

TRANSMITTER EQUALIZATION FOR 4-GBPS SIGNALING

William J. Dally

Massachusetts Institute of
Technology

John Poulton

University of North
Carolina at Chapel Hill

NOTICE: THIS MATERIAL MAY BE
PROTECTED BY COPYRIGHT LAW
SEE TITLE PAGE



Incorporating a 4-GHz FIR equalizing filter into a differential transmitter allows a serial channel to operate over copper wires at 4 Gbps.

For many digital systems, the major performance-limiting factor is the interconnection bandwidth between chips, boards, and cabinets. As VLSI technology continues to scale, system bandwidth will become an even more significant bottleneck, because the number of I/Os will scale more slowly than the bandwidth demands of on-chip logic. Also, off-chip signaling rates have historically scaled more slowly than on-chip clock rates.

Most digital systems today use full-swing, unterminated signaling methods that are unsuited for data rates over 100 MHz on 1-meter wires. Even good current-mode signaling methods with matched terminations and carefully controlled line and connector impedance are limited to about 1 GHz by the frequency-dependent attenuation of copper lines. Without new approaches to high speed signaling, bandwidth will stop scaling alongside technology when we reach these limits.

However, we can apply the density and speed of modern VLSI technology to overcome this bottleneck. Sophisticated I/O circuitry can compensate for the characteristics of the physical interconnect and cancel dominant sources of timing and voltage noise. Such optimized circuitry can achieve I/O rates an order of magnitude higher than those commonly used today—while operating at lower power levels.

We are currently developing 0.5-micron CMOS transmitter and receiver circuits that use active equalization to overcome the frequency-dependent attenuation of copper lines. The circuits will operate at 4 Gbps over up to 6 meters of 24AWG twisted pair or up to 1 meter of 5-mil 0.5-oz. PC trace.

In addition to frequency-dependent attenuation, timing uncertainty (skew and jitter) and receiver bandwidth are also major obstacles to high data rates. To address all of

these issues, we've given our system the following characteristics:

- An active transmitter equalizer compensates for the frequency-dependent attenuation of the transmission line.
- The system performs closed-loop clock recovery independently for each signal line in a manner that cancels all clock and data skew and the low-frequency components of clock jitter.
- The delay line that generates the transmit and receive clocks (a 400 MHz clock with 10 equally spaced phases) uses several circuit techniques to achieve a total simulated jitter of less than 20 ps in the presence of supply and substrate noise. (Several of our techniques draw on those described in Maneatis and Horowitz.)
- A clocked receive amplifier with a 50-ps aperture time senses the signal during the center of the eye at the receiver.

The availability of 4-Gbps electrical signaling will enable the design of low-cost, high-bandwidth digital systems. Designers will be able to replace the wide, slow buses around which many contemporary digital systems are organized with point-to-point networks using a single—or at most a few—high-speed serial channels. This will result in significant reduction in chip and module pinouts and in power dissipation.

A network based on 400-Mbyte/s serial channels, for example, has several times the bandwidth of a 133-Mbyte/s PCI bus that requires about 80 lines. Also, depending on its topology, the network permits several simultaneous, full-rate transfers. A group of eight parallel channels would provide sufficient bandwidth (3.2 Gbytes/s) for the CPU-to-memory connection of today's fastest




Figure 1. Frequency-dependent attenuation causes intersymbol interference: a simulation of a 4-Gbps signal (a) passed through a 3-meter, 24AWG line (b). An isolated high-frequency pulse barely reaches the midpoint of signal swing because of interference from unattenuated low-frequency components of the signal.

processors. For modest distances (up to 30 meters with 18AWG wire), high-speed electrical signaling is an attractive alternative to optical communication in terms of cost, power, and board area for peripheral connection and building-size local-area networks.

Frequency-dependent attenuation

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. This interference degrades noise margins and reduces the maximum frequency at which the system can operate.²

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 4-Gbps signal (Figure 1a) and the simulated result of passing this signal across 3 meters of 24AWG twisted pair (Figure 1b). The highest frequency of interest (2 GHz) is attenuated by -7.6 dB (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.

The problem here is not the magnitude of the attenuation, but rather the interference caused by the attenuation's frequency-dependent nature. 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. Later, 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.

Skin depth determines line attenuation

At high frequencies (above 100 MHz), current is carried primarily on the surface of the conductor, dropping off to a value of e^{-1} at a depth of $\delta = (\pi f \mu \sigma)^{-1/2}$, where σ is the conductivity of the material (5.8×10^7 mhos/m for copper).³

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

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

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

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

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

Thus,

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

where d is the linear dimension (radius or width) of the conductor (in meters) and K_r is 4.15×10^{-8} ohms- $s^{1/2}$ for a round conductor and 1.3×10^{-7} ohms- $s^{1/2}$ for a thin rectangular strip guide.

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

$$dV_r(x) = I_r(x) R(f) dx = V_r(x) \frac{R(f) dx}{Z_0}$$

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

$$dV_r(x) = \exp\left(-\frac{R(f)}{Z_0} x\right) \quad (1)$$

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.

Attenuation examples

Figure 2 (next page) shows the resistance per meter and the attenuation per meter as a function of frequency for a 30AWG ($d = 128 \mu\text{m}$) twisted-pair cable with a differential impedance of $Z_0 = 100$ ohms and for a 5-mil ($d = 125 \mu\text{m}$), half-ounce (0.7-mil-thick) 50-ohm strip guide. For the 30AWG pair, the skin effect begins increasing resistance at 267 kHz. This results in an attenuation to 56% of the original magnitude (-5 dB) per meter of cable at our operating frequency of 2 GHz, corresponding to a bit rate of 4 Gbps. To account for resistive drops in both elements of the pair, we double resistance $R(f)$ in computing attenuation according to Equation 1.




Figure 2. Resistance (top) and attenuation (bottom) curves for 1 meter of 30AWG, 100-ohm twisted pair (a) and 1 meter of 5-mil, 0.5-oz., 50-ohm strip guide (b).




Figure 3. Without equalization, 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 (a). This intersymbol interference also causes trailing edge jitter (b). With equalization, attenuation reduces the height of the eye by A , and leaves the width of the eye unchanged compared to the original signal (c).

Skin effect does not begin to affect the 5-mil PC trace until 43 MHz because of its thin vertical dimension. This line's high dc resistance (6.8 ohms/m) gives it a dc attenuation of 80% (-1.2 dB). Above 70 MHz, the attenuation rolls off rapidly, reaching 40% (-8 dB) at 2 GHz. The important parameter, however, is the difference between the dc and high-frequency attenuation, which is 45% (-6.8 dB).

Attenuation reduces signal quality

The cartoon eye diagrams of Figure 3 illustrate the effect of frequency-dependent attenuation. As shown in the waveform in Figure 3a, without equalization, a high-frequency attenuation factor of A reduces the height of the eye opening to $2A - 1$, with the eye completely disappearing at $A \leq 0.5$. This height is the amount of effective signal swing available to tolerate other noise sources such as receiver offset, receiver sensitivity, cross talk, reflections of previous bits, and coupled supply noise. Because the waveforms cross the receiver threshold offset from the center of the signal swing, attenuating the high-frequency components also reduces the width of the eye.

As illustrated in Figure 3b, the leading edge of the attenuated pulse crosses the threshold at the normal time. The trailing edge, however, advances 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 in Figure 3c 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.

Preemphasizing signal transitions

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 3c, rather than being offset by an unattenuated dc component, as shown in Figure 3a. Narrowband voice, video, and data modems have long used equalization to compensate for the linear portion of the line characteristics. To date, however, broadband, short-distance digital signaling has not applied this technique.

We equalize the line using an 4-GHz FIR filter built into the current mode transmitter. The arrangement is similar to the

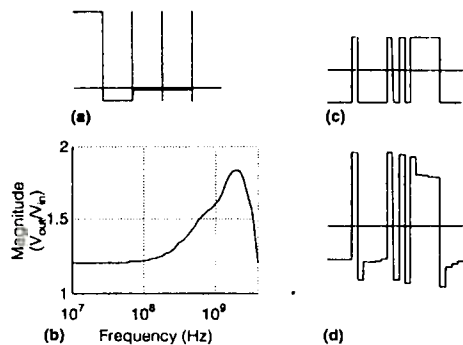


Figure 4. Impulse response (a), frequency response (b), example sequence of a five-tap FIR equalizing filter matched to 1 meter of 30AWG 100-ohm line (c), and response to the example sequence (d).

use of Tomlinson precoding in a modem.⁵ 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 4 GHz. Equalizing at the receiver would require an analog-to-digital converter of at least a few bits of resolution or a high-speed analog delay line, both difficult circuit design problems. We prefer a discrete-time FIR equalizer to a continuous-time passive or active filter, because it is more easily realized in a standard CMOS process.

High-pass frequency response

After much experimentation, we selected a five-tap FIR filter that operates at the bit rate. We trained the weights to match the filter to the frequency response of the line. Figure 4a shows the impulse response for a 1-meter 30AWG line. Each vertical line delimits a time interval of one bit-cell or 250 ps. Figure 4b shows the filter's high-pass response.

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

The filter results in all transitions being full swing, while attenuating repeated bits. Figure 4d shows the filter's response to an example data sequence shown in Figure 4c (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 (next page) illustrates the application of equalization to the example of Figure 1. Figure 6a repeats the earlier figure showing the response of a 3-meter, 24AWG line and receiver parasitics to a 4-Gbps sequence. The isolated pulses are undetectable. Figure 6b shows the filtered version

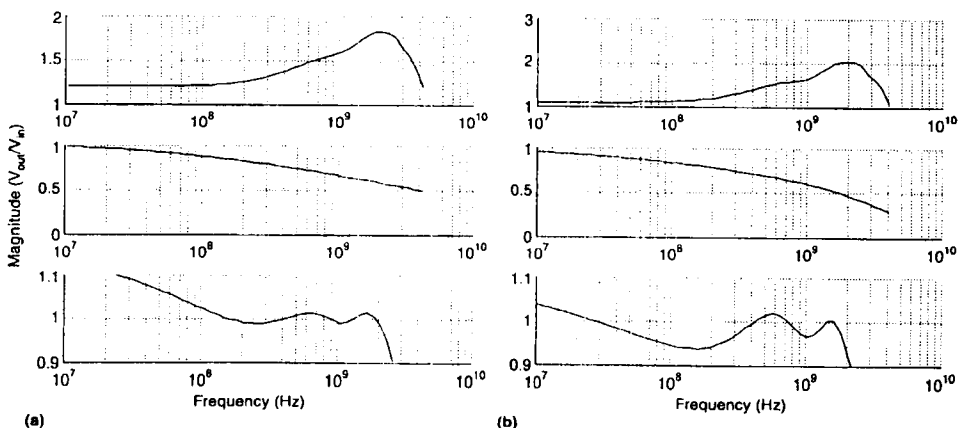


Figure 5. Frequency response of filter (top), line (middle), and combination (bottom), for 1 meter of 30AWG cable (a) and the same cable followed by a 1-pF load capacitor (b). We have compressed the scale on the bottom panels to exaggerate the effect.




Figure 6. Response of equalizing filter to waveform from Figure 1. Repetition of Figure 1 showing the original 4-Gbps signal (top) and the received waveform after a 3-meter, 24AWG line without equalization (bottom) (a); the signal after being equalized (top) and the resulting received waveform (bottom) (b).

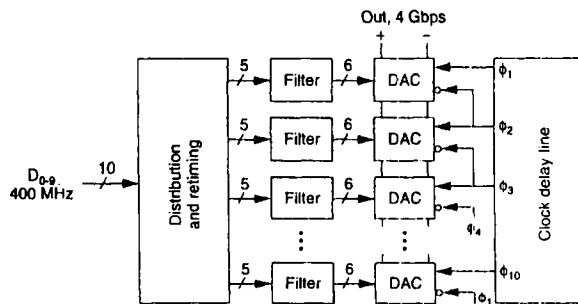


Figure 7. We used 400-MHz current-steering circuitry to build the transmitter. A 10-phase clock sequences 10 DACs that drive measured 250-ps current pulses onto the differential output.




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

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.

4-Gbps transmitter with 400-MHz circuitry

Figure 7 shows the transmitter's block diagram. The transmitter accepts 10 bits of data, D_{0-9} , at 400 MHz. A distribution block delivers 5 bits of data to each of the 10 FIR filters. The i th 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 block also retimes the filter inputs to the clock domain of the filter. Each filter is a five-tap transition filter that produces a 4-bit output encoded as three bits of positive drive and three bits of negative drive. These six bits from the filter directly select which of six pulse generators to enable in the DAC connected to that filter. The enabled pulse generators are sequenced by the 10-phase clock. The i th 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 i th filter operates from clock ϕ_{i+1} .

To simplify the implementation, we approximated each FIR filter by a transition filter implemented with a lookup table, as illustrated in Figure 8. The transition filter compares 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 filter uses the result to look up a 3-bit drive strength for the current bit from a 15-bit serially loaded RAM. It uses six NAND gates to multiply the drive strength by the current bit to generate 3-bit high and low drive signals for the DAC. While the transition filter is a nonlinear 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 96-bit RAM would be required to implement a full five-tap FIR filter via a lookup table.

Circuit details

We have designed a prototype equalizing transceiver chip in an 0.6-micron drawn process, HP14, using scalable rules. Figure 9 shows the layout of the chip's transmitter section. In addition to the elements shown in Figure 7, this chip also includes a pattern generator module and eight on-chip sampling amplifiers. The pattern generator generates test patterns for the transmitter and consists of a 20-bit pseudorandom-number generator, an 80-bit serially loaded

S
r
0
5
l-
d
it
s
e
l-
e
r-
it
e
s
y
s
o
r-
e
is
e
of
r

1.
y
h
n
r-
of
t-
of
n.
ip
nt
es
te
n-
es
d-
x-
a
ip

ip
s.
n.
so
n-
W-
at
ed

pattern RAM, and a pattern ROM containing the synchronization sequence. The on-chip samplers 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.

Figure 10 shows the circuit design of the DAC. Figure 10a shows how each DAC module is composed of three progressively sized differential pulse generators. Each generator is enabled to produce a current pulse on D_{out} (D_{out-}) if the corresponding H(L) line is low. If neither line is low, the generator produces no pulse. Depending on the current bit and the 3-bit value read from the RAM in the filter module, 15 different current

values are possible (nominally from -8.75 to $+8.75$ mA in 1.25-mA steps). A pair of clocks controls the pulse timing. A low-going on clock, ϕ_n , gates the pulse on its falling edge. A high-true off clock, ϕ_{n+1} , gates the pulse off 250 ps later.

We implemented each of the three differential pulse generators as shown in Figure 10b. A predrive 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 turns 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 11 shows the results from HSPICE simulation of the extracted transmitter layout. Figure 11a shows the transmitter output (top) and the receiver input (bottom) with equalization enabled. The top waveform shows the preemphasis of transitions and isolated pulses. The bottom waveform shows how this preemphasis results in a clean bitstream at the receiver with equal amplitude (about 300 mV) for high- and low-frequency components of the signal.




Figure 11. Simulated waveforms with equalization on (a), top trace at transmitter, bottom trace at receiver; waveforms with equalization off (b); and differential eye diagrams (c) of received waveform with (top) and without equalization (bottom).




Figure 9. Layout of the transmitter circuitry in a 0.6- μ m (drawn) CMOS process. The transmitter, less the pattern generator, measures about 0.9 mm x 0.55 mm.




Figure 10. Circuit design for a DAC module: Three pulse generators are enabled by the H and L signals and gated by two clocks to generate a precise 250-ps pulse with one of 15 selectable current levels (a). Each of the three generators uses a qualifying predriver followed by a series final driver that shares a common tail transistor (b).

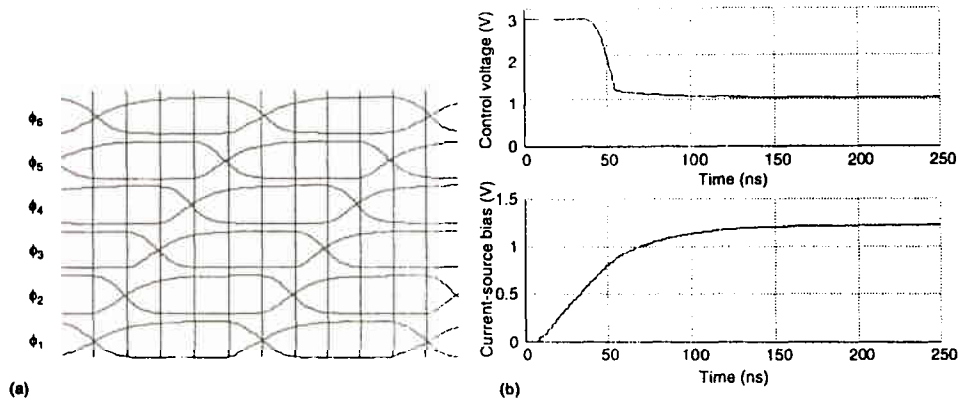


Figure 12. Waveforms from the 10-phase clock generator: the generated clock phases (a) and control voltages during power up (b).

Figure 11b shows waveforms for the transmitter operating with equalization disabled. The transmit waveform shows some attenuation of the high-frequency components due to the driver's slew rate limitations. The bottom waveform in Figure 11b is highly distorted by the high-frequency attenuation of the package parasitics and transmission line. The low-frequency components appear with minimal attenuation (about 600-mV levels), while isolated pulses are severely attenuated (about 300 mV). The result is a signal where several bits are clearly undetectable.

Figure 11c shows differential eye diagrams constructed from the two receiver waveforms. The waveform with equalization (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 12 shows the waveforms from the 10-phase (five-phase complementary) clock generator that controls the transmitter timing. We built the generator 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. Figure 12a shows the clock outputs when the loop is in steady state. For comparison, the vertical lines are spaced at 250-ps intervals. Figure 12b illustrates the dynamics of the loop converging by showing the two signals that control delay during power up. The feedback loop directly drives the current-source bias voltage (top), and a replica-bias circuit³ generates the load control voltage (bottom). As shown, the loop converges to a stable state after less than 250 ns.

Receiver

Figure 13 shows a block diagram of our 4-Gbps receiver.

A demultiplexing receiver samples the differential input stream every 125 ps with sequencing controlled by a 20-phase clock. Each 400-MHz major cycle, the receiver takes 20 samples: 10 data samples D_{0-9} taken from the centers of bit-cells, and 10 edge samples E_{0-9} taken from the boundaries between bit-cells. The receiver inputs the data samples to a funnel shifter that concatenates the current 10 samples with the previous nine samples and then selects a contiguous 10-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 cable delay to be a multiple of 10 bit-cells. The clock control unit uses the edge samples along with the data samples 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 14 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_n being sampled by ϕ_{2n} . The odd clocks sequence the edge samples, with E_n being sampled by ϕ_{2n+1} . To keep loads balanced and lines short, sample D_n is repeated at the end of the line. Each sample is in a separate clock domain. A stage of retiming latches, not shown, aligns all the samples into a single clock domain.

Two timing loops control the 20-phase clock, as illustrated in Figure 15a. The 400-MHz input clock drives a variable delay line with a dynamic range of three bit-cells. This line sets the phase relationship between the 400-MHz input clock and ϕ_0 , as determined by the variable *phase*. The output of this line, ϕ_0 , drives a 10-stage, differential tapped delay line that generates the 20 precisely spaced clock phases. The i th stage generates complementary signals ϕ_i and ϕ_{i+10} . An analog control voltage (set by a phase comparator that aligns ϕ_{10} with ϕ_0) sets the delay of each stage of this line to exactly $1/20$ the period of the input clock. This line is colocated

for a
titled
trick
is lo
signi

The
rat sig
pc qu
pro in
pre rec
ed op
A sys
mit
freq
Weatin
filter
phana
gati
es—
dein
app
le
ter c
para
strip
TI desc
Gbp
rent

400 A

(a)

Figure
adjust

250

250

input
(20-
akes
rs of
anes
to a
with
s 10-
on is
par-
be a
edge
djust
data)

multi-
cked
ples.
e the
loads
end
un. A
uples

strat-
s line
clock
out of
y line
he rth
a anal-
aligns
exact-
scated

with the receive amplifiers, and the loads on all traces are carefully balanced to match delays.

The circuit shown in Figure 15b adjusts the phase control for the first delay line. If there is a transition between the *i*th and (*i*+1)th data samples, D_i and D_{i+1} , signal trans, will be true. On a transition, an analog summing network examines the state of edge sample E_i between these two data samples 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 these two adjacent samples agree and trans, is high, then the clock is early, before the transition. The analog summing network combines and integrates the 10 early signals and the 10 late signals to generate a differential signal, *phase*, that controls the variable delay line.

TRANSMITTER EQUALIZATION EXTENDS the data rates and distances over which we can use electronic digital signaling reliably. 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 are implementing equalization for a 4-Gbps signaling system by building a 4-GHz, 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. We have built the filter in 0.5-micron CMOS circuitry operating at 400 MHz, using a bank of 10 filters and DACs sequenced by a 10-phase 400-MHz clock. We realized narrow drive periods using series gating to combine two clock phases—on and off—in each DAC. To demonstrate the feasibility of this approach, we simulated the extracted layout of the equalized transmitter driving a load through package parasitics and 1 meter of differential strip guide.

The equalizing transmitter we describe is one component of a 4-Gbps signaling system we are currently developing for implementation

in a 0.5-micron 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 4-Gbps 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, 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 4-Gbps serial channel provides adequate bandwidth for most system components, and multiple channels can be grouped in parallel for higher bandwidths.




Figure 13. Receiver block diagram: A demultiplexing receiver sequenced by a 20-phase clock samples the input stream each 125 ps. 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.

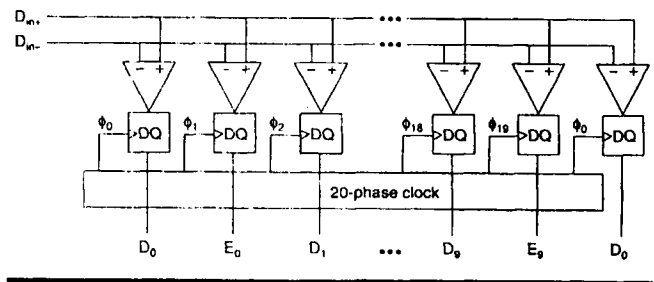


Figure 14. The demultiplexing receiver consists of 21 clocked sense amplifiers sequenced by the 20-phase clock.




Figure 15. Clock generation: two timing loops control the 20-phi clock (a), and a hybrid analog/digital control circuit adjusts the clock phase (b).

A 4-Gbps 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 forty 100-MHz pins. The interface of a 4-Gbyte/s CPU to a level-two cache, for example, could be implemented with just eight serial channels. At the system level, high-speed electrical serial channels can directly replace expensive optical interconnects. Using 18AWG wire, these channels will operate over lengths up to 30 meters, enabling high-bandwidth, low-cost peripheral connections and local-area networks. Inexpensive electrical repeaters can be used for operation over substantially longer distances.

Even with 4-Gbps channels, system bandwidth remains a major problem for system designers. On-chip logic bandwidth (gates \times speed) is increasing at a rate of 90% per year (60% gates and 20% speed). The density and bandwidth of system interconnects increase at a much slower rate—about 20% per year—as they are limited by mechanical factors that are on a slower growth curve than that of semiconductor lithography.² A major challenge for designers is to use scarce system interconnect resources effectively, both through the design of sophisticated signaling systems that use all available wire bandwidth and through system architectures that exploit locality to reduce the demands on this bandwidth. \square

Acknowledgments

This research was supported in part by the Defense Advanced Research Projects Agency 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. The National Science Foundation provided additional support under grant MIP-9306208.

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

References

1. J. Maneatis and M. Horowitz, "Precise Delay Generation Using Coupled Oscillators," *IEEE J. Solid-State Circuits*, Vol. 28, No. 12, pp. 1273-1282.
2. W. Dally and J. Poulton, *Digital Systems Engineering*, Cambridge Univ. Press, Cambridge, UK, 1997.
3. R.E. Matick, *Transmission Lines for Digital and Communication Networks*, McGraw-Hill, New York, 1969.
4. E.A. Lee and D.G. Messerschmitt, *Digital Communication*, second ed., Kluwer, Boston, 1994.
5. M. Tomlinson, "New Automatic Equalizer Employing Modulo Arithmetic," *Electronic Letters*, Mar. 1971.



William J. Dally is a professor of electrical engineering and computer science at the Massachusetts Institute of Technology, where he directs the Concurrent VLSI Architecture Group. This group is currently designing the M Machine, a computer that explores mechanisms for parallel computation. His research interests include computer architecture, parallel computing, computer graphics, computer-aided design, and VLSI design.

Dally received a BS degree in electrical engineering from Virginia Polytechnic Institute, an MS degree in electrical engineering from Stanford University, and a PhD degree in computer science from Caltech.

John Poulton is a research professor in the Department of Computer Science at the University of North Carolina. His projects include several generations of Pixel-Planes and PixelFlow high-performance experimental graphics systems. His research interests are VLSI and system design, high-speed signaling, and special-purpose memory ICs.

Poulton's academic training is in physics, with a BS degree from Virginia Tech, an MS from SUNY-Stony Brook, and a PhD from University of North Carolina, Chapel Hill. He is a senior member of the IEEE.

Direct questions concerning this article to William Dally, MIT AI Laboratory, 545 Technology Sq., Room 620A, Cambridge, MA 02139; billd@ai.mit.edu.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number on the Reader Service Card.

Low 168

Medium 169

High 170

Moving?

Please notify us four weeks in advance

Name (Please print) _____

New Address _____

City _____

State/Country _____

Zip _____

Mail to:
IEEE Computer Society
Circulation Department
PO Box 3014
10662 Los Vaqueros Circle
Los Alamitos, CA 90720-1314

**ATTACH
LABEL
HERE**

- List new address above.
- This notice of address change will apply to all IEEE publications to which you subscribe.
- If you have a question about your subscription, place label here and clip this form to your letter.

vi
yo
Rc
Hs
Kr
Ali
Lu
Fo
Go
An
Ke
Pa
An
Da
Ca
Ka
Da
An
Jeri
Raj
Nor
Elle
Yon
Pau
Kor