Transmitter Equalization for 4Gb/s Signalling

William J. Dally Artificial Intelligence Laboratory Massachusetts Institute of Technology billd@ai.mit.edu John Poulton Microelectronic Systems Laboratory University of North Carolina - Chapel Hill jp@cs.unc.edu

Abstract

To operate a serial channel over copper wires at 4Gb/s, we incorporate an 4GHz FIR equalizing filter into a differential transmitter. The equalizer cancels the frequency-dependent attenuation caused by the skin-effect resistance of copper wire giving a frequency response that is flat to within 5% over the band from 200MHz to 2GHz even over wires with 6dB of high-frequency attenuation. All but the last stage of the transmitter operates at 400MHz. The transmitter output stage uses a stable 10-phase 400MHz clock to sequence an array of drivers that implement the FIR filter. This paper introduces the concept of digital-signal equalization, describes the system design and circuit design of our equalizing transmitter, and presents simulation results from a 4Gb/s 0.5µm CMOS transmitter.

1. Introduction

The performance of many digital systems is limited by the interconnection bandwidth between chips, boards, and cabinets. As VLSI technology continues to scale, system bandwidth will become an even more significant bottleneck as the number of I/Os scales more slowly than the bandwidth demands of on-chip logic. Also, off-chip signalling rates have historically scaled more slowly than on-chip clock rates. Most digital systems today use full-swing unterminated signalling methods that are unsuited for data rates over 100MHz on 1m wires. Even good current-mode signalling methods with matched terminations and carefully controlled line and connector impedance are limited to about 1GHz by the frequency-dependent attenuation of copper lines. Without new approaches to high-speed signal-ling, bandwidth will stop scaling with technology when we reach these limits.

The density and speed of modern VLSI technology can be applied to overcome the I/O bottleneck they have created by building sophisticated I/O circuitry that compensates for the characteristics of the physical interconnect and cancels dominant sources of timing and voltage noise. Such optimized I/O circuitry is capable of achieving I/O rates an order of magnitude higher than those commonly used today while operating at lower power levels.

We are currently developing 0.5µm CMOS transmitter and receiver circuits that use active equalization to overcome the frequency-dependent attenuation of copper lines. Our circuits are designed to operate at 4Gb/s over up to 6m of AWG24 twisted pair or up to 1m of 5mil 0.5oz PC trace. In addition to frequency-dependent attenuation, timing uncertainty (skew and jitter) and receiver bandwidth are also major obstacles to operating at high data rates. To address all of these issues, our system includes the following components:

- 1. An active transmitter equalizer is used to compensate for the frequency-dependent attenuation of the transmission line.
- 2. Closed-loop clock recovery is performed independently for each signal line in a manner that cancels all clock and data skew and the low-frequency components of clock jitter.
- 3. The delay line used to generate the transmit and receive clocks (a 400MHz clock with 10 equally spaced phases) uses several circuit techniques to achieve a total simulated jitter of less than 20ps in the presence of supply and substrate noise. Several of our techniques are motivated by those described in [ManHor 93].
- 4. A clocked receive amplifier with a 50ps aperture time is used to sense the signal during the center of the eye at the receiver.

The availability of 4Gb/s electrical signalling will enable the design of low-cost, high-bandwidth digital systems. The wide, slow buses around which many contemporary digital systems are organized can be replaced by point-topoint networks using a single, or at most a few, high-speed serial channels resulting in significant reduction in chip and module pinouts and in power dissipation. A network based on 400MBytes/s serial channels, for example, has several times the bandwidth of a 133MBytes/s PCI-bus that requires about 80 lines. Also, depending on its topology, the network permits several simultaneous transfers to take place at full rate. A group of eight parallel channels would provide sufficient bandwidth (3.2GBytes/s) for the CPU to memory connection of today's fastest processors. For modest distances (up to 30m with 18AWG wire), high-speed electrical signalling is an attractive alternative to optical communication in terms of cost, power, and board area for peripheral connection and building-sized local-area networks.

This paper focuses on the design of the equalizer for our 4Gb/s CMOS signalling system. Section 2 discusses the problem of frequency-dependent signal attenuation in more detail. The use of equalization to compensate for this attenuation is described at the system level in Section 3 where we present the impulse and frequency response of our equalizing filter and discuss block diagrams of the implementation. Section 4 presents the circuit design and layout of the equalizing transmitter in an $0.5\mu m$ CMOS process along with simulated signal waveforms.

2. Frequency-dependent attenuation causes intersymbol interference

Skin-effect resistance causes the attenuation of a conventional transmission line to increase with frequency. With a broadband signal, as typically used in digital systems, the superposition of unattenuated low-frequency signal components with attenuated high-frequency signal components causes intersymbol interference that degrades noise margins and reduces the maximum frequency at which the system can operate.

This effect is most pronounced in the case of a single 1 (0) in a field of 0s (1s) as illustrated in Figure 1. The figure shows a 4Gb/s signal (top) and the simulated result of passing this signal across 3m of 24AWG twisted pair. The highest frequency of interest (2GHz) is attenuated by -7.6dB (42%). The unattenuated low-frequency component of the signal causes the isolated high-frequency pulse to barely reach the midpoint of the signal swing giving no eye opening and very little probability of correct detection.



Figure 1: Frequency dependent attenuation causes intersymbol interference. This figure shows a simulation of a 4Gb/s signal passed through a 3m 24AWG line. An isolated high-frequency pulse barely reaches the midpoint of signal swing because of interference from unattenuated low-frequency components of the signal.

The problem here is not the magnitude of the attenuation, but rather the interference caused by the frequencydependent nature of the attenuation. The high-frequency pulse has sufficient amplitude at the receiver for proper detection. It is the offset of the pulse from the receiver threshold by low-frequency interference that causes the problem. In Section 3, we will see how using a transmitter equalizer to preemphasize the high-frequency components of the signal eliminates this problem. However, first we will characterize the nature of this attenuation in more detail.

2.1 Skin depth determines line attenuation

At high frequencies (above 100MHz), current is carried primarily on the surface of the conductor, dropping off to a value of e^{-1} at a depth of

$$\delta = \left(\pi f \mu \sigma\right)^{-1/2} \tag{1}$$

where σ is the conductivity of the material (5.8E7 mhos/m for copper) [Matick 69].

For a round conductor with radius *r*, this gives a resistance per unit length (ohms/m) of

$$R(f) = \frac{1}{2r} \left(\frac{\mu f}{\pi \sigma}\right)^{1/2} \tag{2}$$

A thin strip-guide with width w has a resistance per unit length of

$$R(f) = \frac{1}{2w} \left(\frac{\pi\mu f}{\sigma}\right)^{1/2} \tag{3}$$

In both cases the resistance is proportional to the square root of the frequency and the inverse of the linear dimension of the conductor.

$$R(f) = K_R d^{-1} f^{1/2}$$
(4)

where *d* is the linear dimension (radius or width) of the conductor (in meters) and K_R is 4.15E-8 ohms-s^{1/2} for a round conductor and 1.3E-7 ohms-s^{1/2} for a thin rectangular stripguide.

Over an infinitesimal section of line, with length dx, an incident wave with magnitude V_i , drops a voltage across the resistance R(f)dx of

$$dV_{i}(x) = I_{i}(x)R(f)dx = V_{i}(x)\frac{R(f)dx}{Z_{0}}$$
(5)

Solving this differential equation gives the attenuation for a line of length x.

$$A(f, x) = \exp\left(-\frac{R(f)}{Z_0}x\right)$$
(6)

Attenuation is also caused by absorption in the dielectric of the transmission line, by radiation of signal energy, by the frequency response of the package parasitics, and by any lumped capacitance at the load. In many applications, however, the skin-effect attenuation dominates these effects.

2.2 Attenuation examples



Figure 2: Resistance (top) and attenuation (bottom) curves for 1m of 30AWG 100 Ω twisted pair (left) and 1m of 5mil 0.5oz 50 Ω stripguide (right).

Figure 2 shows the resistance per meter and the attenuation per meter as a function of frequency for a 30AWG ($d = 128\mu$ m) twisted pair with a differential impedance of $Z_0=100\Omega$ and for a 5mil ($d = 125\mu$ m) half-ounce (0.7mil thick) 50 Ω stripguide. For the 30AWG pair¹, the skin effect begins increasing resistance at 267KHz and results in an attenuation to 56% of the original magnitude (-5dB) per meter of cable at our operating frequency of 2GHz corresponding to a bit rate of 4Gb/s.

^{1.} To account for resistive drops in both elements of the pair, we double the resistance R(f), in computing attenuation according to (6).

Skin effect does not begin to effect the 5mil PC trace until 43MHz because of its thin vertical dimension. The high DC resistance ($6.8\Omega/m$) of this line gives it a DC attenuation of 88% (-1.2dB). Above 70MHz the attenuation rolls off rapidly reaching 40% (-8dB) at 2GHz. The important parameter, however, is the difference between the DC and high-frequency attenuation which is 45% (-6.8dB).

2.3 Attenuation reduces signal quality



Figure 3: Without equalization (left), attenuating high-frequency components by a factor A reduces the height of the data eye by a factor of 2A-1 and reduces the width of the eye. This inter-symbol interference also causes trailing-edge jitter (center). With equalization (right) the height of the eye is reduced by a factor of A and the width of the eye is unaffected.

The effect of frequency dependent attenuation is graphically illustrated in the 'cartoon' eye-diagrams of Figure 3. As shown in the waveform on the left, without equalization, a high-frequency attenuation factor of *A* reduces the height of the eye opening to $2A \cdot I$ with the eye completely disappearing at $A \le 0.5$. This height is the amount of effective signal swing available to tolerate other noise sources such as receiver offset, receiver sensitivity, crosstalk, reflections of previous bits, and coupled supply noise. Because the waveforms cross the receiver threshold offset from the center of the signal swing, the width of the eye is also reduced. As illustrated in the center of Figure 3, the leading edge of the attenuated pulse crosses the threshold at the normal time. The trailing edge, however, is advanced by $t_j = (1-A)t_r$. This data-dependent jitter causes greater sensitivity to skew and jitter in the signal or sampling clock and may introduce noise into the timing loop.

The waveform on the right of Figure 3 illustrates the situation when we equalize the signal by attenuating the DC and low frequency components so all components are attenuated by a factor of A. Here the height of the eye opening is A, considerably larger than 2A-1, especially for large attenuations. Also, because the waveforms cross at the midpoint of their swing, the width of the eye is a full bit-cell giving better tolerance of timing skew and jitter.

3. Preemphasizing signal transitions equalizes line attenuation

Equalization eliminates the problem of frequency-dependent attenuation by filtering the transmitted or received waveform so the concatenation of the equalizing filter and the transmission line gives a flat frequency response. With equalization, an isolated 1 (0) in a field of 0s (1s) crosses the receiver threshold at the midpoint of its swing, as shown in Figure 3 (right), rather than being offset by an unattenuated DC component, as shown in Figure 3 (left). Narrow-band voice, video, and data modems have long used equalization to compensate for the linear portion of the line characteristics [LeeMes 94]. However, it has not been used to date in broadband short-distance digital signalling.

We equalize the line using an 4GHz FIR filter built into the current-mode transmitter. The arrangement is similar to the use of Tomlinson precoding in a modem [Tomlin 71]. In a high-speed digital system it is much simpler to equalize at the transmitter than at the receiver, as is more commonly done in communication systems. Equalizing at the transmitter allows us to use a simple receiver that just samples a binary value at 4GHz. Equalizing at the receiver would require an A/D of at least a few bits resolution or a high-speed analog delay line, both difficult circuit design problems. A discrete-time FIR equalizer is preferable to a continuous-time passive or active filter as it is more easily realized in a standard CMOS process.

3.1 The equalizing filter has a high-pass frequency response





After much experimentation we have selected a five-tap FIR filter that operates at the bit rate. The weights are trained to match the filter to the frequency response of the line. For a 1m 30AWG line, the impulse response is shown in Figure 4(a). Each vertical line delimits a time interval of one bit-cell or 250ps. The filter has a high-pass response as shown in Figure 4(b).

As shown in Figure 5, this filter cancels the low-pass attenuation of the line giving a fairly flat response over the frequency band of interest (the decade from 200MHz to 2GHz). We band-limit the transmitted signal via coding to eliminate frequencies below 200MHz. The equalization band is limited by the length of the filter. Adding taps to the filter would widen the band. We have selected five taps as a compromise between bandwidth and cost of equalization. Each panel of Figure 5 shows (a) the response of the line, (b) the response of the filter, and (c) the overall response of the system (the product of (a) and (b)). The filter cancels the response of parasitics as well as the response of the line. The left panel of Figure 5 depicts the equalization of 1m of 30AWG twisted pair. The right panel shows the result of training the filter on the same line but with an additional 1pF parasitic load at the receiver. In both cases the response is flat to within 5% across the band of interest. (Note that the scale on the bottom panels is compressed to exaggerate this effect).



Figure 5: Frequency response of filter (top), line (middle) and combination (bottom) for 1m of 30AWG cable (left) and the same cable followed by a 1pF load capacitor (right). The scale on the bottom panels is compressed to exaggerate the effect.

The filter results in all transitions being full-swing, while attenuating repeated bits. Figure 4(d) shows the response of the filter to an example data sequence shown in Figure 4(c) (0000100000101011110000). The example

shows that each signal transition goes full swing with the current stepped down to an attenuated level for repeated strings of 1s (0s).



Figure 6: Response of equalizing filter to waveform from Figure 1. The left panel repeats Figure 1 showing the original 4Gb/s signal and the received waveform after a 3m 24AWG line without equalization. The right panel shows the signal after being equalized (top) and the resulting received waveform (bottom).

Figure 6 illustrates the application of equalization to the example of Figure 1. The left half of the figure repeats the previous figure showing the response of a 3m 24AWG line and receiver parasitics to a 4Gb/s sequence. The isolated pulses are undetectable. The right side of the figure shows the filtered version of the original signal and the received waveform. With equalization the isolated pulses and high-frequency segments of the signal are centered on the receiver threshold and have adequate eye openings for detection.



3.2 The 4Gb/s transmitter is realized with 400MHz circuitry

Figure 7: The transmitter is realized using 400MHz current-steering circuitry. A 10-phase clock sequences 10 DACs that drive measured 250ps current pulses onto the differential output.

A block diagram of the transmitter is shown in Figure 7. The transmitter accepts 10 bits of data, $D_{0.9}$, at

400MHz. A distribution block delivers 5 bits of data to each of the 10 FIR filters. The ith filter receives bit D_i and the four previous bits. For the first four filters this involves delaying bits from the previous clock cycle. The distribution also retimes the filter inputs to the clock domain of the filter. Each filter is a 5-tap *transition filter* that produces a 4-bit output encoded as 3 bits of positive drive and 3 bits of negative drive. These six bits from the filter directly select which of six pulse generators in the DAC connected to that filter are enabled. The enabled pulse generators are sequenced by the 10-phase clock. The ith pulse generator is gated on by ϕ i and gated off by ϕ i+1. To meet the timing requirements of the pulse generator, the ith filter operates off of clock ϕ i+1.

To simplify the implementation each FIR filter is approximated by a transition filter implemented with a lookup table as illustrated in Figure 8. The transition filter compares the current data bit, D_i , to each of the last four bits and uses a find-first-one unit to determine the number of bits since the last signal transition. The result is used to look up a 3-bit drive strength for the current bit from a 15-bit serially-loaded RAM. The drive strength is multiplied by the current bit with six NAND gates to generate three-bit high and low drive signals for the DAC. While the transition filter is a non-linear element, it closely approximates the response of an FIR filter for the impulse functions needed to equalize typical transmission lines. Making this approximation greatly reduces the size and delay of the filter as a 96bit RAM would be required to implement a full 5-tap FIR filter via a lookup table.



Figure 8: A transition filter approximates the FIR filter by looking up a magnitude depending on the number of bits since the last transition.

4. Circuit details

We have designed a prototype equalizing transceiver chip in an 0.6µm drawn process, HP14 using scalable rules. The layout of the transmitter section of this chip is illustrated in Figure 15 (attached to the paper). In addition to the elements shown in Figure 7, this chip also includes a pattern generator module and seven on-chip sampling amplifiers. The pattern generator is used to generate test patterns for the transmitter and consists of a 20-bit pseudo-random number generator, an 80-bit serially loaded pattern RAM, and a pattern ROM containing the synchronization sequence. The on-chip samplers are used to probe repetitive high-speed on-chip waveforms by comparing the on-chip signal to an externally generated analog reference level at a time determined by an externally provided differential clock signal. The transmitter, less the pattern generator, measures 550µm x 900µm.

The circuit design of the DAC is shown in Figure 9. Figure 9(a) shows how each DAC module is composed of three progressively sized differential pulse generators. Each generator is enabled to produce a current pulse on Dout+ (Dout–) if the corresponding Π (Σ) line is low. If neither line is low no pulse is produced. Depending on the current bit and the three-bit value read from the RAM in the filter module, 15 different current values are possible (nominally from –8.75mA to +8.75ma in 1.25mA steps). The timing of the pulse is controlled by a pair of clocks. A low-going on-clock, $\overline{\phi}_i$, gates the pulse on its falling edge. The high-true off clock, ϕ_{i+1} , gates the pulse off 250ps later.

Each of the three differential pulse generators is implemented as shown in Figure 9(b). A pre-drive stage inverts the on-clock and qualifies the off-clock with the enable signals. A low (true) enable signal, which must be stable while the off-clock is low, turns on one of the two output transistors priming the circuit for the arrival of the on-clock. When the on-clock falls, the common tail transistor is turned on starting the current pulse. When the off-clock rises, the selected output transistor terminates the current pulse. The qualifying NOR-gate is carefully matched against the on-clock inverter to avoid distorting the pulse width.



Figure 9: Circuit design for a DAC module - (a) three pulse generators are enabled by the H and L signals and gated by two clocks to generate a precise 250ps pulse with one of 15 selectable current levels, (b) each of the three generators is implemented with a qualifying pre-driver followed by a series final driver that shares a common tail transistor.

Results from HSPICE simulation of the extracted transmitter layout are shown in Figure 10. The left panel shows the transmitter output (top) and the receiver input (bottom) with equalization enabled. The top waveform shows the pre-emphasis of transitions and isolated pulses. The bottom waveform shows how this preemphasis results in a clean bit-stream at the receiver with equal amplitude (about 300mV) for high- and low-frequency components of the signal.

The center panel shows waveforms for the transmitter operating with equalization disabled. The transmit waveform shows some attenuation of the high-frequency components due to slew-rate limitations of the driver. The bottom waveform of this panel is highly distorted by the high-frequency attenuation of the package parasitics and transmission line. The low-frequency components appear with minimal attenuation (about 600mV levels) while isolated pulses are severely attenuated (about 300mV). The result is a signal where several bits are clearly undetectable.

The right panel of Figure 10 shows differential eye diagrams constructed from the two receiver waveforms. The waveform with equalization on the top shows a clean eye opening that encompasses about 50% of the received signal swing and, before adding clock jitter, about 70% of the bit cell. The bottom trace, without equalization, has no opening at all. Equalization has clearly improved both the voltage and timing margins of the received waveform.



Figure 10: Simulation Results - (a) simulated waveforms with equalization on, top trace is at transmitter, bottom trace is at receiver, (b) waveforms with equalization off, (c) differential eye diagrams of received waveform with equalization (top) and without (bottom).

Figure 11 shows the waveforms from the 10-phase (5-phase complementary) clock generator that controls the timing of the transmitter. The generator is realized as a six-stage differential delay line with the delay of each stage controlled by a feedback loop to keep ϕ 1 and ϕ 6 180 degrees out of phase. The left panel of the figure shows the clock outputs when the loop is in steady-state. For comparison, the vertical lines are spaced at 250ps intervals. The right panel illustrates the dynamics of the loop converging by showing the two signals that control delay during powerup. The feedback loop directly drives the current-source bias voltage (top) and the load control voltage (bottom) is

generated by a replica-bias circuit [ManHor 93]. The figure shows that the loop converges to a stable state after less than 250ns.



Figure 11: Waveforms from the 10-phase clock generator (a) the generated clock phases, (b) control voltages during powerup.

5. Receiver

A block diagram of our 4Gb/s receiver is shown in Figure 12. A demultiplexing receiver samples the differential input stream every 125ps with sequencing controlled by a 20- ϕ clock. Each 400MHz major cycle the receiver takes 20 samples, ten data samples d_{0:9} taken from the centers of bit cells, and ten edge samples, e_{0:9} taken from the boundaries between bit cells. The data samples are input to a funnel shifter that concatenates the current ten samples with the previous nine samples and then selects a contiguous ten-bit field of this 19-bit sequence to output. The selection is set up during training to restore proper framing to the parallel output. In effect it rounds up the delay of the cable to be a multiple of 10 bit cells. The edge samples are used, along with the data samples, by the clock control unit to continuously adjust the phase of the 20-sample clocks to keep the even (data) samples centered on the eyes of the incoming stream.



Figure 12: Receiver block diagram. A demultiplexing receiver sequenced by a 20- ϕ clock samples the input stream each 125ps. The even samples are output as data after shifting to restore framing. The odd samples are used to align the clock with the data eye.

Figure 13, shows a more detailed view of the demultiplexing receiver. The 20-phase clock sequences 21 clocked sense amplifiers. The even clocks generate the data samples with d_i being sampled by ϕ_{2i} . The edge samples are sequenced by the odd clocks with e_i being sampled by ϕ_{2i+1} . To keep loads balanced and lines short, sample d_0 is repeated at the end of the line. Each sample is in a separate clock domain. A stage of retiming latches, not shown, is used to align all of the samples into a single clock domain.



Figure 13: The demultiplexing receiver consists of 21 clocked sense amplifiers sequenced by the 20-phase clock.

The 20-phase clock is controlled by two timing loops as illustrated in Figure 14 (left). The 400MHz input clock drives a digitally controlled delay line with a dynamic range of three bit cells. This line sets the phase relationship between the 400MHz input clock and $\phi 0$ as determined by the digital variable, phase. The output of this line, ϕx , drives a ten-stage differential tapped delay line that generates the 20 precisely spaced clock phases. The ith stage generates the complementary signals ϕ_i and ϕ_{i+10} . The delay of each stage of this line is set to exactly 1/20 the period of the input clock using an analog control voltage set by a phase comparator that aligns $\phi 10$ ' with $\phi 0$. This line is colocated with the receive amplifiers and the loads on all traces are carefully balanced to match delays.



Figure 14: Clock generation: (left) Two timing loops control the 20- ϕ clock, (right) Clock phase is adjusted by a hybrid analog/digital control circuit.

The phase control for the first delay line is adjusted using the circuit shown in Figure 14 (right). If there is a transition between the ith and $i+1^{st}$ data samples, d_i and d_{i+1} , signal trans_i will be true. On a transition, the state of the edge sample, e_i , between these two data samples is examined to see if the transition is early or late. If e_i and d_i differ, the transition has occurred before the edge clock, and thus the clock is late. If the these two adjacent samples agree and trans_i is high, then the clock is early, before the transition. An analog summing network combines the ten early signals and the ten late signals and produces a single up/down command pair to drive a counter that controls the clock phase.

6. Conclusion

Transmitter equalization extends the data rates and distances over which electronic digital signalling can be reliably used. Preemphasizing the high-frequency components of the signal compensates for the low-pass frequency response of the package and transmission line. This prevents the unattenuated low-frequency components from interfering with high-frequency pulses by causing offsets that prevent detection. With equalization an isolated pulse at the receiver has the same amplitude as a long string of repeated bits. This gives a clean received signal with a good eye opening in both the time and voltage dimensions.

We implement equalization for a 4Gb/s signalling system by building an 4GHz, five-tap FIR filter into the transmitter. This filter is simple to implement yet equalizes the frequency response to within 5% across the band of interest. The filter is realized using 0.5μ m CMOS circuitry operating at 400MHz using a bank of 10 filters and DACs sequenced by a 10-phase 400MHz clock. Narrow drive periods are realized using series gating to combine two clock phases, an on-phase and off-phase, in each DAC. We have simulated extracted layout of the equalized transmitter driving a load through package parasitics and 1m of differential strip guide to demonstrate the feasibility of this approach.

The equalizing transmitter described here is one component of a 4Gb/s signalling system we are currently developing for implementation in an 0.5µm CMOS technology. The system also relies on low-jitter timing circuitry, automatic per-line skew compensation, a narrow-aperture receive amplifier, and careful package design.

The availability of 4Gb/s serial channels in a commodity CMOS technology will enable a range of system opportunities. The ubiquitous system bus can be replaced by a lower-cost yet higher-speed point-to-point network. A single *hub* chip with 32 serial ports can directly provide the interconnection for most systems and can be assembled into more sophisticated networks for larger systems. A single 4Gb/s serial channel provides adequate bandwidth for most system components and multiple channels can be ganged in parallel for higher bandwidths.

A 4Gb/s serial channel can also be used as a replacement technology at both the component and system level. At the component level, a single serial channel (two pins) replaces 40 100MHz pins. A 4GByte/s CPU to L2 cache interface, for example, could be implemented with just eight serial channels. At the system level, high-speed electrical serial channels are a direct replacement for expensive optical interconnect. Using 18AWG wire, these channels will operate up to lengths of 10m enabling high-bandwidth, low-cost peripheral connections and local-area networks. Inexpensive electrical repeaters can be used to operate over substantially longer distances.

Even with 4Gb/s channels, system bandwidth remains a major problem for system designers. On-chip logic bandwidth (gates x speed) is increasing at a rate of 90% per year (60% gates and 20% speed). The density and bandwidth of system interconnect is increasing at a much slower rate of about 20% per year as they are limited by mechanical factors that are on a slower growth curve than that of semiconductor lithography. A major challenge for designers is to use scarce system interconnect resources effectively, both through the design of sophisticated signal-ling systems that use all available wire bandwidth and through system architectures that exploit locality to reduce the demands on this bandwidth.

7. Acknowledgments

This research was supported in part by the Defense Advanced Research Projects Agency (DARPA) under ARPA order 8272 monitored by the Air Force Electronic Systems Division under contract F19628-92-C-0045, in part by DARPA under ARPA order E253 contract DABT63-96-C-0039, and in part by DARPA under ARPA order A410, with additional support from the National Science Foundation under Grant No. MIP-9306208.

The authors are indebted to Mark Horowitz and Tom Knight for many helpful comments and suggestions.

8. References

[LeeMes 94] Lee, Edward A., and Messerschmitt, David G., *Digital Communication, Second Edition*, Kluwer, 1994. [ManHor 93] Maneatis, J. and Horowitz, M., "Precise Delay Generation Using Coupled Oscillators," *IEEE JSSC*, Vol 28, No. 12, pp 1273-1282.

[Matick 69] Matick, Richard E. *Transmission Lines for Digital and Communication Networks*, McGraw-Hill, 1969. [Tomlin 71] Tomlinson, M., "New Automatic Equalizer Employing Modulo Arithmetic," *Electronic Letters*, March 1971.