required for the modulator, and this in turn may impose temperature constraints. On the other side, there are several benefits to a modulator-based approach. Benefits of a centralized optical power supply include the following: (i) all excess power required to run the laser is dissipated off-chip; (ii) high-efficiency lasers can be used because we are not constrained by the requirements of integration; (iii) the centralized laser can be used to provide centralized clocking of the entire system by pulsing the laser; (iv) such wavelength control as is necessary need only be done on one laser (wavelength control is very important in systems employing holographic optics); (v) laser isolation only has to be done once. Modulators are not threshold devices, which means they can be run efficiently at arbitrarily low powers; they also do not require sophisticated processing in the plane of the layers (as is typical for both waveguide and surface-emitting lasers), making them consequently easier to integrate in principle. A final issue concerns reliability. If we are to make large scale use of optics in digital systems, we must expect that we will have a large number of optical devices, perhaps many thousands. This imposes strict lifetime criteria. A system (without redundancy) consisting of 10^4 devices each with a mean time to failure of 10^5 hours would have a mean time to system failure of about 10 hours, which is clearly unacceptable. Modulators have not so far shown any reliability problems in laboratory use, although there has been little formal study of their long term reliability. Large arrays of SEEDs have not shown failure problems in systems experiments. Fundamentally modulators may be more reliable because they are not threshold devices, and can therefore be run at low power densities to increase lifetime if necessary. GaAs/AlGaAs quantum well modulators have been successfully grown on silicon substrates,[13] operating continuously for over 10000 hours without apparent degradation. Lasers and light-emitting diodes on silicon typically have very short lifetimes. One issue with light emitters on silicon substrates is that some of the recombination in the device is non-radiative, occurring instead through traps, and hence transferring energy into the defects in the crystal. This may cause the dark-line defects to propagate, killing the device. Quantum well diode modulators have essentially no recombination inside the active region of the device because the photocarriers are swept out of the diode region, and so this failure mechanism may be missing in such devices.

5. CONCLUSIONS

I have argued that the basic physical differences between optics and electronics give strong reasons for the use of optics in digital systems, especially for interconnection. The challenge for devices is significant. We need devices that in some cases are quite different from our current ones. In particular, for any large scale use of optics in electronic systems, we need a good integration technology. Quantum well modulator-based devices are attractive because of the relative ease with which they can be made, their good physical performance, and their strong prospects for integration with electronics.

REFERENCES

[1] For an extensive discussion of the physics of quantum well linear and nonlinear optical and electrooptic properties see S. Schmitt-Rink, D. S. Chemla and D. A. B. Miller, "Linear and Nonlinear Optical Properties of Semiconductor Quantum Wells", Advances in Physics, 38, 89-188 (1989).

[2] For an introductory discussion of quantum well optical processes and devices, see D. A. B. Miller, "Optoelectronic applications of quantum wells", Optics and Photonics News, 1, 7-15 (Feb. 1990).

[3] For a recent longer discussion of quantum well optoelectronic processes and devices, see D. A. B. Miller, "Quantum well optoelectronic switching devices", Int. J. High Speed Electronics, 1, 19-46 (1990).

[4] D. A. B. Miller, "Quantum well self-electrooptic-effect devices", Opt. Quantum Electron., 22, S61-S98 (1990).

[5] M. E. Prise, "Optical computing using self-electro-optic effect devices", in "Digital Optical Computing", e.d. R. A. Athale, SPIE Critical Reviews of Optical Science and Technology, CR35, 3-27 (1990).

[6] D. A. B. Miller, "Optics for low-energy communication inside digital processors: quantum detectors, sources, and modulators as efficient impedance converters", Optics Lett., 14, 146-148 (1989).

[7] D. A. B. Miller, "Device requirements for digital optical processing", in "Digital Optical Computing", ed. R. Athale, SPIE Critical Reviews of Optical Science and Technology, CR35, 68-76 (1990).

[8] M. N. Islam, "All-optical cascadable NOR gate with gain", Optics Lett., 15, 417-419 (1990).

[9] A. L. Lentine, F. B. McCormick, R. A. Novotny, L. M. F. Chirovsky, L. A. D'Asaro, R. F. Kopf, J. M. Kuo and G. D. Boyd, "A 2 kbit array of symmetric self-electrooptic effect devices", IEEE Photonics Tech. Lett., 2, 51-53 (1990).

[10] G. D. Boyd, A. M. Fox, D. A. B. Miller, L. M. F. Chirovsky, L. A. D'Asaro, J. M. Kuo, R. F. Kopf and A. L. Lentine, "33 ps optical switching of symmetric self-electro-optic effect devices", Appl. Phys. Lett., 57, 1843-1845 (1990).

[11] N. Streibl, K.-H. Brenner, A. Huang, J. Jahns, J. Jewell, A. W. Lohmann, D. A. B. Miller, M. Murdocca, M. E. Prise and T. Sizer, "Digital optics", Proc. IEEE, 77, 1954-1969 (1989).

[12] D. A. B. Miller, M. D. Feuer, T. Y. Chang, S. C. Shunk, J. E. Henry, D. J. Burrows and D. S. Chemla, "Field-effect transistor self-electrooptic effect device: integrated photodiode, quantum well modulator and transistor", IEEE Photonics Technol. Lett., 1, 62-4 (1989).

[13] K. W. Goossen, G. D. Boyd, J. E. Cunningham, W. Y. Jan, D. A. B. Miller, D. S. Chemla and R. M. Lum, "GaAs-AlGaAs Multiquantum Well Reflection Modulators Grown on GaAs and Silicon Substrates", IEEE Photonics Technol. Lett., 1, 304-306 (1989).