JOURNAL OF PARALLEL AND DISTRIBUTED COMPUTING **41,** 42–52 (1997)
ARTICLE NO. PC961285

Limit to the Bit-Rate Capacity of Electrical Interconnects from the Aspect Ratio of the System Architecture

D. A. B. Miller* and H. M. Ozaktas†

*Ginzton Laboratory, Stanford University, Stanford, California 94305-4085; and
†Bilkent University, Department of Electrical Engineering, TR-06533 Bilkent, Ankara, Turkey

We show that there is a limit to the total number of bits per second, B, of information that can flow in a simple digital electrical interconnection that is set only by the ratio of the length l of the interconnection to the total cross-sectional dimension $\sqrt{A}$ of the interconnect wiring—the "aspect ratio" of the interconnection. This limit is largely independent of the details of the design of the electrical lines. The limit is approximately $B \sim B_0 A/l^2$ bits/s, with $B_0 \sim 10^{15}$ (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. Because the limit is scale-invariant, neither growing nor shrinking the system substantially changes the limit. Exceeding this limit requires techniques such as repeatering, coding, and multilevel modulation. Such a limit will become a problem as machines approach Tb/s information bandwidths. The limit will particularly affect architectures in which one processor must talk reasonably directly with many others. We argue that optical interconnects can solve this problem since they avoid the resistive loss physics that gives this limit.

1. INTRODUCTION

Electrical interconnections currently dominate the communication of information within digital machines; they are often the most inexpensive approach, and are known to be reliable. At long distances and high bit rates, optical techniques are, however, increasingly being used, in large part because they do not suffer from the signal attenuation and distortion that arise from the skin effect loss in normal metal cables [27]. The skin effect is the physical phenomenon [9] that, at high frequencies, current is carried only in a thin layer near the surface of the conductors (for example, the "skin depth" in copper at 1 GHz is about 2 μm). In practical cables, the skin effect is unavoidable [1] [3]. There are also several other reasons for the use of optics in interconnections [17, 18, 21]. Optics dominates long-distance telecommunications, has growing use in some networks and connections to peripheral devices, and is being introduced for frame-to-frame interconnects and for backplanes.

It is obvious that longer lines are more lossy and have more pulse distortion. A key point of the present paper is that it is the "aspect ratio"—the ratio of the length of the (longest) interconnection line to the cross-sectional dimension of the interconnect—rather than the length that better characterizes limits from loss and distortion, especially for simple digital interconnections in which signals are to be sent directly as "on" or "off" voltages without coding or multilevel modulation schemes. The resulting limits are independent of size scale, being set instead by the proportions of the interconnections, and, ultimately by the proportions of the underlying system architecture; neither growing nor shrinking the size of the system substantially changes the limit (though in practice, lines on chip exhibit slightly better performance than other lines). As we will show below, for example, the bit-rate capacity, B, of a broad range of electrical cables and lines is given by $B \sim 10^{15} A/l^2$ bits/s. Hence we can conclude that certain architectures become substantially more difficult or expensive to implement with electrical interconnections once we wish to pass more than a certain amount of information through them per unit time.

By contrast, optical interconnections have no comparable scaling limits. Loss in optical media is essentially independent of the modulation bit rate, even into the terahertz regime. Optical attenuation can be relatively very small; it is negligible in "free-space" propagation over scales of meters, and in optical fibers over kilometers (e.g., 0.2 dB/km loss at 1.55 micrometer wavelength). Optical systems readily exceed the electrical scaling limit derived here by multiple orders of magnitude. For example, a single-mode optical fiber, 125 micrometer in diameter and 15 km long (corresponding to $A/l^2 \sim 5 \times 10^{-17}$), would have only 3 dB of distance-dependent loss. Sending signals at, for example, 100 Mb/s is simple and routine in such a system. A simple electrical interconnect with the same length and cross section would, according to the scaling analysis of this paper, be able to carry less than 0.1 bit/s of data. More sophisticated optical systems can handle much greater distances and bit rates, with, for example, 160 Gb/s transmitted over 232 km of optical fiber

[1] In principle, the skin effect could be avoided up to certain frequencies by making cables with a very large number of thin layers, but no practical fabrication technique has emerged. See Clogston [3].

0743-7315/97 $25.00

[25]. Though optical interconnection systems doubtless have their own practical problems, it is clear that they do not suffer from the same kind of scaling limits as electrical interconnects. Consequently, we will concentrate in the rest of this paper on the analysis of electrical interconnects, and on the knowledge that optics can solve the resulting problems of interconnect bit-rate capacity, at least at some size scales.

This aspect ratio limit to electrical bit-rate capacity is relatively universal. Because capacity depends on total cross-sectional area, replacing a large number of thin cables with a small number of thick cables (or *vice versa*) has no effect on the limit. The detailed design of the cable cross-section is also largely irrelevant because of an underlying logarithmic scaling in the physics of cable properties; as a result, all well-designed cables and lines give similar limits (though bad design could certainly make the problem worse). Somewhat surprisingly, we also get essentially similar limits regardless of whether the loss results from skin-effect or conventional "bulk" resistive effects.

Obviously, a single long wire has a very large aspect ratio, and as a result has a low bit-rate capacity (this is equally true for a long coaxial cable or a long wire on a chip). The wiring layers in boards and chips have relatively small cross-sectional areas, and so can also encounter this limit. In systems with large numbers of processors, each of which wishes to be connected to every other reasonably directly, large aspect ratio interconnects are inevitable, since at least some of the processors must necessarily be many "processor sizes" away from some of the other processors. Switching fabrics that must connect a large number of input sources to a large number of output systems (as in telecommunications switching or possibly interconnection switching for large multiprocessor parallel computing machines) have similar problems.

Aspect ratios of interconnects inside machines will often be lower than those of long cables between machines, but they can readily be in the range of 10:1 to 100:1. For such aspect ratios, we will likely start to run into the "aspect ratio limit" at total bit-rate capacities in the range of 10s of Gb/s to 1 Tb/s. Though large by current standards, such machines may not be very far off. For example, an experimental telecommunications switching fabric with 320 Gb/s capacity has been demonstrated using state-of-the-art electronic packaging [26], and it is interesting to note that the architecture chosen for this demonstration (an active crossbar, i.e., a simple mesh interconnect) is one that avoids high-aspect-ratio interconnects, though its cost would increase quadratically with its capacity. Other known multistage switching architectures scale more nearly linearly,[2] but have inherently higher aspect ratios. Optics in various forms has been suggested for use in high-capacity switching systems [4, 26]. Experimental optically interconnected multistage systems have been demonstrated [16]. The use of optics for very-high-aspect-ratio interconnections within a switching system is being researched now; in one such experimental system, a single silicon chip forms the core switching fabric interconnecting a large number of high-capacity switching machines [14]. In this system, optical connections are made directly from the surface of the silicon chip to the other switching machines, avoiding entirely the usual hierarchy of repeatered electrical interconnections. Such radical approaches are possible because of emerging technology that allows large numbers of high-speed optical interconnections directly in and out of silicon chips [7].

The argument of this paper is an extension of more general arguments previously given by one of us [21, 22], and the central point has also been briefly discussed recently by Smith [24] (and also orally by one of us [18] at the same meeting). In this paper, we give the physical arguments explicitly and in greater detail, including also the case of equalized lines, and extend the discussion of practical electrical interconnects to show specific realistic numbers for limits.

In Section 2, we present theoretically this "aspect ratio" limit on the bit-rate capacity of both skin-effect-limited lines and conventional resistive lines. We discuss the transition between skin-effect-limited behavior and conventional resistive behavior in Section 3; importantly, we show here that there really are only two modes of operation of lines—either skin-effect-limited "propagating-wave," "LC" lines, or bulk-resistance-limited, "diffusive-conduction" "RC" lines. This conclusion makes the results simpler and more universal. In Section 4, we discuss how to scale the theoretical limits from ideal lines to the case of more practical lines, and give explicit scaling results for coaxial lines, strip lines in multichip modules, and lines on chips. In Section 5, we extend the analysis to the case of equalized lines. We discuss the results in Section 6, and draw conclusions in Section 7.

2. THEORETICAL BIT-RATE CAPACITIES OF ELECTRICAL LINES

Though there are many different designs of electrical transmission lines, in practice there is little variation in the capacitance per unit length, C_l, and the inductance per unit length, L_l, of different lines. Most lines have about 1–3 pf/cm of capacitance and about 2 nH/cm of inductance. The reason for these relative constancies is that both capacitance and inductance scale only logarithmically with the ratio of conductor size to conductor separation. For example, for coaxial lines, the capacitance is

$$C_l = \frac{2\pi\varepsilon_r\varepsilon_o}{\ln(r_b/r_a)}, \tag{1}$$

where ε_r is the relative dielectric constant of the material between the conductors, ε_o is the permittivity of free space, r_b is the inner radius of the outer conductor of the coaxial cable, and r_a is the radius of the inner conductor. The inductance is

$$L_l = \frac{\mu_r\mu_o}{2\pi}\ln(r_b/r_a), \tag{2}$$

[2]Strictly, multistage architectures tend to scale as $N \log N$, where N is the total capacity.

where μ_r is the relative permeability of the material between the conductors, and μ_o is the permeability of free space. For all practical low-loss cable, $\mu_r = 1$.

For any line (coaxial or of other design) operated at a frequency where the inductive impedance per unit length substantially exceeds the resistive impedance per unit length (i.e., LC lines), the characteristic impedance, Z_o, is

$$Z_o = \sqrt{\frac{L_l}{C_l}}. \tag{3}$$

For a coaxial line, Z_o becomes

$$Z_o = \frac{1}{2\pi}\sqrt{\frac{\mu_r\mu_o}{\varepsilon_r\varepsilon_o}}\ln\left(\frac{r_b}{r_a}\right). \tag{4}$$

For polyethylene dielectric (commonly used in coaxial cable), ε_r = 2.33 at 100 MHz, which would give r_b/r_a = 3.57 for Z_o = 50-Ω cable. Air- or vacuum-spaced 50-Ω coaxial cable would have r_b/r_a = 2.30 (expanded PTFE with $\varepsilon_r \sim 1.3$ is also commonly used).

The signal propagation velocity, v, is, for any LC cable

$$v = \frac{1}{\sqrt{L_l C_l}}, \tag{5}$$

which in practice for coaxial cables becomes $c/\sqrt{\varepsilon_r}$, where c (=3 × 10^8 m/s) is the velocity of light in free space.

Lines other than coaxial lines (e.g., twisted pair, conductor above ground plane (narrow) strip lines) all have behavior similar to coaxial lines as the relative dimensions are changed because the logarithmic behavior comes from the underlying physics. Because of the logarithmic scaling, redesigning for substantially lower capacitance or higher impedance results in at least one of the conductors becoming very small relative to the overall area, which leads to high resistance and, hence, high loss. Lines, such as very wide strips close together on top of one another, have higher capacitance and lower inductance; such lines are useful for power distribution, but their low impedance or high capacitance leads to high required driver powers [3] if used to send signals. Increasing the dielectric constant also increases line drive power and slows down signal propagation. Increasing the permeability, in principle, could increase line impedance, but is generally not practical and would reduce propagation velocity.

Hence, to minimize capacitance (and maximize impedance in LC lines) without incurring too much loss or occupying too much area, practical lines designed for signal transmission all tend to have at least one conductor dimension comparable to the separation between the conductors, and to have rather similar inductance, capacitance, and impedance (for LC lines); $\ln(r_b/r_a)$ (or the equivalent factor for other designs of line) tends to have a value of ∼1. Other specific analyses support this conclusion [15, 21, 22]. Broadly speaking, only the resistive loss changes substantially between different cable or line designs, and it tends to scale with overall area. This relative universality of transmission line properties allows us to draw rather general conclusions about the bit-rate capacity of electrical signal transmission lines. As we will show below, the effect of resistive loss is to give a bit-rate capacity for a given line that is proportional to the cross-sectional area and inversely proportional to the square of the cable length, with a proportionality constant that is essentially similar for all practical signal transmission line designs.

2.1. Skin-Effect-Limited Lines

The skin effect is a well-known consequence of Maxwell's equations for the case of wave propagation in the presence of conductors. All of the conduction tends to occur within about a "skin depth," δ, of the surface of the conductor. δ is given by

$$\delta = \frac{1}{\sqrt{\pi f \mu_r \mu_o \sigma}}, \tag{6}$$

where f is the frequency and σ is the conductivity. At room temperature, the conductivity of copper is about 5.80 × 10^7 Ω^{-1} m^{-1} and that of aluminum is about 3.78 × $10^7\Omega^{-1}$ m^{-1}, which leads to skin depths in the range of micrometers at the clock frequencies of electronic processors.

When skin-effect-limited resistance loss dominates (e.g., neglecting radiation loss and dielectric loss), the response, $h(\tau)$, at the (properly terminated) [4] end of an LC line to a unit step function input is [27]

$$h(\tau) = \operatorname{erfc}(\sqrt{1/\tau}), \tag{7}$$

where "erfc" is the complementary error function, and τ is time measured in normalized units β, where (for a coaxial line) [5]

$$\beta = \frac{\mu_r\mu_o l^2}{16\pi^2 r_a^2 \sigma Z_o}, \tag{8}$$

with l being the length of the line ($\tau = 0$ corresponds to the time at which the step function edge would have arrived at the end of the line in the lossless case).

[3] Optics can also eliminate the problem of driving the unavoidably low impedance of even well designed lines. See Miller [19].

[4] When we consider LC lines in this paper, we always assume that they are terminated. If lines are not terminated, then there are additional problems with reflections on the lines that in general only make the eye diagrams worse. Also, such reflections do not, in general, speed up the rise time on the line. Hence maximum practical bit-rate capacities will be found for terminated lines. The termination could be either "series" (i.e., at the transmission end) or "parallel" (i.e., at the receiving end), or both. The formulas are given for parallel termination; series termination may have a different magnitude (by a factor of 2 depending on definitions), but has the same form and rise time.

[5] Here we assume, for simplicity, that the loss in the coaxial line is dominated by the loss from the smaller, inner conductor. The actual loss will be slightly higher if the loss from the outer conductor is also included.

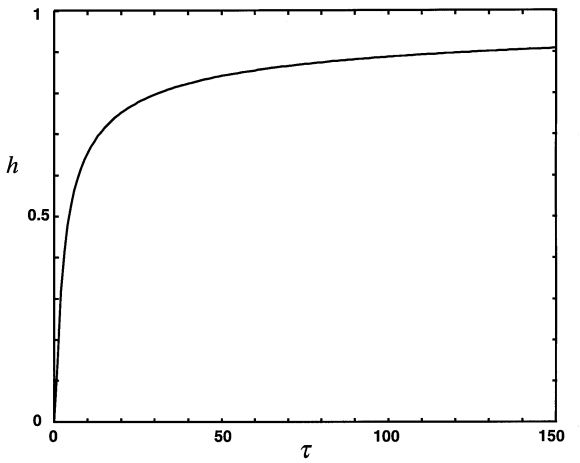

FIG. 1. Rise of a voltage at the end of a terminated, skin-effect-limited transmission line when driven by a unit step function (in dimensionless units of time).

The form of the function $h(\tau)$ is shown in Fig. 1. Note that, though the initial rise of $h(\tau)$ is rapid (in a few units of time), the function takes a very long time to rise after that. The 10–90% rise time is actually about 120 units of time. This long overall rise time significantly limits the number of bits per second that can be sent down the line unless sophisticated coding and signal recovery techniques are used.[6]

Resulting "eye diagrams" are shown in Fig. 2. In an eye diagram, a random bit stream drives the system, and the results from each bit period are plotted on top of one another, showing the unwanted influence of previous bits ("pattern-dependent" effects) as a "closing" of the "eye."

The long total rise time of the skin-effect-limited line leads to substantial "pattern-dependent" effects, as can be seen in Fig. 2. For a bit period of 5 time units, making the digital decision as to whether a given bit is a "zero" or a "one" is not quite impossible (the eye is still "open"), but a relatively sophisticated "decision" circuit would be required to do so, and the immunity to noise is clearly significantly degraded. To meet a standard digital design criterion of a bit period significantly larger than the 10–90% rise time would mean a bit period of 100s of time units. Here we presume that, with care in design, we can work with an effective "rise time," $\tau_{\rm rise}$, of 50 time units, and we also presume that we can set the bit period equal to this $\tau_{\rm rise}$ and get sufficiently reliable digital communication with relatively simple decision circuits. With this presumption, therefore, the bit-rate capacity, B, of a given cable would be

$$B < 1/\tau_{\rm rise} = 1/50\beta. \tag{9}$$

The practical meaning of the unit of time, β, is more easily understood by rewriting it. Using the fact that the dc resistance per unit length of a coaxial line is

$$R_l = \frac{1}{\pi r_a^2 \sigma} \tag{10}$$

(where we neglect the slight additional resistance from a larger, outer conductor), we find

$$\beta = \frac{1}{4\log(r_b/r_a)}\, R_l C_l l^2 = \frac{1}{4\log(r_b/r_a)}\, RC, \tag{11}$$

where R and C are the total dc resistance and capacitance of the line, respectively. Hence the effective rise time, $\tau_{\rm rise}$, of the skin-effect-limited coaxial line is [from Eqs. (9) and (11)]

$$\tau_{\rm rise} \cong \frac{12.5}{\ln(r_b/r_a)}\, RC. \tag{12}$$

This formula will be useful for comparison with the behavior of the RC line below (it is, though, an algebraic accident that the unit of time here appears to depend on the dc resistance, R, of the line, and it has no basic physical significance).

Another way of rewriting the unit of time is

$$\beta = \frac{\mu}{16\pi Z_o^2 \sigma} \left(\frac{r_{bo}}{r_a}\right)^2 \frac{l^2}{A}, \tag{13}$$

where A is the cross-sectional area of the line, and r_{bo} is the outer radius of the outer conductor (which may include both the thickness of the outer conductor and that of any insulating

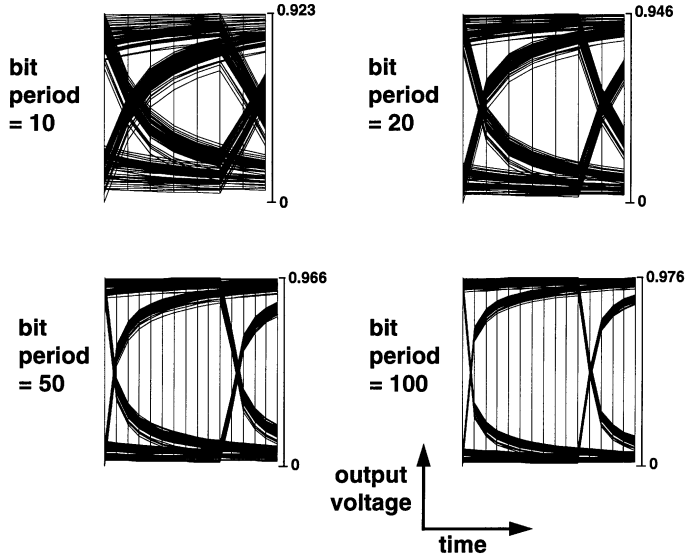

FIG. 2. Eye diagrams for bit periods corresponding to 10, 20, 50, and 100 normalized time units, calculated for a random 256-bit pattern.

[6]The long tail on the rise time is a consequence of the $1/\sqrt{f}$ frequency response of the skin effect. The more common exponential time response (e.g., the rise of a simple resistance–capacitance circuit) corresponds to a $1/f$ frequency response. Though the $1/\sqrt{f}$ has less attenuation at higher frequencies than the $1/f$ response (corresponding to a relatively rapid initial rise in the time response for the $1/\sqrt{f}$ case), it has more attenuation at lower frequencies, hence leading to the relatively longer tail in the time response compared to an exponential rise.

or protective sleeving). Based on this formulation, and taking as a limit $r_{bo} = r_b$ (the inner radius of the outer conductor), we obtain for such an "ideal" copper coaxial cable with polyethylene dielectric and 50-Ω impedance

$$B < 9.1 \times 10^{15}\, A/l^2 \tag{14}$$

(this relation is plotted as the LC line limit in Fig. 4).

Note that this bit-rate capacity depends essentially only on the aspect ratio $l/\sqrt{A}$. Note also that we could, instead of using one large cable of cross-sectional area A, use several small cables of the same *total* cross-sectional area A, and obtain the same *total* bit-rate capacity B. Each of the individual small cables would be run at a proportionately lower bit rate because the rise time would be proportionately larger, but the aggregate capacity would be the same. Hence A really refers to the total cross-sectional area. Changing the cable design makes essentially no difference to the overall capacity as long as the total cross-sectional area remains the same.

2.2. Bulk Resistive Lines

In the case of lines of dimensions such that the resistive impedance dominates over the inductive impedance at the operating frequency range of interest, the 0–90% rise time at the end of the (unterminated) line happens to be numerically equal to the total RC time constant of the line.[7] We choose this time as the effective rise time, τ_{rise}, of such lines, giving

$$\tau_{\text{rise}} \cong RC = R_l C_l l^2. \tag{15}$$

R will be the dc resistance of the line (though there may be a transitional frequency region where it starts to be affected by the skin effect—we discuss this transition below).

This formula can be compared directly with Eq. (12) above for the coaxial line case. We see immediately that the bit-rate capacity of the RC line is actually larger (for the same geometry) than that of the LC line by approximately an order of magnitude.

The dc resistance decreases with the area of the line. As we will discuss below, for high-performance systems, the RC line behavior will usually only hold in practice on chip. The geometry of a well-designed on-chip line will have an area approximately 8 times the area of the (smaller) conductor[8] [1]. Again, there is relatively little that can be done to improve such a design. For example, reducing the resistance by increasing the fraction of the area occupied by the smaller conductor will generally increase capacitance to give little or no net benefit. The resistance of the smaller conductor is therefore

$$R_l \cong \frac{8}{\sigma A}, \tag{16}$$

and so we have, for the bit-rate capacity of the RC lines,

$$\begin{aligned} B &\cong \frac{1}{8 R_l C_l}\,\frac{A}{l^2} = 2.4 \times 10^{16}\,\frac{A}{l^2}\,(\text{copper}) \\ &= 1.6 \times 10^{16}\,\frac{A}{l^2}\,(\text{aluminum}). \end{aligned} \tag{17}$$

Again, we find that it does not matter whether we use a large number of small RC lines or a small number of large RC lines—the *total* bit-rate capacity is the same, and depends only on the *total* area A.

3. TRANSITION BETWEEN BULK RESISTIVE RC LINES AND SKIN-EFFECT-LIMITED LC TRANSMISSION LINES

We have analyzed only the above two cases: (i) LC lines with skin-effect loss and signal distortion, and (ii) RC lines, where the resistance is not dominated by the skin effect; i.e., the resistance is that of the whole cross-sectional area of the conductor. For the generality of the analysis, it is important to understand the other two possibilities, i.e., LC lines with resistive loss from the whole conductor cross section or RC lines with skin-effect loss. In fact, as we demonstrate below, there actually are essentially only the two cases we have already analyzed; the frequency at which a line of a given size changes from bulk resistive behavior to skin-effect behavior is also essentially the frequency at which the line changes from

[7] Note that the rise time of a distributed RC line is faster than that of a single resistance and capacitance of the same total values; this can be rationalized by noting that the capacitance near the transmitter end is charged through a relatively small resistance, for example. Only the capacitance at the very far end of the line experiences the total line resistance. The general form of the response of a distributed RC line is sometimes known as "diffusive conduction," since the line obeys an equation identical in form to the diffusion equation. There is, unfortunately, no known analytic form for the solution to the diffusion equation in the time domain for the kind of boundary and initial conditions implied by a step function voltage drive, though simulation is straightforward.

[8] Total line capacitance for on-chip lines minimizes when the ratio of conductor width, W, to the separation between the conductor and the ground plane (dielectric height, H) is ~ 1.75, i.e., $W \sim 1.75H$ (see Bakoglu [1], p. 140). The conductor is taken to have a square cross section. The separation between adjacent conductors in a collection of lines is taken as $2W$ to achieve reasonably low cross talk, and to avoid too much coupling capacitance between the lines, giving a center-to-center spacing of $3W$. For a single conductor above a single ground plane, the height is not simply $H + W$ because the electric field of the line extends well above the conductor. To estimate the effective total height of the line, we presume that we place another ground plane a distance H above the conductor to make a strip line. Such a structure can be stacked to make a multilayer wiring structure if desired. The line therefore has an overall thickness of at least $2H + W$. We should also include the thickness of one ground plane conductor in the effective height, and we take this thickness to be at least $W/2$. Hence the effective area of one line is at least $3W(2H + W + W/2) \sim 8W^2$. The capacitance of the simple line without the upper ground plane is about 2 pF/cm for the silicon oxide or silicon nitride dielectrics used in present integrated circuits [1]. Adding the upper ground plane will increase the capacitance, so we presume a line capacitance of about 3 pF/cm.

RC to LC behavior (essentially the same conclusion is reached in Ref. [22] by a different, but equivalent, argument).

The frequency, f_{RL}, for the transition from RC to LC behavior for a given line is that at which the inductive impedance equals the resistive impedance, i.e., using Eqs. (2) and (10) and the fact that the inductive impedance is $2\pi fL$,

$$f_{\mathrm{RL}} = \frac{1}{\pi r_a^2 \sigma \mu_r \mu_o \ln(r_b/r_a)}, \tag{18}$$

where, for the sake of definiteness, we consider the coaxial line geometry.

The frequency, f_{SE}, at which the skin effect becomes important is when the skin depth is comparable to the conductor radius (or half the larger cross-sectional dimension of the smaller conductor in the line), i.e., for a coaxial line, using Eq. (6),

$$f_{\mathrm{SE}} = \frac{1}{\pi r_a^2 \sigma \mu_r \mu_o}. \tag{19}$$

It is clear that these two frequencies, f_{RL} and f_{SE}, are essentially identical; they differ only by the factor $\ln(r_b/r_a)$, which we have argued above is approximately unity for practical signal lines. If we consider lines other than the coaxial line, the logarithmic factor $\ln(r_b/r_a)$ may be replaced by a slightly different factor, but, for well-designed signal lines, as discussed above, this will make little difference. If we have lines in which the smaller conductor is substantially larger in one cross-sectional dimension than in the other, the formulas will be slightly different again, but well-designed lines tend to have relatively square or circular cross sections since otherwise they acquire excess capacitance or low impedance. Hence, Eqs. (18) and (19) are essentially correct for all well-designed signal lines. Taking the conductor size $d = 2r_a$ as the effective cross-sectional dimension of the smaller conductor of the line, we have, for the "cut-over" frequency, f_{C}, between bulk RC behavior and LC skin-effect behavior,

$$f_{\mathrm{C}} = \frac{4}{\pi d^2 \sigma \mu_r \mu_o}. \tag{20}$$

This relation is graphed in Fig. 3.

Note from Fig. 3 that, at the frequencies of 100 MHz and above that might be found in high-performance interconnections, only the smallest lines (e.g., <10 μm) are RC lines, and such small lines are usually found only on chips (rather than on boards or in cables).

4. SCALING THE PERFORMANCE OF PRACTICAL CABLES AND LINES

The LC line limit given by Eq. (14) is obviously idealized as far as cable cross-sectional area is concerned, since it presumes no outer conductor or sheathing thickness. We can scale to practical cables by noting the actual cross-sectional area of the real cable of the same inner conductor size as the "ideal" cable of Eq. (14); the "ideal" cable (which, by assumption, has a 50-Ω geometry and a polyethylene dielectric) has an area of $\pi(3.57/2)^2 = 10.0$ square units for an inner conductor one unit in diameter. The bit-rate capacity of the real cable is derated in proportion to its actual area compared to this "ideal" area.

Practical lines also typically have more loss than the ideal cable of the same inner conductor size. There are several possible reasons for this difference. One is surface roughness on the conductors. As discussed above, the skin depth is often of the order of micrometers; if the surface is rough on this scale or larger, the current actually must flow over a longer path than if the surface had been smooth, resulting in greater loss. Lines in geometries other than coaxial (e.g., strip lines) will generally have more current flowing on some parts of the conductor surface than others (usually those parts of the conductor nearer to the "return path" on the other conductor), which results in more skin-effect loss for the same total signal current. There are also other (non-skin-effect) losses, such as dielectric loss, anomalous skin effect [2], and radiation loss, though we will neglect these here.

The effect of additional "skin-effect-related" loss reduces the bit-rate capacity of a real line essentially as the square of that additional loss factor, as can be seen from the following arguments. Suppose that the loss, α_{p} (in, for example, dB per unit length), of a practical cable is greater than that, α_{i}, of the ideal cable by some factor η at some frequency f_a. The loss from the skin effect scales as $f^{1/2}$. Hence, to reduce the loss to the same as that of the ideal cable, we should reduce the operating frequency (and hence the bit-rate capacity) by a factor η^2. Equivalently, we can keep the frequency the same

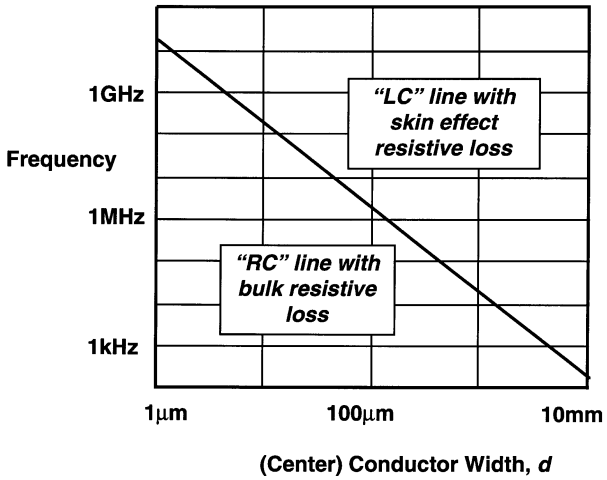

FIG. 3. Graph showing transition frequency between bulk resistive-capacitive (RC) diffusive conduction behavior of a line and inductive-capacitive (LC) line behavior with loss dominated by the skin effect, as a function of the cross-sectional dimension, d, of the smaller conductor in the line.

but scale up the cable cross section. Since the current only flows in a layer of thickness δ (which remains the same as we scale up the cable cross section), to reduce the loss by a factor η means we must increase the *perimeter* of the conductors by a factor η, which in turn means scaling up the area by a factor η^2 (we increase the conductor dimensions by η in both cross-sectional dimensions). The cable bit-rate capacity decreases in proportion as we have to increase its area beyond that of the ideal cable, as discussed above.

In fact, one good way to deduce the bit-rate capacity of an arbitrary cable or line is to find the area of the ideal cable with the same loss at some given frequency, and simply to derate the bit-rate capacity by the ratio of the areas, as discussed above. This approach has the advantage that it gives a meaningful result for any form of cable or line, including other geometries, such as strip lines, and automatically includes the effects of larger areas and excess skin-effect losses discussed above. The loss of the ideal coaxial cable is

$$\alpha_{Vi} = \frac{R_{l\mathrm{SE}}}{2Z_o}\,\mathrm{nepers/m}, \tag{21}$$

where $R_{l\mathrm{SE}}$ is the skin-effect-limited resistance of the cable per unit length. The loss stated in this way is the voltage attenuation coefficient. Cable losses are usually quoted in power attenuation coefficients (which are twice as large because power is proportional to the square of the voltage), and in decibels per unit length. Hence, for the ideal 50-Ω copper line with polyethylene dielectric, the power attenuation coefficient in decibels becomes

$$\alpha_i = 2.28 \times 10^{-8}\,\frac{\sqrt{f}}{\sqrt{A}}\,\mathrm{dB/m}. \tag{22}$$

For each practical cable we can define a "bit-rate capacity constant" B_o through

$$B < B_o\,\frac{A}{l^2}. \tag{23}$$

Using the above scaling argument, therefore, for a chosen frequency, f, and using the area A of the real cable to calculate α_i from Eq. (22),

$$B_o = 9.1 \times 10^{15}\left(\frac{\alpha_\mathrm{i}}{\alpha_\mathrm{p}}\right)^2. \tag{24}$$

In Fig. 4, we show the "bit capacity constants," B_o, for various practical cables and lines. The parameters of cable or strip line performance used for the estimations in Fig. 4 are shown in Table I. In the table, d is the diameter of the center conductor in a coaxial cable, or the width of the (center) conductor in a strip line. Note, for the case of planar wiring (i.e., connections on chip and on MCMs) in Table I, that actual connections in planar wiring require two levels of interconnections to allow for crossing wires. The best possible case is to neglect the area required for this second layer, calculating based on the cross-sectional area of one wiring layer. A more realistic case is that two layers of wiring are required, doubling the cross-sectional area needed for a given bit-rate capacity.

Note from Fig. 4 the extent to which cables and lines of very different sizes and designs all have very similar bit-capacity constants. Even coax lines and MCM-D multichip module strip lines have overall similar performance (lines on chip, being more likely RC lines, have somewhat higher calculated bit-rate capacity constants).

As a "rule of thumb," we can conclude from Fig. 4 that, for practical cables and lines operating at or above the typical speeds of modern processors, the practical interconnect aspect ratio limit of electrical interconnect total bit-rate capacity is

$$B \sim 10^{15}\,\frac{A}{l^2} \tag{25}$$

with slightly better performance than this possible on chip (i.e., $B \sim 10^{16}\,A/l^2$).

One simple practical example is to consider the coaxial cabling required to connect cabinets within a room. The resulting cable length could be $l \sim 10$ m. The maximum cross-sectional area of cable that we might consider coming out of the back of one cabinet might be $A \sim 0.1$ m^2 (i.e., $\sim 30 \times 30$ cm total cable cross-section, which would typically correspond to $> 60 \times 60$ cm cross-section of connectors). For such a system, $l^2/A \sim 1000$ (corresponding to an aspect ratio ~ 32), and so the total bit-rate capacity that we can send into or out of one cabinet, using simple (unequalized) digital interconnects, is ~ 1 Tb/s on the basis of the "rule of thumb" relation, Eq. (25).

Another example is to understand how much information we could send internally from one side of a multichip module (MCM) to another (note that we are not considering the problem of interconnections onto or off of the module). We

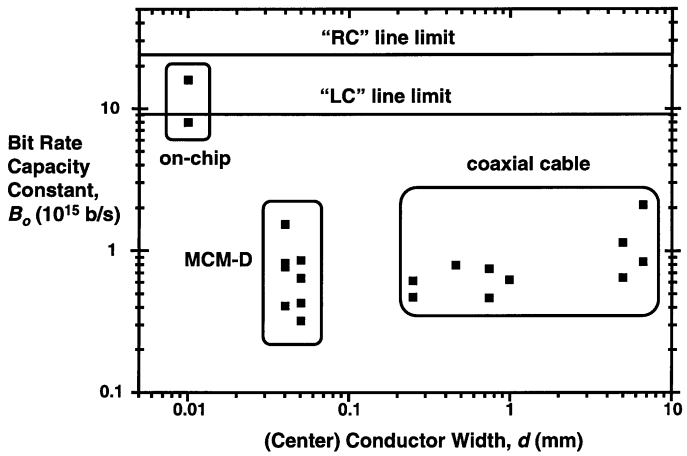

FIG. 4. Bit-rate capacity constant for various cables and transmission lines, as a function of the cross-sectional dimension, d, of the smaller conductor in the line.

presume a 5-cm square module in MCM-D technology, with a signal-carrying wiring layer 75 μm thick (including conductors

TABLE I
Parameters for Practical Interconnect Wiring

Cable type	d (mm)	A (mm^2)	α_p (dB/m) f (GHz)	f (GHz)	α_i(dB/m)	B_o(10^{15} b/s)	Notes and references
On-chip Al (1 layer)	0.01	0.0008	—	—	—	16.0	[a]
On-chip Al (2 layers)	0.01	0.0016	—	—	—	8.0	[a]
TI MCM-D (1 layer)	0.04	0.0075	35.00	3.00	14.42	1.54	[11]
	0.04	0.0075	97.50	12.40	29.32	0.82	[11]
TI MCM-D (2 layers)	0.04	0.015	35.00	3.00	10.20	0.77	[11]
	0.04	0.015	97.50	12.40	20.73	0.41	[11]
GE MCM-D (1 layer)	0.05	0.0079	26.40	1.00	8.12	0.86	[8]
	0.05	0.0079	52.90	3.00	14.07	0.64	[8]
GE MCM-D (2 layers)	0.05	0.016	26.40	1.00	5.74	0.43	[8]
	0.05	0.016	52.90	3.00	9.95	0.32	[8]
RG178	0.25	2.85	0.59	0.10	0.14	0.48	[5, 23]
	0.25	2.85	1.64	1.00	0.43	0.62	[5, 23]
RG316	0.46	5.27	1.84	3.00	0.54	0.80	[10, 23]
RG58	0.74	19.3	0.18	0.10	0.05	0.75	[23]
	0.74	19.3	0.72	1.00	0.16	0.47	[23]
RG142	0.99	19.3	2.20	12.40	0.58	0.63	[10, 23]
RG218/219	4.95	384	0.03	0.10	0.01	1.15	[23]
	4.95	384	0.14	1.00	0.04	0.65	[23]
RG220	6.60	453	0.02	0.10	0.01	2.10	[23]
	6.60	453	0.11	1.00	0.03	0.84	[23]

and dielectrics), and we take the line length for the calculation to be 10 cm, corresponding to the longest line based on "Manhattan" routing (i.e., lines in x or y directions only, with no diagonal lines allowed). The cross-sectional area of this interconnect is $A = 5 \times 0.0075 = 0.0375$ cm^2. Using a bit capacity constant of 2×10^{15} b/s (corresponding to a very good MCM-D technology) gives a bit-rate capacity for this on-MCM interconnect of 750 Gb/s.

As an example of the aspect ratio limit on chip, we consider a 1×1 cm chip, with a wiring layer effectively ~ 10 μm thick (reasonable for the final "long-distance" wiring layer on a hypothetical chip). Hence, the cross-sectional area A is $\sim 0.001 \times 1$ cm = 0.001 cm^2. As for the MCM case, we assume no diagonal interconnecting lines, so the longest line is $l = 2$ cm. Hence, using a bit capacity constant of 1.6×10^{16} b/s, we obtain a bit-rate capacity of 4 Tb/s for interconnections within the chip. This number does point out that there is likely less of a problem with high-aspect-ratio interconnects on chip than there is off chip. There would, of course, be many other problems in sending such a bit rate across chip, including the difficulty of driving them.

Note that in these examples we are considering only a specific part of an overall interconnect, not the additional problems of connecting any of these systems to other levels of the interconnect hierarchy. For example, in the coaxial cable case, we are presuming the drivers and receivers are right beside the cable ends; we make no allowance for any interconnect required to connect to the cable ends.

5. EQUALIZED CABLES

One technique commonly used to increase the bit-rate capacity of cables is equalization. A relatively simple passive network can compensate for the frequency dependence of the loss in the cable over some frequency range, at the expense of attenuation of the overall signal (see, e.g., Foster and Van Duesen [6] for an example of modern commercial practice). A simple rule of thumb for equalized cable is that the signal attenuation in the equalized cable is the loss that the unequalized cable would have at the signal clock frequency.

This can be understood intuitively. We could imagine that ideal equalization simply attenuates all frequencies below the clock frequency to have the same loss as the clock frequency signal. Hence we would have a flat frequency response up to the clock frequency, and thus the ability to carry signals with a bandwidth of the clock frequency without distortion. Since with non-return-to-zero signaling the bit rate can be twice the bandwidth, we calculate the loss in an equalized cable to be the same as the loss in the unequalized cable at a frequency $B/2$ (we calculate, for the moment, as if all of the information flows over one cable, but because the scaling will turn out to be linear in area, the same answer will be obtained if the cable is divided into many smaller cables with the same total area, just as above for the unequalized cable analysis).

We presume the receiver at the end of the equalized cable has a sensitivity such that it can tolerate a cable loss $G = \alpha_p l$. Hence, using α_i from Eq. (22),

$$B = 3.85 \times 10^{15} \left(\frac{\alpha_{\mathrm{i}}}{\alpha_{\mathrm{p}}} \right)^2 G^2 \frac{A}{l^2}, \tag{26}$$

where G is in units of decibels of power loss.

Note from Eq. (26) that, for a given receiver sensitivity (i.e., choice of G), the bit-rate capacity of an equalized cable (or set of equalized cables) has the same scaling with area A and length l as unequalized cables. Hence equalization does not solve the basic aspect ratio scaling problem of electrical interconnects.

Equalization can improve the overall bit-rate capacity of the system. For example, for a receiver that can tolerate a factor of 10 voltage loss in the signal (i.e., $G = 20$ dB), for ideal coaxial cable $B = 1.54 \times 10^{18} A/l^2$ bits/s, and for practical cable, the capacity will be approximately an order of magnitude less because of the larger loss, i.e., $B \sim 10^{17} A/l^2$ bits/s. The limits for equalized and unequalized cable are intuitively consistent; the two limits are approximately equal for the case where the (voltage) loss in the cable at the clock frequency is a factor of $\sim$2, which is approximately the level at which equalization has little benefit.

In principle, the bit-rate capacity of these equalized interconnects can be improved by increasing the sensitivity of the receiver. Note, however, from Eq. (26), that the bit-rate capacity scales only as the square of the logarithm, G, of the receiver gain, g, i.e., $B \propto (\log g)^2$. Hence, to increase B by a factor n requires that the gain of the receiver be raised to the power $\sqrt{n}$, i.e., $g \rightarrow g^{\sqrt{n}}$. For example, we might imagine increasing the receiver gain from 20 dB (corresponding to 100 mV received for 1 V transmitted) to 40 dB (corresponding to 10 mV received for 1 V transmitted) in a noisy digital environment, which would increase the bit-rate capacity by about a factor of 4, though even this would require relatively sophisticated and complex receiver circuits.

6. DISCUSSION

It is true, of course, that there are ways of trying to engineer to allow higher capacities in electrical interconnects, even when the underlying interconnect has a high aspect ratio. One way is simply to break a long interconnect up into several shorter parts. It is also possible to use more sophisticated receiver techniques to extract digital information reliably from poorer eye diagrams. The data could also be coded, for example using a Manchester code, to minimize the pattern-dependent effects, or by using multilevel modulation to increase the bit rate for a given bandwidth. Such techniques inevitably increase cost, however, and can eventually only lead to diminishing returns. They also lead to other problems, such as increased transmission delay, circuit complexity, and power dissipation. The present analysis does at least show the conditions where such techniques must be considered, and where, consequently, the cost of electrical interconnection will tend to rise faster than a simple linear scaling of cost with capacity.

It is also true that there are many other limits in electrical interconnects, many of which will set in before this aspect ratio limit in particular cases, and so it is by no means certain that we can even reach the aspect ratio limit in many cases. For example, we have neglected delay and delay variation, which is often a more important problem than rise time (though the delay is the same as rise time for the case of RC lines), and we have completely neglected the additional problems associated with connecting lines, including wave reflections from discontinuities, and the effects of the finite inductance of pins. The aspect ratio limit does, however, give an upper bound that applies at any size scale, and as a result makes a rather general case for looking at optical interconnections for certain classes of system architectures.

The argument for optics is particularly strong in that, despite some initial engineering cost for a new technology, it is nowhere near any "aspect-ratio-limited" bit-rate capacity, nor is it likely to be at any point in the foreseeable future. This argument is very clear for fiber optical interconnection. Free-space imaging interconnection systems also have essentially no problems of signal dispersion or distance-dependent loss on the length scale of electronic processors.

Optics does require driver and receiver circuits to convert between optical and electrical signals. Most such circuits in use today do consume substantial amounts of power, area, and circuit complexity. Recent work on circuits for "smart pixel" systems has, however, shown that small, simple, low-power, circuits can be made when the circuits are optimized for the specific conditions of short-range digital interconnections [12]. Recent scaling analysis of smart pixel technologies [13] (in particular, receiver power dissipation) also indicates that the capacity of optical interconnects in and out of chips may be able to scale with the increasing logical capability of silicon, and may allow optical interconnect capacities >1 Tb/s for single chips.

7. CONCLUSIONS

We have argued that there is a limit on the bit-rate capacity of electrical interconnects that depends only on the aspect

ratio of the interconnect, i.e., the ratio of interconnect length l to the square root of the total cross-sectional area of the cables or interconnect lines, $\sqrt{A}$. Ultimately, the aspect ratio is really a property of the architecture of the system; architectures with large numbers of processing units that must communicate relatively directly with many or most of the other units intrinsically have high aspect ratio interconnects, for example. Because this limit depends only on a ratio of dimensions, such high-aspect-ratio interconnects or systems will be hard to implement electrically regardless of the physical size of the system, i.e., miniaturization will not help.

We have shown that there are essentially only two kinds of behavior of lines: either the line is LC with loss dominated by the skin effect, or it is RC with bulk resistive loss. For the clock frequencies of high-performance electronic processors, all lines are LC except those on chip, which will generally be RC unless they are particularly long. The physics of LC and RC lines is substantially different, but, somewhat surprisingly, the aspect ratio limit on bit-rate capacity shows the same scaling in both cases, and also in the case of equalized LC lines. The aspect ratio limit is approximately $B \sim B_o A/l^2$ bits/s, with $B_o \sim 10^{15}$ (bit/s) for unequalized LC lines, $\sim 10^{16}$ for RC lines, and $\sim 10^{17} - 10^{18}$ for equalized LC cables. These limits are relatively independent of the details of the design of the line because of the logarithmic scalings involved in the underlying physics (though bad design could make the limits worse).

Though it may seem counter-intuitive, the numerical capacity of RC lines is somewhat higher than LC lines for the same aspect ratio of interconnect. One reason for the relatively poor performance of LC lines is that, though LC lines show a fast initial rise for a step drive, they also have a particularly long "tail." It is also interesting to note that, despite the better performance of on-chip lines, groups working on predicting the future scaling of silicon circuits expect that interconnection limits on chips will become so severe that they see a strong need to devise architectures that avoid so much interconnection [20], which is essentially a request for low-aspect-ratio architectures.

The numerical magnitudes of the aspect ratio limits are sufficiently high so that most current processors are not yet limited. Some high-performance processors are, however, likely to be experiencing problems today that ultimately arise from this aspect ratio limit. As processors move toward aggregate bit rates in the 10 Gb/s^{-1} Tb/s range flowing over their interconnection networks, such aspect ratio problems will become significant.

It should be emphasized that the aspect ratio limit can be exceeded through various engineering techniques, such as repeatering, coding, and multilevel modulation, in which cases the capacity ultimately becomes bounded by the Shannon limit. The limit discussed here indicates where such additional techniques must start to be used with electrical interconnects, showing the point where the cost of electrical interconnections will likely start to rise faster than linearly, creating practical opportunities for other approaches such as optics. For all practical purposes, optics has no such aspect ratio limit; long, thin optical interconnects work especially well. Research demonstrations of high-aspect-ratio, optically-interconnected systems do now exist. Precisely when and how optics is introduced will, however, also depend on how fast the optical technology evolves to allow systems of low enough cost and proven reliability.

ACKNOWLEDGMENT

D. A. B. Miller is pleased to acknowledge the support of Bell Laboratories, where much of this work was performed.

REFERENCES

1. Bakoglu, H. B., *Circuits, Interconnections, and Packaging for VLSI.* Addison–Wesley, New York, 1990.
2. Casimir, H. B. G., and Ubbink, J., The skin effect: II. The skin effect at high frequencies. *Philips Tech. Rev.* **28,** 300–315 (1967).
3. Clogston, A. M., Reduction of skin effect losses by the use of laminated conductors. *Bell Syst. Tech. J.* **30,** 491–529 (1951).
4. Egawa, T., Yukimatsu, K., and Yamasaki, K., Recent research trends and issues in photonic switching technologies. *NTT Rev.* **5** (1), 30–37, (1993).
5. Fink, D. G., and Christiansen, D. (Eds.), *Electronics Engineers' Handbook*, 3rd ed. McGraw–Hill, New York, 1989.
6. Foster, C., and Van Duesen, H., Extending the range of copper for fiber channel interconnects. *Electronic Packaging and Production*, 59–64 (Apr. 1996).
7. Goossen, K. W., Walker, J. A., D'Asaro, L. A., Hui, S. P., Tseng, B., Leibenguth, R., Kossives, D., Bacon, D. D., Dahringer, D., Chirovsky, L. M. F., Lentine, A. L., and Miller, D. A. B., GaAs MQW modulators integrated with silicon CMOS. *IEEE Phot. Tech. Lett.* **7,** 360–362 (1995).
8. Haller, T. R., Whitmore, B. S., Zabinski, P. J., and Gilbert, B. K., High-frequency performance of GE high-density interconnect modules. *Proc. 42nd Electronic Components and Technology Conference*. IEEE, New York, 1992, pp. 136–143.
9. Hwang, L., and Turlik, I., A review of the skin effect as applied to thin film interconnections. *IEEE Trans. Components Hybrids Manufacturing Tech.* **15,** 43–55, (1992).
10. ITT Pomona Electronics.
11. Kacines, J. J., Schaefer, T. M., and Gilbert, B. K., Performance of TI high density interconnect for 1 GHz digital MCM applications. *Proc. 45th Electronic Components and Technology Conference.* IEEE, New York, 1995.
12. Krishnamoorthy, A. V., Lentine, A. L., Goossen, K. W., Walker, J. A., Woodward, T. K., Ford, J. E., Aplin, G. F., D'Asaro, L. A., Hui, S. P., Tseng, B., Leibenguth, R., Kossives, D., Dahringer, D., Chirovsky, L. M. F., and Miller, D. A. B., 3D integration of MQW modulators over active submicron CMOS circuits: 375 Mb/s transimpedance receiver–transmitter circuit. *IEEE Phot. Tech. Lett.* **7,** 1288–1290 (1995).
13. Krishnamoorthy, A. V., and Miller, D. A. B., Scaling optoelectronic-VLSI circuits into the 21st century: A technology roadmap. *IEEE J. Selected Topics Quant. Electron.*, to be published.
14. Lentine, A. L., Novotny, R. A., Reiley, D. J., Morrison, R. L., Sasian, J. M., Beckman, M. G., Buchholz, D. B., Hinterlong, S. J., Cloonan, T. J., Richards, G. W., and McCormick, F. B., Demonstration of an experimental single chip optoelectronic switching system. Postdeadline paper PD2.5, 8th Annual Meeting, IEEE Lasers and Electro-Optics Society. IEEE, Piscataway, 1995.

15. Masaki, A., Electrical resistance as a limiting factor for high-performance computer packaging. *IEEE Circuits Devices Magazine* 22–26 (May 1989).
16. McCormick, F. B., Cloonan, T. J., Lentine, A. L., Sasian, J. M., Morrison, R. L., Beckman, M. G., Walker, S. L., Wojcik, M. J., Hinterlong, S. J., Crisci, R. J., Novotny, R. A., and Hinton, H. S., Five-stage free-space optical switching network with field-effect transistor self-electro-optic-effect-device smart-pixel arrays. *Appl. Opt.* **33,** 1601–1618 (1994).
17. Miller, D. A. B., Dense optical interconnects for silicon electronics. In A. Consortini, (Ed.), *Trends in Optics Research, Development and Applications.* Academic Press, San Diego, 1996, pp. 207–222.
18. Miller, D. A. B., Hybrid SEED massively parallel optical interconnections for silicon ICs. *Proc. 2nd Int. Conf. Massively Parallel Processing Using Optical Interconnections (MPPOI'95).* IEEE Computer Society Press, Los Alamitos, California, 1995, pp. 2–7.
19. Miller, D. A. B., Optics for low-energy communication inside digital processors: Quantum detectors, sources, and modulators as efficient impedance converters. *Optics Lett.* **14,** 146–148 (1989).
20. *National Technology Roadmap for Semiconductors.* Semiconductor Industry Association, San Jose, 1994.
21. Ozaktas, H. M., A physical approach to communication limits in computation. Ph.D. thesis, Stanford University, Stanford, CA, 1991.
22. Ozaktas, H. M., and Goodman, J. W., The limitations of interconnections in providing communication between an array of points. In S. K. Tewksbury (Ed.). *Frontiers in Computing Systems Research,* Vol. 2 Plenum Press, New York, 1991, pp. 61–130.
23. *Reference Data for Engineers: Radio, Electronics, Computer, and Communications,* 8th ed. Prentice–Hall, Carmel, Indiana, 1993.
24. Smith, B. J., Interconnection networks for shared memory parallel computerss. *Proc. 2nd Int. Conf. Massively Parallel Processing Using Optical Interconnections (MPPOI'95).* IEEE Computer Society Press, Los Alamitos, California, pp. 255–256.
25. Tkach, R. W., Derosier, R. M., Gnauck, A. H., Vengsarkar, A. M., Peckham, D. W., Zyskind J. L., Sulhoff, J. W., and Chraplyvy, A. R., Transmission of eight 20-Gb/s channels over 232 km of conventionalsingle-mode fiber. *IEEE Phot. Tech. Lett.* **7,** 1369–1371 (1995).
26. Yamanaka, N., Endo, K., Genda, K., Fukuda, H., Kishimoto, T., and Sasaki, S., 320 Gb/s high-speed ATM switching system hardware technologies based on copper–polyimide MCM. *IEEE Trans. Components Packaging Manufacturing Tech. B* **18,** 83–91 (1995).
27. Wigington, R. L., and Nahman, N. S., Transient analysis of coaxial cables considering skin effect. *Proc. IRE* **45,** 166–174 (1957).

Received December 15; revised February 15, 1996; accepted October 21, 1996

DAVID A. B. MILLER received a B.Sc. in Physics from St. Andrews University, and Ph.D. in 1979 from Heriot–Watt University. From 1981 to 1996, he worked at Bell Laboratories, in the latter period as a department head. He is currently the W. M. Keck Foundation professor of electrical engineering at Stanford University. His research interests include quantum well optics and optoelectronics and fundamentals and applications of optics in switching, interconnection, and computing systems, including "smart pixel" technologies. He has published over 180 technical papers and holds more than 30 patents. He is a Fellow of the Royal Society of London, IEEE, OSA, and APS, and was President of the IEEE Lasers and Electro-Optics Society in 1995. He was awarded the 1986 Adolph Lomb Medal of the OSA, was corecipient of the 1988 R. W. Wood Medal, and received the 1991 Prize of the International Commission for Optics.

HALDUN M. OZAKTAS was born in 1966 in Ankara, Turkey. He received a B.S. from Middle East Technical University, Ankara, in 1987, and M.S. and Ph.D. from Stanford University, California, in 1988 and 1991, respectively, all in electrical engineering. During his graduate studies he was a research assistant to Joseph W. Goodman. He joined Bilkent University, Ankara in 1991, where he is presently associate professor of electrical engineering. In 1992 he was with Adolf W. Lohmann at the University of Erlangen–Nurnberg, Bavaria as an Alexander von Humboldt Foundation Postdoctoral Fellow. Over the summer of 1994 he worked as a consultant for AT&T Bell Laboratories, New Jersey. He is the author of over 45 refereed journal articles and numerous conference presentations, 8 of which have been invited. Two of his articles have been reprinted as milestone works. His academic interests include optical information processing, signal and image processing, optoelectronic and optically interconnected computing systems, and, to a lesser extent, optical communications.