

Physical Reasons for Optical Interconnection

David A. B. Miller
 Department of Electrical Engineering
 Ginzton Laboratory
 Mail Code 4085
 Stanford University
 Stanford, CA 94305

There are many physical reasons for using optics for interconnection within otherwise electronic information processing machines. These reasons are summarized and categorized based on the fundamental physical differences of the higher frequency, shorter wavelength, and larger photon energy of optics compared to electrical interconnections. Optics arguably solves or mitigates most of the physical problems that limit electrical interconnections, such as signal and clock distortion, skew and attenuation, impedance matching, cross-talk, power dissipation, wave reflection phenomena, interconnect density limitations, and voltage isolation. Optics also creates radical opportunities that have no analog in existing electrical interconnections, such as free-space parallel interconnections, wavelength-division multiplexing, and the use of short optical pulses for timing and improved interconnect performance.

1. Introduction

Optics is undeniably useful for communicating information, as is clear from fiber optics and vision. Optical interconnections in information processing machines are, however, presently still relatively rare. That we have not used optics more in that application is usually for sound reasons: optics may be currently impractical because appropriate devices, though physically possible, may not yet have been realized or are not yet manufacturable; it may be unfamiliar to designers; the architectures that may exploit it most effectively may not yet be sufficiently developed or understood; machines or applications requiring interconnect capacities sufficiently large to benefit from optics may not yet have emerged; electrical interconnections work very effectively in a broad range of situations, are well understood, and represent a mature and cost-effective technology; optics may simply cost too much today.

It is not the purpose of this article to debate the current practicalities and the details of the technology in optical versus electrical interconnections. Such issues have been well reviewed elsewhere [1]. For the immediate future, those practicalities will determine whether optical interconnections can be used. Many important research challenges remain to make practical systems that could exploit many of the benefits of optics for interconnection [1]. Rather the goal is to discuss reasons that can be deduced from basic physical differences between optical and electrical physics why we would want, or even be compelled, to use optical interconnections in information processing in the future. Physics can provide some absolute answers to limitations or benefits of one approach or another, in some cases independent of technology or at least technological details. It can help delineate broad areas where one approach will work well and another will be increasingly impractical. It can also point out opportunities that could be exploited, hence motivating technological work.

The physics of optical and electrical approaches to interconnection are different in many ways. Optics arguably has many potential benefits to offer, and only a few of these have

been exploited so far. In this article, we use the comparison of the basic electrical and optical physics as a framework to catalog various features, both actual and potential, of optics in interconnection. This approach will show several benefits, and will also help in clarifying some misconceptions about optical advantages. Some of these arguments have previously been presented briefly by the author [2]. Notable prior discussions of optics for interconnections and the reasons for it include Goodman et al. [3] and the comparison of optical and electrical interconnects by Feldman et al. [4]. Some, too, are known, often intuitively, to those involved in optics, but are not readily available or conveniently collected. To make this article useful (and partly tutorial) for the broadest audience, we will discuss these arguments here, though we will try to avoid unnecessary detail.

There are some other notable arguments that we will not discuss here, such as the discussions by Ozaktas [5] of reasons for optical interconnection based on the scaling of thermal transfer.




Fig. 1 Illustration of types of optical and electrical propagation and their velocities.

2. Electrical and Optical Interconnection Physics - Similarities

A useful way to approach the potential benefits of optics for interconnection is to start with the basic similarities and differences in optical and electrical physics. At the most basic level, optical and electrical physics are very closely linked. All

optical signals are electromagnetic, and the same electromagnetic physics underlies electrical systems. In practice, in both the electrical and optical cases, it is electromagnetic waves that carry the signals (Fig. 1).

This particular point is sometimes misunderstood; it is not electrons or other charge carriers that carry the signals in wires, rather it is electromagnetic waves (of one kind or another). To see this point rather simply, we only need to note that signals in wires do not propagate at the electron velocity ($\sim 10^6$ m/s). In typical electrical cables, the signals move essentially at the velocity of light (or somewhat smaller if the cables are filled with a dielectric). (We refer to such cables as "LC" lines below.) Hence it is not in general correct to say that signals propagate faster in optics because they move at the speed of light - a point that is relatively obvious physically, but often misstated, especially in popular accounts. In fact, signals typically travel slightly slower in optical fibers than they do in coaxial cables because the dielectric used in cables has a lower dielectric constant than glass.

In the case of electrical interconnection lines on chips, the signals do move at a slower rate, but here the rate is determined by the overall resistance and capacitance of the interconnect line (a dissipative wave propagation), a rate that is only indirectly related to the electron velocity. The practical bit rate of communications on chip is about $1/RC$ [6] for simple digital circuits, where R is the line resistance and C is the line capacitance. (We will call such lines "RC" lines below.) This gives a propagation time that is proportional to the square of the line length, and hence cannot be characterized by a simple velocity. The effective velocity is, however, often much slower than the velocity of light (e.g., $\sim 10^7$ m/s for long lines on a chip).

It is very rare in electronics that signals are carried directly by the electrons themselves, except within devices (with a good example being the cathode ray tube). Of course, the electrons in the wires do perform a useful function, which is that the resulting electrical currents are crucial in guiding the wave in the wires, a point to which we will return below. Hence in general we can conclude that in both the optical and electrical cases, information is carried by electromagnetic waves. We could equally well say that the information is carried by photons in both cases since we can always describe an electromagnetic wave quantum mechanically in terms of photons. (To be somewhat pedantic, the description in terms of photons is more correct since it can cover the nonclassical states such as photon number states and squeezed states that can also be used to carry information but that cannot be adequately described classically.)

The point of the above discussion is to clarify that the difference between electrical and optical interconnections is not that one sends information on electrons and the other on photons - both use photons.

Optical logic is not the subject of this paper, but it is also worth pointing out that all ways of performing optical logic work through the use of electrons (or, conceivably, charged nucleons, as in phonon scattering of light). This is true of nonlinear optics in its various forms, and of optoelectronics. Obviously, all

electronic logic devices use electrons to effect the necessary switching. So it is also practically true that all logic interactions, optical or electronic, are based on electrons.

Because of the fact that both optical and electrical interconnections use photons (and all logical interactions in both use electrons), the common statement that the difference between electronics and optics is that electrons are fermions and photons are bosons is arguably at best irrelevant and at worst misleading in understanding the real differences between optics and electronics and the benefits of either.

3. Electrical and Optical Interconnection Physics - Differences

The real and important basic differences between optical and electrical physics for the purposes of interconnecting electronics can be summed up as in Fig. 2. The three differences, the shorter wavelength, the higher (carrier) frequency, and the larger photon energy, are all aspects of the same difference, since choice of any one of wavelength, frequency, or photon energy uniquely determines the other two. (Wavelength $\lambda = c/v$, where c is the velocity of light, and v is the frequency; photon energy (in the convenient energy units of electron-volts) $= hv / e$ where h is Planck's constant and e is the electronic charge.) For the electrical case, we have shown numbers for photon energies and electromagnetic wavelengths corresponding to a typical practical range of clock frequencies for electronic digital systems, 10 MHz to 10 GHz. Arguably most of the potential advantages of optics come from one or other of these fundamental areas of difference, and this will be a useful way to categorize the differences and their consequences. In some cases, it is somewhat arbitrary whether consequences are ascribed to the high frequency or to the short wavelength, though the consequences of the large photon energy are clearly distinct.




Fig. 2. Fundamental differences between optics and electronics for communication, expressed in terms of wavelength, frequency and photon energy.

3.1. High frequency of light

In electrical interconnections, we generally work at "base band", i.e., we typically do not use a "carrier", but simply turn the voltage on and off. Converting the information so that it modulates a high-frequency (e.g., microwave) electrical carrier is usually sufficiently cumbersome that it is impractical for interconnection (though can be used for telecommunications). In optics, on the other hand, we are generally modulating a very high frequency carrier, and with modern optoelectronic devices and integration techniques, the modulation of the optical carrier and its detection are becoming increasingly practical. The high frequency of light has several consequences.

3.1.1. Absence of signal loss and distortion at high modulation frequencies

The carrier frequency of light is very high compared to any frequency at which we can modulate. As a result, modulating the light beam makes essentially no difference to the propagation of light. Only over large distances in fibers do we see dispersive effects resulting from high speed modulation, and optics has negligible additional propagation loss from large bandwidth signals. Hence there is no frequency-dependent loss or frequency-dependent crosstalk when using light as the information carrier (whether in fibers or free-space) within machines.

By contrast, electrical interconnections have very substantial problems of signal distortion at high modulation frequencies [7][8][9][10]. This is true whether we consider "RC" lines or "LC" lines. The distortion is a consequence of the signal loss becoming higher at higher frequencies. In RC lines, this is essentially a filtering effect from the RC circuit that corresponds to the line. In the LC case, it is a consequence of the skin effect, the phenomenon by which conduction takes place only in an increasingly thin layer (~ 2 microns at 1 GHz) near the surface of a conductor at higher frequencies. (Somewhat surprisingly, there are only really two kinds of lines - RC lines dominated by the bulk resistance of the wire, and LC lines dominated by the skin effect resistance, because the cross-over from RC to LC behavior occurs at essentially the same frequency as that at which the skin effect becomes important [7]). The distortion tends to lead to "inter-symbol interference" (ISI) - a carry-over of the effects of previous bits of information that makes impossible the simple discrimination of the current bit being transmitted. The problems of loss and distortion in electrical lines lead to several difficulties in system design, including the following two specific consequences.

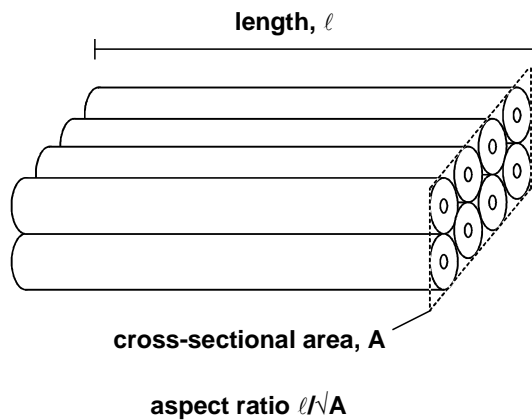


Fig. 3. Illustration of the "aspect ratio" of a set of electrical cables.

3.1.1 (a) Difficulty of "high aspect ratio" architectures

It has recently been realized [7][8][9][10] that there is a relatively general formula for characterizing the number of bits per second, B , that can be sent down simple electrical interconnects given the ISI from frequency-dependent loss and distortion. This limit is set only by the ratio of the length ℓ to the cross-sectional dimension \sqrt{A} of the interconnect wiring — the

"aspect ratio" of the interconnection (see Fig. 3). (A is the total cross-sectional area of the wiring). The limit is approximately $B \sim B_o A/\ell^2$ bits/s, with $B_o \sim 10^{15}$ (bit/s) for high performance strip lines and cables, $\sim 10^{16}$ for small on-chip lines, and $\sim 10^{17} - 10^{18}$ for equalized lines. Such a limit will certainly become a problem as machines approach Tb/s information bandwidths, and is already easily seen as a practical limit on sending signals down long cables, for example. This limit is scale-invariant - that is, neither growing or shrinking the size of the system substantially changes the number of bits that can be sent - because it only depends on the ratio of length and cross-sectional size. It also makes no difference to the capacity of the interconnect whether we use many small wires running at low bit rates per wire, or a few large wires running faster - the total number of bits per second possible down the interconnect depends only on the length and the total cross-sectional area of the wiring.

The aspect ratio limit is also largely independent of the details of the design of the electrical lines. This is partly because all well-designed lines, from lines on chips through to large coaxial cables, have approximately the same capacitance per unit length ($\sim 1 - 3$ pF/cm), inductance per unit length (~ 2 nH/cm), and (if they are LC lines) characteristic impedance ($\sim 30 - 100$ ohms). These numbers follow from fundamental constants (i.e., the permittivity of free space is 8.85 pF/m, the permeability of free space is 1.26 μ H/m, and the impedance of free space is 377 ohms), and from the fact that capacitance, inductance, and impedance scale essentially logarithmically with the ratio of conductor size to conductor separation in an electrical line, hence changing little with the design of the line. Attempts to make low capacitance lines or high impedance lines would involve making at least one of the conductors in the line very small, which would exponentially increase the loss as the capacitance is reduced or impedance increased. Making large capacitance (and low inductance) lines is possible (and useful for applications such as power supply lines), but the impedance of the line becomes very low, leading to high power dissipation if used as a signal line. The practical optimum signal line design always ends up having the conductor separation comparable to the (smaller) conductor cross-sectional size. These properties of electrical lines have other consequences, to which we will return below.

This "aspect ratio" limit can be exceeded in electrical interconnects through the use of techniques such as repeater, coding, and multilevel modulation; it is worth noting, too, that equalization, repeater, coding, and multilevel modulation could be implemented using electronic circuits on chips. In fact, digital equalization techniques for interconnections are being researched now [11] precisely because of this difficulty. Also, it should be emphasized that the "aspect ratio" limits are not the fundamental, Shannon limit on interconnect capacity, a limit that can be approached using the techniques mentioned above. The point here, however, is that there is a simply characterized boundary at which the concept of a simple digital electrical "wire" breaks down - that is, the wire can no longer be regarded as reliably transmitting the logic voltage from a logic gate at one end into the logic voltage into a logic gate at the other end. Pushing past this point requires increasing use of equalization,

repeating, coding, and multilevel modulation techniques, at increasing cost in terms of delay, circuit size and complexity, and power dissipation.

Optical interconnections, however, simply do not have this "aspect ratio" problem at all. First, as mentioned above, they do not have modulation-frequency-dependent loss, e.g., changing the modulation frequency of a signal on a light beam from 1 MHz to 1 GHz makes no difference to the loss experienced by the signal. Second, the loss in optical fibers can be extremely low in absolute magnitude, e.g., 0.2 dB/km in fiber used for long distance communications, leading to negligible distance-dependent loss over the scale of interconnect distances. Third, pulse dispersion, though it does exist in optical fiber, is relatively weak compared to metallic cable, and can be compensated anyway. In typical (uncompensated) long-distance fiber, there is essentially zero pulse dispersion at about 1.3 microns wavelength, and at 1.5 microns wavelength (where the loss is minimum), the dispersion is about 15 ps for every nanometer of modulation wavelength bandwidth (corresponding to about 130 GHz of frequency bandwidth) and every kilometer of length. This means, for example, less than 1/10 of a clock period of dispersion for a 6 GHz bandwidth signal over 1 km fiber length. Fourth, optical fiber can be very small in diameter (e.g., 125 microns). As a result, the optical interconnections can readily exceed the bit rate capacity of simple electrical interconnects by at least 9 orders of magnitude for the same cross-sectional area and length [7].




Fig. 4. Illustration of one kind of high-aspect-ratio architecture, in which one processor needs to communicate directly with many others, hence requiring long, thin interconnections (i.e., high aspect ratio interconnects).

High "aspect ratio" interconnects are likely to be encountered in situations where one processor or unit must communicate directly with many others, since then at least some of the processors must be physically relatively distant (Fig. 4). Such architectures are likely to be encountered, for example, in telecommunications switching and multiprocessor information

processing machines, and possibly in other situations such as processor-to-memory communications. With increasing clock rates, numbers of processors, and sizes of memory, the "aspect ratio" problem becomes increasingly severe for electrically interconnected machines, creating a strong motivation for exploring optical alternatives.

3.1.1 (b) Signal and clock timing

A signal propagating down an electrical line may start out with sharply rising and falling "edges", but these will gradually lengthen from the loss-related distortion discussed above. This "softening" of the edges makes precise extraction of timing information progressively more difficult. This can be a significant problem, for example, when trying to communicate the system clock signal. Loss and distortion have no effect on the clock frequency, though with strong attenuation some circuitry may be required to extract the frequency reliably. Reliably extracting the clock phase can, however, be much more difficult. Obviously, if the clock receiving circuit triggers at some specific voltage, then varying loss in the line will lead to a variation in the time at which the circuit triggers, especially if the arriving clock signal has a relatively slow rise and fall. Even if we compensate for the loss (by, for example, arranging to trigger when the clock signal rises above 50% of its peak received value), any variation in the rise or fall times of the clock signal will change the trigger points and the detected clock phase. One source of variability in both loss and signal rise and fall time is the temperature dependence of the resistivity of the metals used as conductors in electrical wiring. For both copper and aluminum, for example, the resistance of a line changes at a fractional rate of about 0.004 / °C, leading to a ~ 40% change over a 100 °C range. The delay on an RC line and the rise time on an LC line are both simply proportional to the resistivity in the line, and so would both change by ~ 40 % over the 100 °C range, leading to substantial potential variation in detected clock phase.

By contrast, optical systems have relatively little problem with such variations. As discussed above, the dispersion and loss in optical fiber are small, so there is little distance-dependent signal distortion or loss. There is some variation in the propagation speed of signals with temperature in optical fiber because of the change of refractive index with temperature, which is ~ -10^{-5} / °C. For a 10 m optical fiber cable, the corresponding change in delay over a 100 °C temperature range is only about 30 ps (about 0.07% of the propagation delay).

3.1.2. Absence of frequency dependent cross-talk at high frequencies

Electrical wires become increasingly good antennas at high frequencies, both for transmitting and receiving. This is true whether we consider true electromagnetic transmission and reception, or simply capacitive coupling between lines (as is readily encountered on chips, for example [6]). Again, because of the high carrier frequency of optics, it essentially makes no difference to any cross-talk in optics what the modulation frequency is, so there is essentially no frequency-dependent cross-talk with optics, a significant feature for high-speed, dense interconnections.

3.1.3. Impedance matching using resonant transformers

There are several differences between optics and electronics as far as impedance matching is concerned. We will return to consider several others below when we discuss the benefits of short wavelength and large photon energy.

One particular feature of impedance matching with optics results from the fact that optical signals have very small modulation bandwidth compared to the optical carrier frequency; the impedance matching necessary in optics, for example as a light beam has to transition from propagating in air to propagating in a semiconductor or in glass, can be handled relatively effectively using a very simple resonant impedance transformer. The "resonant impedance transformer" in the optical case is an anti-reflection coating. The classic simple anti-reflection coating is a dielectric layer a quarter wavelength thick, with a refractive index that is the geometric mean of the indices being matched. In principle, being a resonant device, it is only a perfect impedance transformer for exactly the right carrier frequency (or wavelength). Because the modulation bandwidth is so small compared to the carrier frequency in the optical case, however, the modulation has no practical effect on the functioning of the antireflection coating.

Simple resonant impedance transformers work well also in electrical systems with small modulation bandwidth on a carrier frequency, with an example being a stub resonator used to match a transmission line to an antenna. But such a simple passive electrical impedance transformer will not work well with broadband modulation, and certainly does not work with "base-band" (i.e., no carrier) modulation that is normal in simple digital electrical interconnections; unencoded digital logic signals may go down nearly to d.c. in their frequency spectrum because they may have arbitrarily long strings of "zeros" or of "ones". Even non-resonant impedance transformers, such as the common inductive transformer, will not work down to d.c., and for that reason among others, are not commonly used in logic systems. We will return to such impedance transformers when we discuss the consequences of large photon energy.

3.1.4. Use of short optical pulses

It is relatively straightforward in optics to generate short optical pulses. The technique of "mode-locking" of lasers can give a repetitive stream of pulses, with pulse lengths in the range of ~ 100 ps to ~ 10 fs, and pulse repetition rates of ~ 100 MHz to ~100 GHz. Such pulse sources are now routine in the laboratory, and are likely to become increasingly practical for other applications because of advances in the high power laser diodes that can be used to drive the short pulse laser operation. Generation of short pulses by mode-locking is possible because of the high frequency of light.

The possibility of using short optical pulses creates some novel opportunities, even when the electronic devices in the rest of the system operate at speeds much longer than the pulse length. One use would be in clock distribution (a concept for which optics is interesting even without the use of short pulse lasers [12]). A central, mode-locked laser could serve as the system clock [13].

The short pulses arriving at clock receivers would give the best possible clock signal to the clock receiver; rather than a slowly rising clock "edge", the clock receiver would be driven essentially by an impulse function, giving the best possible definition of the clock phase. Based on numbers discussed above, for example, it would appear to be possible to distribute clock signals with less than 30 ps variation over a 100 °C temperature range in a system of the order of 10 m in size using optical fiber. (Such a system would require cutting the length of the fibers to within millimeters of accuracy, which should present no basic problem.)

In systems in which an optical modulator (rather than a light emitter) is used as the optical output device driving the interconnect, it is possible to take advantage of short pulses for the signal transmission as well. Though the electrical signal driving the modulator may have relatively slow rising and falling edges, and the modulator may also only respond on such relatively slow time scales, still the modulator can modulate a short optical pulse from a mode-locked optical power supply beam. As a result, the signal is now carried as the amplitude of a short optical pulse.

The use of output modulators with short optical pulses has two benefits. The first benefit is in the performance of the interconnect link [14][15][16][17][18]. At the receiver end, the receiver is driven by an impulse, which will generally give much better performance out of the receiver than if it is driven by the usual slowly rising and falling signals [18]. For a given amount of optical energy, the impulse drive will generally give the largest output voltage swing, for example, and the fastest rising and falling edges out of the receiver [18]. At the transmitter end, we need only drive optical power through the output modulator when it has completed its transition to its desired output state [17]. This gives the most efficient use of optical power since no optical power is wasted driving the modulator while it is still transitioning from one state to another. Some modulators (e.g., quantum well electroabsorptive modulators [19]) dissipate electrical power based on the absorbed optical power, and hence the use of short optical pulses would minimize that power dissipation also.

The second benefit is that the use of short pulses with output modulators can eliminate signal skew. If all of the output signals are read out based on the same short pulse optical clock source, they can all be read out synchronously. Furthermore, the use of the short pulse effectively "samples" the outputs; as long as the outputs are all valid at the time of the arrival of the short pulse, any underlying variation in the actual rise and fall times, or the precise timing, of the drive signals to the modulators is irrelevant, and is effectively removed by the "sampling" operation. This additional feature of short pulse operation raises the possibility of running entire large systems (e.g., 10 m in size) entirely synchronously, with the same clock and signal phase throughout.

Another possibility with short pulse systems is the use of ultrafast devices for time-multiplexing an interconnect for higher capacity. This is currently not yet practical for interconnects, but devices operating on picosecond or faster time scales are

feasible in the laboratory [20][21], and do represent a longer term possibility for the use of optics.

3.1.5. Wavelength-division multiplexing

The very high carrier frequency of light also allows the use of multiple different frequency carriers on the same light beam or in the same optical fiber. (In the terminology of optics, it is more common to refer to the carriers as being on different wavelengths rather than on different frequencies, hence the term "wavelength-division multiplexing".) There is no problem in principle with the use of multiple wavelengths in optics, and laboratory techniques exist for combining and separating them. At the time of writing, various techniques are being developed that may allow practical use of this concept in applications, and the technique is in use in long-distance communications systems. Such wavelength-division multiplexing could increase the capacity of the interconnection system or reduce the amount of cabling required in the system. For example, it could allow interconnection between two-dimensional arrays of devices using only one-dimensional arrays of optical fibers; one-dimensional fiber arrays are currently much easier to align and connectorize.

One issue with the use of wavelength-division multiplexing is constructing arrays of sources with the correct set of output wavelengths. This is, however, likely a solvable problem, either through carefully graded or controlled arrays of devices [22], or by "slicing" a broad spectrum source (such as a short pulse laser [23]).

3.2. Short wavelength of light

The short wavelength of light leads to the following consequences.

3.2.1. Low-loss dielectric waveguides and optical components

In electrical interconnections, the wavelength corresponding to the frequency of the signal is generally large compared to the cross-sectional size of the wiring that must route the signals within the system. In optics, however, because the wavelength of light is so small, the structures that guide the optical waves can be made larger in cross-sectional dimensions than the optical wavelength (e.g., a 10 micron diameter core in a single mode optical fiber is much larger than the approximately 1 micron wavelength of light).

In general, waves are confined and guided using boundaries between materials. (Usually the boundaries are abrupt, though sometimes they are smoothly graded.) At or near the boundary with the guiding material, the guiding material responds sufficiently strongly to the incident wave amplitude to reflect the wave in the desired direction. (Waveguiding can be regarded as successive reflections off the sides of the waveguide.) If we require abrupt changes in wave direction or we try to confine the wave in dimensions comparable to or even smaller than its wavelength, a large response is required from the guiding material to satisfy the very abrupt changes (per wavelength) required in the spatial derivatives of the waves.

For electromagnetic waves, the response of the guiding material is essentially a current (sometimes expressed in terms of a polarization). When a wave is incident on a dielectric material, small oscillatory currents can flow, essentially without loss, as temporary, small distortions (or polarizations) of the electron clouds in the material. Such effects are strong enough to confine waves in dielectric waveguides (such as optical fiber) that are large compared to a wavelength. As a result, we can have extremely low loss propagation of signals in optical fiber, and we can also make low-loss lenses and other optical components that route optical signals in free space. But for the base-band electrical case, where waves must be confined and directed over dimensions small compared to the wavelength, only conducting materials can in practice provide enough current response for the guiding. Conventional conducting materials are lossy, leading to high loss in electrical lines. The loss is usually frequency-dependent because of the skin effect, and so also leads to pulse dispersion. Even superconductors, though technically loss-less conductors at d.c., can have significant loss when carrying high frequency signals because the inductive voltage (which unavoidably appears across the line when currents are changing) leads to conventional, lossy conduction as a parasitic process.

3.2.2. "Free-space" multi-channel imaging interconnects

In electrical systems it is usually unthinkable not to control carefully the information path from source to destination using a waveguide. Certain exceptions exist, where we may make a few wireless connections through free space, but for interconnections at any significant density such free space electrical interconnections are impractical. One reason for this impracticality is fundamentally that the wavelengths of the electrically-driven signals are too long. In fact, the wavelengths are typically larger than a circuit board, and certainly larger than a chip in most situations. The laws of diffraction tell us that it is difficult to focus a wave to a dimension smaller than a wavelength. Hence, we could not focus two interconnecting "beams" to different points on a chip or board, allowing us only one interconnection.

Various other reasons make such "free-space" electrical interconnects impractical. For example, it is difficult to design antennas that have a broad enough bandwidth to operate with base band modulation (including operating down to d.c.). Free space electrical interconnections could also be very sensitive to cross-talk and to picking up extraneous signals. It is possible in principle to use carrier frequencies to alleviate some of the problems, but modulating and demodulating carriers, as mentioned above, is relatively complex, and may involve significant delay, power dissipation, and circuit area.

In optics, by contrast, it is routine simultaneously to image multiple sources on one plane to multiple receivers on another; this is the function performed by a simple lens in a camera or in the human eye. (See Fig. 5.) The fundamental reason that makes this possible is the short wavelength of light; even with relatively simple optics, it is possible to image thousands of outputs on one surface to thousands of inputs on another, with spot sizes on the order of several (e.g., 10) wavelengths in size.




Fig. 5. Concept of imaging arrays of outputs on one plane to arrays of inputs on another.

Despite the fact that all of the beams may overlap at the lens, the lens separates out all of the information again to image the multiple beams at the receiving "inputs". Optics therefore allows very large numbers of connections from one plane to another through "free space". An incidental benefit is that all of the connections in such imaging-based systems are automatically very well synchronized with one another, hence avoiding another practical problem of "signal skew" common in electrical systems. Another related consequence is that it is possible to make very global interconnect topologies (such as so-called "perfect shuffles" [24], crossover networks [25], banyans [26], and sliding banyans [27]) in which many of the "beams" cross through one another. Such patterns can be useful for operations such as telecommunications switching fabrics and fast Fourier transforms.

The concept of letting the information flow almost unguided through free space is generally rather foreign to those used to the electrical world. It does have its practical difficulties. There is, for example, the issue of generating and aligning multiple light beams to relatively fine mechanical tolerances. This is not, however, a fundamental problem, and working digital optical systems with more than 60,000 light beams have been demonstrated in the laboratory [28]. It is also true that, with such free space optics, aligning a regular array of many light beams is not much more difficult than aligning one light beam, so there is a substantial economy of scale in working with thousands of light beams in an array. There is undoubtedly technological work to be done here to make such optics practical and inexpensive, but there is a substantial base of existing and novel optical technologies that can be exploited. For a collection of work on optical systems, see, e.g., Ref. [29].

It is also worth noting that free space interconnections need not actually be in open space; they could take place essentially entirely within solid, rigid glass structures, for example. Objections that free space optics would fundamentally be subject to atmospheric dust and turbulence are therefore not fundamental, and any such practical difficulties are likely relatively easy to overcome.

Electrical interconnections do have the strong practical feature that it is relatively straightforward to construct essentially arbitrary interconnection patterns. In free-space optics, on the other hand, such arbitrary patterns are not so readily implemented. Techniques such as volume holograms do allow such patterns in principle, though in practice the efficiency drops off when large numbers of different interconnections are implemented; in practice, such volume holographic interconnects are much more suited to some interconnect topologies than others [30]. We should note, however, that optics also has no fundamental problem with making arbitrary interconnect patterns between two surfaces. As an existence proof, we can imagine imaging each output point in a surface array onto its own fiber in a similar array of optical fibers. The fibers can then be rearranged any way we wish before collecting them again into another regular array that can be imaged onto an array of input points on another array, hence implementing an arbitrary permutation of outputs to inputs. The practical problem here is that there is no known simple way to mass-produce such a complex interconnection pattern in optics. Simple patterns can be mass-produced, for example, using diffractive optical elements that can be stamped out on a surface, but such approaches run out of "degrees of freedom" if we try to implement truly arbitrary patterns [31]. One relative strength of optics is that highly regular global patterns (such as the perfect shuffle), which must have many relatively long and crossing "lines", are by contrast usually easy with free-space optics. In free space, light beams can pass through one another with no interference of any kind.

3.2.3. Beamsplitters without back reflection

It is very often desirable to be able to make multiple connections to a given signal line so that the same data can be made available to multiple parts of a system. This situation is encountered routinely on backplanes, for example, where we may wish to plug in a number of different circuit cards. Once we are working at clock frequencies sufficiently high that the wavelength associated with the clock frequency is comparable to or smaller than the size of the system (e.g., the length of the backplane), we cannot, however, neglect wave reflections. Any simple attempt to plug in additional connections leads to wave reflections from the connection. Often this is describable as an impedance matching problem; two 50 ohm lines connected in parallel appear to have an impedance of 25 ohms, so connecting two 50 ohm output lines to one 50 ohm input line leads to an impedance mismatch. The only way that all of the boundary conditions can be met is for a wave to reflect back up the input line.

In free space optics, however, it is straightforward to use a beamsplitter to split out any fraction of the beam without any back reflections (though, of course, the power transmitted through the beam splitter is reduced accordingly). One reason why this works is that the physical processes that divide the beam (usually partial wave reflection off of an interface) effectively divide both the electric and magnetic components of the wave by the same factor, retaining the correct impedance ratio between them; no back reflected wave is required to satisfy boundary conditions. By contrast, at the simple junction of several conducting lines, the voltage is constrained to be

identical on all of the lines because the conductors are joined, leading to the necessity of the backward wave to satisfy the voltage and current boundary conditions.

In principle, this kind of perfect beamsplitting is also possible in certain kinds of waveguides, both optical and microwave. In practice in waveguides there is the additional subtlety that the modes must be matched as well as the overall impedances. (Strictly, different modes are liable to have slightly different impedances as well, though the mode matching problem goes beyond these slight impedance mismatches.) Mode matching means that the spatial form of the electric (and magnetic) field distribution should not be disturbed at any junction, which is obviously not easy to achieve in general (though free-space beamsplitters achieve this naturally by replicating the input beam form and only changing its angle of propagation).

3.2.4. Fan-in

A problem that is shared by both electrical and optical connections is the difficulty of combining independent signals without fundamental loss. Though there are some differences between optical and electrical methods, which we discuss here, this is an area in which they actually behave similarly despite apparent differences in the physics.

Many schemes that we might imagine would allow us to fan in multiple signals to the same input port without loss, do not work, and if they did they would allow us to violate the second law of thermodynamics. The underlying problem is that, if we were able to combine without loss two uncorrelated signals into one physical channel (or "mode") we would be able to devise schemes that would allow us to heat up one warm body with two cooler ones, violating the second law. This point is discussed by Goodman [32] and Goodman and Lam [33]. In both the electrical and optical cases, it is difficult to combine N channels to one input without sustaining a $1/N$ power loss. In electrical systems, it is therefore usual to perform the fan-in logically rather than physically, sending each input channel to a separate logical gate input, and rather than trying to perform a simple "wired-OR" function physically without logic gates.

The presence of the same $1/N$ loss has sometimes not been understood in optics, with a presumption that optics has effectively arbitrarily large "fan-in". It is indeed possible to combine N inputs without power loss into one photodetector, for example by bringing the N beams at N distinct angles or onto N distinct positions on the photodetector. The catch is that the photodetector then needs to be N times larger in area than it would have to be for only one beam. This is obvious in the case of bringing different inputs onto different parts of the photodetector. In the case of different angles, diffraction theory tells us that if we want to have N complete beams incident at different angles on a surface, without having any of the angles overlap with one another, we need to make the surface N times as large also because narrowly converging beams focus to larger spots.

The larger detector area required for "loss-less" beam combination tends to reduce the electrical response of the detector proportionately also. For example, the larger area of

detector means that its capacitance, C , has likely increased by the same factor N . As a result, to get a particular frequency response ($\sim 1/RC$) from the detector and its electrical load circuit, the load impedance, R , on the photodetector has to be reduced by a factor N . Consequently, the voltage induced on the photodetector for a given input beam power (and hence photodetector current) will also be reduced by a factor N . In this case we would have done just as well by taking the $1/N$ beam combination loss and working with a smaller detector. Hence, at least in this example case, there is little or no benefit in terms of useful power transfer efficiency by trying to use a larger photodetector with "loss-less" beam combination.

It is, or course, quite legal to combine different wavelengths without loss since the different frequencies represent different physical modes, and use of such wavelength techniques is a possible advantage of optics.

3.3. Large photon energy of light

The most important single physical consequence of the large photon energy in optics is that, for essentially all of the situations of importance here, light is generated and detected quantum mechanically, whereas electrical signals use classical generation and detection. (The main exceptions to the quantum mechanical generation and detection of light are thermal generation of light from hot bodies, as in tungsten light bulbs, and thermal detection of light energies, neither of which is important for optical interconnection.) For example, detection of light in practice involves counting photons, not measuring electric field amplitudes. A typical semiconductor photodiode will generate one electron of current for every absorbed photon, with the electron of current resulting from the quantum-mechanical absorption of a photon to create an electron-hole pair. Similarly, a relatively efficient laser diode will generate one photon for every few electrons passed through the diode, with each photon resulting from the recombination of an electron and a hole. In contrast, electrical signals are carried on voltages or currents, which effectively are respectively the statistical average potential energies and flow rates of classical "gasses" of electrons. Changes of these average potentials or flow rates at the receiving device cause changes of average potentials or flow rates in "gasses" of electrons inside the receiving device. Describing the electric field in an electrical line in terms of photons, though correct, is of little use, there being no simple relation between arrival of the very low energy photons and the response of any particular electron.

The fact that the photon energy of light is so large has two specific consequences for optical interconnections, one already extensively exploited, and another that has only relatively recently been understood.

3.3.1. Voltage isolation

Detecting photons allows us to generate currents and voltages without any direct electrical connection with the light source, yet still with a band width that extends down to d.c. as required for logical interconnections. The photons carry no information about the d.c. levels in the source circuit. In other words, optical interconnects give us perfect electrical isolation between two

circuits, and make the absolute voltage levels in the different circuits irrelevant. This already solves an important problem in electrical systems, and is exploited extensively in so-called "opto-isolators", which usually contain a light-emitting diode (connected to the "transmitting" circuit) and a photodiode (connected to the "receiving" circuit) (Fig. 6). All optical interconnects automatically have this voltage isolation property.

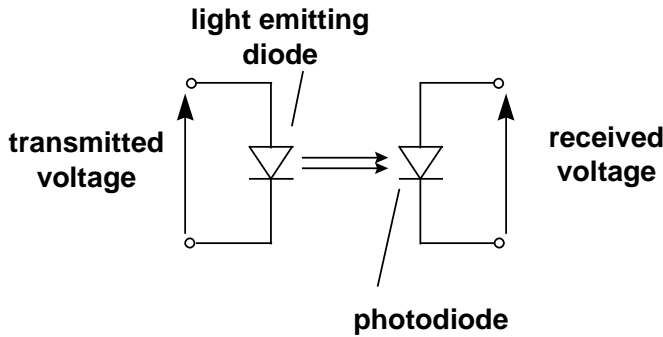


Fig. 6. Opto-isolator circuit. The optical interconnection between the source (usually a light-emitting diode in an opto-isolator) and the photodetector provides signal connection down to d.c. without any connection between the voltage levels in the two circuits. This isolation follows naturally from the quantum nature of the optical detection, and is a property of all optical interconnects.

3.3.2. Quantum impedance conversion

As discussed above, essentially all electrical signal lines have both high capacitance per unit length (1 - 3 pF/cm), and low impedance (~ 30 - 100 ohms). This creates a problem for electronic circuits, illustrated in Fig. 7. Though individual electrical logic devices may be small, low-power-dissipation, high-impedance devices, the interconnection medium between them is unavoidably low impedance and/or high capacitance per unit length.

The classic solution to this problem might seem to be to use impedance transformers. At low (e.g., audio) frequencies, inductive impedance transformers work reasonably well, though even then they do not operate to d.c. At high frequencies, they typically have significant loss and parasitic capacitance effects, and are also not readily integrated. The result is that the typical solution in logic systems to match, for example, to low impedance transmission lines, is to use active impedance matching in the form of low-impedance line drivers, and terminating resistors. The resulting power dissipation (e.g., 10's to 100's of mW per high speed electrical line) and use of chip area for drivers is a significant issue in chip design when high data throughput is required.

Use of optical emitters or modulators and photodiodes fundamentally enables us to avoid the problems of the low impedance of electrical transmission lines [34], as illustrated in Fig. 8. The reason why optics can avoid the low impedance problem is that the voltage generated in a photodetector bears no particular relation to the classical "voltage" in the light beam. It is quite possible, for example, to generate 1 V in a photodetector

from a light beam with 600 microvolts of classical voltage [34] - a consequence, fundamentally, of the photoelectric effect.




Fig. 7. A small, high impedance, low power electronic device that wishes to communicate to another similar device, for example on another chip, is forced to use an electrical line with low-impedance and/or high capacitance per unit length. Line drivers with low output impedance and high power dissipation are the typical solution.




Fig. 8. Optical devices can effectively match the impedances between the electronic logic devices because they use quantum detection and generation. The photodiode generates one electron of current for every absorbed photon. Electroabsorptive modulators tend to pass one electron of current for every photon modulated. Efficient laser diodes would also emit one photon for every electron of current in principle.

This particular benefit has not historically been exploited in practice because the optical sources have usually consumed more power than a 50 ohm resistor anyway, so there was no power saving, and the driver problem was, if anything, worse than that of an electrical line. The emergence of quantum well modulator technology has, however, led to quite practical low-power optical output devices that can demonstrably send digital signals from chip to chip with substantially less power (e.g., < 6 mW total dissipation at 375 Mb/s) than electrical connections [35]. Other approaches, such as low threshold lasers, may be able to offer similar benefits. This feature of optics is likely to be particularly important for large arrays of optical inputs and outputs, and may allow much larger amounts of information to be sent on and off chips optically than is practical electrically. One recent study [36] has predicted (implicitly using the benefits of quantum impedance conversion) that dense optical interconnections directly in and out of silicon chips will have an interconnect capacity that will be able to track the ability of advancing silicon technology to perform logic operations, and to achieve information flows exceeding 1 - 10 Tb/s on and off a single silicon chip.

Fig. 9 shows a comparison of chip layout designs for electrical pad drivers and receivers and optical input and output driver circuits implemented in the hybrid SEED technology [37][38], one of the so-called "smart pixel" technologies [39] that allow very dense optical interconnections to electronics. Because of the quantum impedance conversion, the driver circuit for the optical case is small. It can also operate at much higher speeds than the electrical input and output circuits and pads shown in the figure. Note also that, because of the electrical isolation provided by the optical input, there is no need for the input electrostatic discharge protection circuitry, which helps make the optical input circuitry small. Because of the small size and low power dissipation of the electronic circuitry required for optical inputs and outputs, optics has the potential to provide very high densities of interconnections; more than 4000 optical I/O have already been demonstrated bonded to a silicon chip [37].

One commonly heard criticism of optical interconnections is that they have high latency (total time taken to send and/or receive a signal). Sources of latency in some data links include the multiplexing and buffering circuitry used; such latency results from transmission format conversion, which can be avoided in interconnection over shorter distances. In optical links intended for very long distances, and hence designed to work with the minimum possible received optical power, there may also be some latency in the receiver circuitry. There is, however, no fundamental source of latency in conversion between optics and electronics. In the case of the circuitry illustrated in Fig. 9, in which optical diodes for inputs and outputs are bonded directly to the electronic circuits, there may even be less latency in the optical case because there are fewer stages of output driver circuitry required. It is also sometimes argued that conversion between optics and electronics is inefficient; such inefficiency is a consequence of system design, not of fundamentals. If power efficiency is important (which it generally was not for long distance communications), systems such as those in Fig. 9 can be designed with low overall dissipation (e.g., < 6 mW as mentioned above).




Fig. 9. Comparison of circuit layouts for optical and electrical inputs and outputs. The electrical input and output pads and associated circuitry are a standard cell design. The optical input and output circuits are for differential drive of quantum well input photodiode and output modulator diode pairs.

The reduction in power required for output drivers in the optical case may also reduce some of the high current swings required from the electrical power supply on the chip, which may in turn help avoid problems of pin inductance on chips. (Rapid changes in current cause inductive voltage drops over the power supply (and possibly also signal) pins on the chips, which can become very significant and can lead to the use of multiple power and ground pins on chips.)

4. Conclusions

Optical and electrical physics have much in common, with both communicating through electromagnetic waves and both using electrons for essentially all logical interactions. They differ in the higher frequency, shorter wavelength, and larger photon energy of light compared to electrical signals.

The higher frequency and shorter wavelength allow us to use optical fibers or free space to send optical signals, and thereby to avoid the major loss phenomena that limit the capacity of electrical interconnects because of signal and clock distortion and attenuation. Optics avoids issues of frequency-dependent cross-talk, and most of the problems of impedance matching and wave reflections encountered in electrical systems. Because of the large photon energy of optics, optical generation and detection for interconnection is quantum mechanical (e.g., counting photons) as opposed to the classical sourcing and detection of voltages and currents in the electrical case. Two practical consequences are that all optical interconnections provide voltage isolation (as already exploited in opto-isolators), and optics can offer lower powers for interconnects because it can solve the problem of matching high-impedance low-power devices to the unavoidably low impedance (and/or higher capacitance) of electromagnetic propagation. As a result of the basic differences in physics, optics arguably offers solutions to most of the physical problems now being experienced in electrical interconnections.

Optics also offers several additional opportunities that have essentially no practical analog in the electrical case, including use of short pulses for improved interconnect performance and system synchronization, free space imaging interconnects of thousands of interconnects directly on and off chips, and wavelength division multiplexed interconnections. One general conclusion we can draw is that there are many potentially useful features of optics for interconnection that are not currently exploited.

We could ask what it is that will drive us to use optics more for interconnection. The answer to that question has aspects that go well beyond the physics of optical and electrical interconnections. It is, however, clear from the physics that optics becomes increasingly attractive at high bit rates and higher interconnect densities (e.g., high density edge connectors for boards, or even very high density connections on and off of chips), and arguments for optics become increasingly strong as the interconnect length becomes longer. The aspect ratio limit means existing architectures cannot be simply scaled to higher clock rates and larger numbers of processing units using electrical interconnections, a problem that will become increasingly severe as systems approach interconnection rates of

~ 1 Tb/s; this is an area where optics can completely eliminate the underlying problem.

One benefit of optics could be to simplify design of interconnects, avoiding the growing difficulty of designing high performance electrical interconnection systems. For example, problems of impedance matching, wave propagation effects, pin inductance, electromagnetic pick-up and cross-talk, electrostatic discharge, and voltage isolation on electrical lines are eliminated, and the designs of optical systems for interconnection remain the same independent of bit rate.

At the system level, optics offers opportunities that are not currently available with electrical interconnections. For example, given the small dispersion and weak temperature dependence of delay in optical fibers, and the possibility of short pulse optical power sources for clocking and re-timed signal readout, it would appear to be quite feasible to run a large system (e.g., of the order of 10 m in size) entirely synchronously, with delays being an exact number of clock cycles. Whether such a concept is important can only be answered at the system level, but it would at least eliminate synchronization circuits and buffers, and might allow simpler communications protocols.

Another physical opportunity offered by optics to system design is the ability to avoid making a hierarchical interconnect structure. In electrical interconnects, it is often necessary to go through a hierarchy of interconnections of increasing cross-section (e.g., device to chip, chip to board, board to backplane, backplane to cable) to allow communication between different parts of the system. This hierarchy often leads to bottlenecks in the interconnection system (e.g., in busses). With optics, a connection directly out of a chip can go an arbitrarily long distance (at least on the scale of a processing machine) with no further amplification, regeneration, or increase in cable diameter. Hence, we can contemplate directly connecting chips in different cabinets if there is any advantage to the system. Whether these example system-level opportunities are important remains to be seen, but they do indicate radically different architectures may be feasible through the use of optics.

It is arguably the case that existing system architectures are largely optimized around the strengths and weaknesses of electrical interconnections. As a result, the introduction of optics as a substitute for electrical lines may at first seem to offer only marginal benefit since the electrical interconnect by definition is already performing the task adequately. The real benefits of optics are likely to emerge once the systems evolve to exploit optical features. It is also worth noting that in interconnections, optics is competing with copper and aluminum, not silicon. Though silicon technology advances rapidly, the relatively mature technology of electrical lines is not likely to evolve to keep up with the silicon, and indeed the rapid improvement of silicon is likely to create more demand for optical interconnects to allow the interconnect to keep pace with the logic and memory.

In conclusion, the various features of optics for interconnections discussed above clearly address many of the problems that are being encountered more and more in making electronic information processing systems. These features come directly

from the differences in the physics between optics and electronics. The fact that the beneficial features of optics are so firmly rooted in the physics gives us reason to believe that optics will eventually be used extensively in interconnection in electronic systems.

Acknowledgments

I am pleased to acknowledge many useful and stimulating discussions with Tony Lentine, Ashok Krishnamoorthy, Haldun Ozaktas and Ted Woodward, and extensive discussions and collaborations with Wayne Knox on use of short optical pulses.

References

- [1] F. A. P. Tooley, "Challenges in Optically Interconnecting Electronics", *IEEE J. Sel. Top. Quantum Electron.* **2**, 3 - 13 (1996)
- [2] D. A. B. Miller, "Dense Optical Interconnections to Silicon Electronics", *Trends in Optics*, ed. A. Consortini (Academic Press, San Diego, 1996), pp. 207 - 222
- [3] J. W. Goodman, F. J. Leonberger, S.-Y. Kung, and R. A. Athale, "Optical Interconnections for VLSI", *Proc. IEEE*, **72**, 850 - 866 (1984)
- [4] M. R. Feldman, S. C. Esener, C. C. Guest, and S. H. Lee, "Comparison between Optical and Electrical Interconnect based on Power and Speed Considerations", *Appl. Optics*, **27**, 1742 - 1751 (1988)
- [5] H. M. Ozaktas, "Towards and Optimal Foundation Architecture for Optoelectronic Computing", *Proc. Third Int. Conf. Massively Parallel Processing Using Optical Interconnections*, Maui, 1996, (IEEE Computer Society, 1996), pp. 8 - 15
- [6] H. B. Bakoglu, "Circuits, interconnections, and packaging for VLSI" (Addison-Wesley, New York, 1990).
- [7] D. A. B. Miller and H. M. Ozaktas, "Limit to the Bit Rate Capacity of Electrical Interconnects from the Aspect Ratio of the System Architecture", *J. Parallel Dist. Comp.* (to be published).
- [8] H. M. Ozaktas, *A Physical Approach to Communication Limits in Computation*, Ph.D. Thesis, Stanford University, Stanford, CA (1991).
- [9] H. Ozaktas and J. W. Goodman, "The limitations of interconnections in providing communication between an array of points", *Frontiers in Computing Systems Research*, Vol. 2, 61-130, Ed: S. K. Tewksbury (Plenum Press, New York, 1991).
- [10] B. J. Smith, "Interconnection networks for shared memory parallel computers", *Proc. 2nd Int. Conf. Massively Parallel Processing Using Optical Interconnections (MPPOI'95)*, 255-256 (IEEE Computer Society Press, Los Alamitos, California).
- [11] W. J. Dally and J. Poulton, "Transmitter Equalization for 4 Gb/s Signalling", *Proc. of Hot Interconnects IV Conference*, August 15-17, 1996, Stanford, pp.29-39
- [12] J. W. Goodman, F. J. Leonberger, S.-Y. Kung, and R. A. Athale, "Optical Interconnections for VLSI Systems", *Proc. IEEE* **72**, 850 - 866 (1984)
- [13] P. J. Delfyett, D. H. Hartman, and S. Z. Ahmad, "Optical Clock Distribution using a Mode-locked Semiconductor Laser Diode System", *J. Lightwave Tech.* **9**, 1646 - 1649 (1991)
- [14] G. D. Boyd, G. Livescu, L. M. F. Chirovsky, and A. L. Lentine, "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).
- [15] G. Livescu, G. D. Boyd, L. M. F. Chirovsky, R. A. Morgan, and T. Mullally, "Picosecond switching of GaAs/AlGaAs smart pixels" *Optics Lett.* **19**, 2122-2124 (1994).
- [16] 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).
- [17] A. L. Lentine, L. M. F. Chirovsky, and T. K. Woodward, "Optical energy considerations for diode-clamped smart pixel optical receivers", *IEEE J. Quantum Electron.* **30**, 1167-1171 (1994).
- [18] L. Boivin, M. C. Nuss, J. Shah, D. A. B. Miller, and H. A. Haus, "Receiver Sensitivity Improvement by Impulsive Coding" (submitted to *Optics Letters*)
- [19] D. A. B. Miller, "Quantum Well Optoelectronic Switching Devices" *International J. of High Speed Electronics*, **1**, 19-46, (1990)

- [20] N. S. Patel, K. L. Hall, and K. A. Rauschenbach, "40-Gbit/s cascaded all-optical logic with an ultrafast nonlinear interferometer", *Optics Lett.* **21**, 1466-8 (1996)
- [21] K. H. Ahn, M. Vaziri, B. C. Barnett, G. R. Williams, X. D. Cao, M. N. Islam, B. Malo, K. O. Hill, and D. Q. Chowdhury, "Experimental demonstration of a low-latency fiber soliton logic gate", *J. Lightwave Tech.* **14**, 1768-75 (1996)
- [22] W. Yuen, G. S. Li, C. J. Chang-Hasnain, "Multiple-wavelength vertical-cavity surface-emitting laser arrays with a record wavelength span", *IEEE Photonics Tech. Lett.* **8**, 4-6 (1996)
- [23] E. A. De Souza, M. C. Nuss, W. H. Knox, and D. A. B. Miller, "Wavelength-division multiplexing with femtosecond pulses," *Optics Letters*, **20**, 1166-1168 (1995)
- [24] K. H. Brenner and A. Huang, "Optical implementation of the perfect shuffle interconnection," *Appl. Opt.* **27**, 135-137 (1988)
- [25] J. Jahns and M. J. Murdocca, "Crossover networks and their optical implementation," *Appl. Opt.* **27**, 3155-3164 (1988)
- [26] T. J. Cloonan and M. J. Herron, "Optical implementation and performance of one-dimensional and two-dimensional trimmed inverse augmented data manipulator networks for multiprocessor computer-systems," *Opt. Eng.* **28**, 305-314 (1989)
- [27] R. R. Michael, M. P. Christensen, and M. W. Haney, "Experimental evaluation of the 3-D optical shuffle interconnection module of the sliding banyan architecture", *J. Lightwave Tech.* **14**, 1970-1978 (1996)
- [28] H. S. Hinton, T. J. Cloonan, F. B. McCormick, A. L. Lentine, and F. A. P. Tooley, "Free-space Digital Optical Systems", *Proc. IEEE* **82**, 1632 - 1649 (1994)
- [29] J. Jahns and S. H. Lee (eds.), *Optical Computing Hardware* (Academic, New York, 1994)
- [30] D. Psaltis, D. Brady, X.-G. Gu, and S. Lin, "Holography in artificial neural networks," *Nature*, **343** (#6526), 325-330 (1990)
- [31] H. M. Ozaktas, K.-H. Brenner, and A. W. Lohmann, "Interpretation of the space-bandwidth product as the entropy of distinct connection patterns in multifacet optical interconnection architectures", *J. Opt. Soc. Am. A* **10**, 418-422 (1993)
- [32] J. W. Goodman, "Fan-in and fan-out with optical interconnections", *Optica Acta*, 32 1489 - 1496 (1985)
- [33] J. W. Goodman and J. C. Lam, "Fan-in loss for electrical and optical interconnections", *Trends in Optics*, ed. A. Consortini (Academic Press, San Diego, 1996) pp. 223 - 236
- [34] 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).
- [35] A. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, T. K. Woodward, J. E. Ford, G. F. Aplin, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. D. Kossives, D. Dahringer, L. M. F. Chirovsky, and D. A. B. Miller, "3-D Integration of MQW Modulators Over Active Submicron CMOS Circuits: 375 Mb/s Transimpedance Receiver-Transmitter Circuit", *IEEE Photonics Tech. Lett.* **7**, 1288-90 (1995)
- [36] 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**, 55-76 (1996).
- [37] A. L. Lentine, K. W. Goossen, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. J. Tseng, R. E. Leibenguth, J. E. Cunningham, W. Y. Jan, J.-M. Kuo, D. W. Dahringer, D. P. Kossives, D. D. Bacon, G. Livescu, R. L. Morrison, R. A. Novotny, and D. B. Buchholz, "High-Speed Optoelectronic VLSI Switching Chip >4000 Optical I/O Based on Flip-Chip Bonding of MQW Modulators and Detectors to Silicon CMOS", *IEEE J. Sel. Top. Quantum Electron.* **2**, 77 - 84 (1996)
- [38] T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, and L. M. F. Chirovsky, "Optical Receivers for Optoelectronic VLSI", *IEEE J. Sel. Top. Quantum Electron.* **2**, 106 - 116 (1996)
- [39] H. S. Hinton, "Progress in Smart Pixel Technologies", *IEEE J. Sel. Top. Quantum Electron.* **2**, 14 - 23 (1996)