CMOS High-Speed I/Os — Present and Future

M.-J. Edward Lee[1], William J. Dally[1,2], Ramin Farjad-Rad[1], Hiok-Tiaq Ng[1], Ramesh Senthinathan[1], John Edmondson[1], and John Poulton[1]

[1]Velio Communications, Inc. Milpitas, CA

[2]Stanford University Stanford, CA

Abstract

High-speed I/O circuits, once used only for PHYs, are now widely used for intra-system signaling as well because of their bandwidth, power, area, and cost advantages. This technology enables chips with over 1 Tb/s of I/O bandwidth today and over 10 Tb/s of bandwidth by 2010 as both signaling rates and number of high-speed I/Os increase with process scaling. Key technologies that enable this growth in I/O performance include low-jitter clock circuits and equalized signaling. An analysis of clock jitter and channel interference suggests that signaling rates should track transistor performance to rates of at least 40 Gb/s over boards, backplanes, and short-distance cables.

1. Introduction

High-speed input/output circuits are becoming increasingly critical as technology scales to increase system bandwidth and decrease power dissipation, die area and system cost. Once used primarily for serial PHYs, high-speed I/O circuits are rapidly becoming the technology of choice for all intra-system connections as well. High-speed I/Os integrated in large numbers enable chips with over 1 Tb/s of I/O bandwidth today. Furthermore, the per-pin bandwidth scales with device speed, at ≈20% per year. As this trend continues, chips with many hundreds of 20 Gb/s I/Os will be feasible by 2010.

High-speed I/Os use *incident-wave signaling* in which a signal is detected on its first traversal of the signal line (the incident wave) and absorbed by a receive termination. This enables the data bandwidth to scale with transistor performance, independent of the length of the line. At high data rates, several bits may be in transit at once — pipelined along the length of the line. In contrast, traditional I/O designs, e.g., LVCMOS, have a bandwidth that is limited by the length of the signal line rather than transistor performance. Without matched terminations, these I/O systems have to *ring-up* the signal wire over several round-trip delays to reliably send one bit. Their data bandwidth is tied to the length of the line, independent of transistor performance, resulting in a bottleneck as bandwidth demands increase.

More than half of the power dissipation of many systems today is I/O power, and the fraction of power due to I/O is increasing. The dynamic power of a logic function scales as α^3 (where gate length scales as α) while a portion of I/O power scales only with α, because a certain amount of current must be delivered to a load that is matched to the line impedance to reliably detect the signal. The minimum current required per I/O is nearly constant, independent of bit rate; thus high-speed I/Os give more bandwidth for this fixed power. Furthermore, the additional power required to build a sophisticated high-speed I/O often scales with α^3, like the core logic. Thus, a better process technology not only enables a higher bandwidth per channel but also reduces the energy consumed per bit.

There are two fundamental challenges to continued scaling of high-speed I/Os: band-limited channels and timing uncertainty. As data rates increase, channel bandwidth becomes limited by the frequency-dependent loss (FDL) of the channel. The distance that a signal can be reliably propagated decreases with the square-root of signal bandwidth for cables (where skin-effect dominates) and linearly with signal bandwidth for circuit boards (where dielectric absorption dominates). Equalization can cancel the frequency-dependent part of the attenuation. However, the magnitude of the attenuation is ultimately a limiting factor. Also, as attenuation levels increase, care must be taken to avoid near-end cross-talk, which is becoming a significant problem in legacy systems.

As signal rates scale, the timing jitter of a high-speed I/O must decrease to remain a constant fraction of a bit time or unit interval (UI). Power supply noise, substrate noise and thermal noise are the most important contributors to clock jitter. Fortunately, our analyses show that by increasing reference clock frequencies and devoting a larger fraction of I/O area to clock circuits, timing jitter can be made to scale with bit time. Overall, it appears that there are no major obstacles to achieving 40 Gb/s signaling rates over boards, backplanes, and short-distance cables (tens of meters). Hence signaling rates should continue to scale with transistor performance to at least this speed.

The remainder of this paper describes high-speed I/O cir-

Proceedings of the 21st International Conference on Computer Design (ICCD'03)
1063-6404/03 $ 17.00 © 2003 IEEE

cuits in more detail. Section 2 describes the architecture of a typical high-speed I/O and the details of some of its components. Section 3 discusses the current state-of-the-art in high-speed I/O technology and the future challenges posed by channel attenuation and clock jitter.

2. A Typical High-Speed I/O

2.1. Top Level Architecture

Figure 1 shows a typical high-speed I/O[1]. The transmitter converts N-bits of parallel data from the *core logic* into a two-bit stream, and then a 2:1 multiplexer gates out two symbols per clock cycle with precise timing. The retimer ensures the data are positioned correctly for multiplexing. A higher multiplexing ratio can be implemented with more clock phases to further reduce the frequency requirement. However, in multiplexed systems any phase mismatch between different clock phases results in deterministic jitter in the serialized data. To avoid this, the data can be retimed with a full bit-rate clock before the final output driver, at the expense of higher power consumption and a lower data rate [2]. A pseudo-random bit sequence (PRBS) generator is usually built in for at-speed testing.

A bank of samplers in the receiver samples the bit stream on evenly spaced clock phases to demultiplex the data directly, easing the frequency requirement. This multiphase approach suffers the same deterministic jitter problem as its counterpart at the transmitter. The clock recovery unit adjusts the clock phase to place the data samples in the middle of the bit cell. The adjustment is performed by sampling both the center (the samplers labeled *C*) and edge (the samplers labeled *E*) of each bit cell. On a transition, the value of the edge sample determines if the sampling clock is *early* or *late*[2]. The 2-bit data from the samplers are deserialized into an N-bit parallel data suitable for the digital logics.

2.2. Bandwidth Limitations

The bandwidth achievable by a signaling system is limited by attenuation, interference, and jitter. These factors are illustrated in the conceptual *eye diagram* of Figure 2. Constructed by *folding* the data waveform into a symbol time, an eye diagram shows variations of signal amplitude (voltage noise) and timing (jitter) across bit cells. The rectangle in the middle represents the *eye opening*, which must be wider than the receiver jitter plus aperture[3] and taller than the receive sensitivity. That is, $t_b \geq t_r + t_a + t_{u,all}$, where t_b is the bit time, t_r is the rise time, t_a is the receiver aperture, and $t_{u,all}$ is the total timing uncertainty of the system [4]. For most systems, the dominant component is $t_{u,all}$, caused mostly by clock jitter, intersymbol interference (ISI) and cross-talk.

Input offset is often the largest component of the receiver sensitivity. As shown in Figure 3, a digital calibration scheme can cancel this offset with digitally trimmed current sources trained at startup [6] [11]. With this method, the offset is reduced from > 100 mV to < 10 mV.

2.3. Clock Multiplier

A clock multiplier multiplies the reference clock up to the multiplexing rate. Two common implementations are a phase-locked loop (PLL), shown in Figure 4, and a multiplying delay-locked loop (MDLL), shown in Figure 5. In a PLL, the bandwidth of the feedback loop should be high to reject the oscillator jitter since a low-jitter reference clock is often provided (e.g., from a crystal). In practice, however, the bandwidth of a PLL is limited to about 5% of the reference clock frequency due to the delay around the loop [9] [14]. In contrast, an MDLL, shown in Figure 5, periodically injects the clean reference clock into the oscillator to *reset* the phase error every reference clock cycle [5] [12]. In this implementation, the pulser generates an enable pulse for the multiplexer (so that the clean reference clock can be muxed in) and the phase detector (so that only one oscillator edge per reference clock cycle is compared).

Figure 6 illustrates the response of a PLL and a MDLL to a frequency shift (or a phase ramp), which is a common test for the jitter performance of clock generation circuits since supply noise exhibits a similar behavior. Jitter in a PLL accumulates until the loop is able to respond. Both the peak jitter amplitude and the time it occurs are approximately inversely proportional to the loop bandwidth. In contrast, the clean reference clock resets the jitter in a MDLL every reference clock cycle. The peak of the sawtooth decays as the loop gradually corrects the frequency offset. It can be shown that even with the upper bandwidth limit, a PLL exhibits more than twice the peak jitter amplitude compared with a MDLL, which does not require a high loop bandwidth to achieve low jitter.

2.4. Clock Recovery

The clock recovery block determines where to position the sampling clocks. A PLL locked to the receiver input is often used for this function. Unlike the clock multiplier at

1 In this example, a signal can only go in one direction on a channel. Simultaneous bidirectional signaling allows signals to flow in both directions on one channel but will not be discussed in this paper since it is rarely encountered.

2 This type of phase detector, called a *bang-bang* or *Alexander* phase detector, is popular due to its ease of implementation [1].

3 Receiver aperture is the time the signal must be above the receiver sensitivity to make a correct decision.

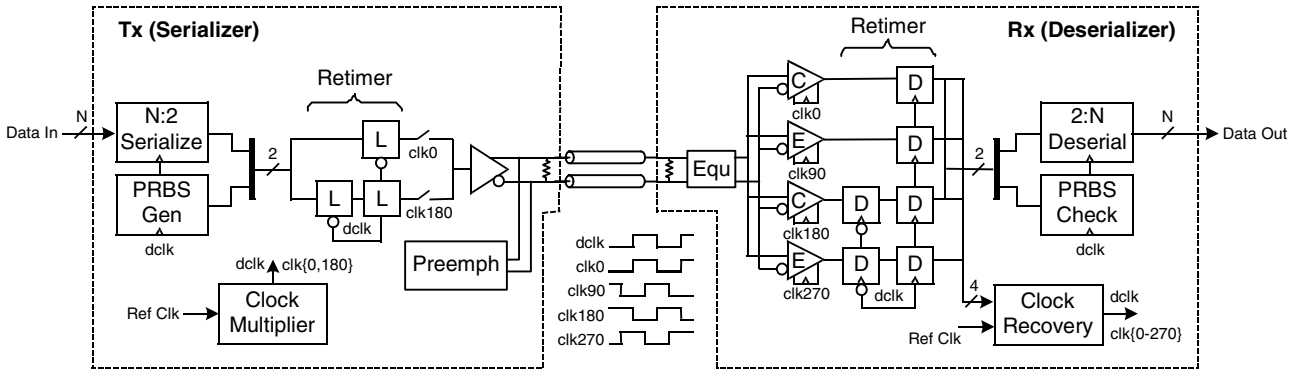

Figure 1: A typical high-speed I/O architecture.

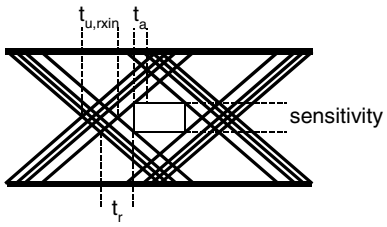

Figure 2: A conceptual eye diagram.

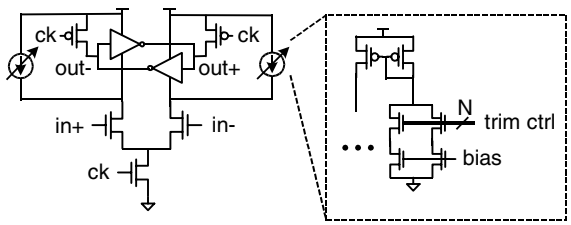

Figure 3: A digitally-trimmed input sampler.

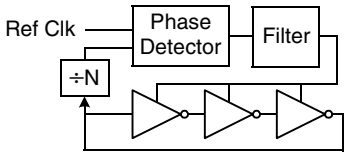

Figure 4: A phase-locked loop.

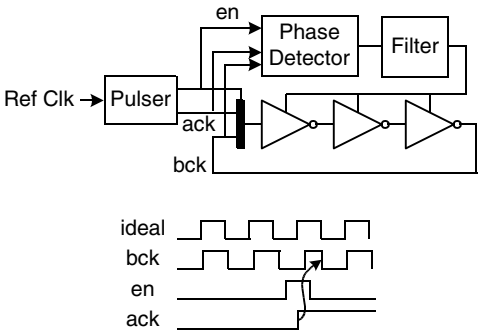

Figure 5: A multiplying delay-locked loop.

the transmitter, the high-jitter receiver input necessitates a low loop bandwidth, which is in conflict with oscillator jitter rejection. A *dual-loop* approach, in which a high-bandwidth clock multiplier is used to multiply a low-jitter reference clock and a separate low-bandwidth loop is used for receiver input tracking, removes this tradeoff [19].

Most dual-loop systems use a first-order receiver tracking loop, as shown in the first dashed box in Figure 7. The binary *early/late* indications from the phase detector are passed through a phase filter to reduce noise due to input jitter. The filter output controls the phase of the sampling clocks through a *timing vernier* that changes the phase of the clock multiplier output. With a plesiochronous input, this results in either phase lag, if the loop is too slow to track the input, or phase wander, if the loop is too fast to filter the input jitter, or both. A second-order receiver tracking loop eliminates these phase errors by estimating the frequency of the input signal. The frequency tracking loop, shown in the second dashed box in Figure 7, integrates the output of the phase filter to estimate the frequency of the received signal and sends a stream of up/dn signals to compensate for any offset from the reference clock frequency. This enables a slow loop to be used to filter input jitter without causing phase lag [13]. The advantage of using a digital implementation is that many loop parameters, such as the length of the phase filter and the frequency filter, can be made pro-

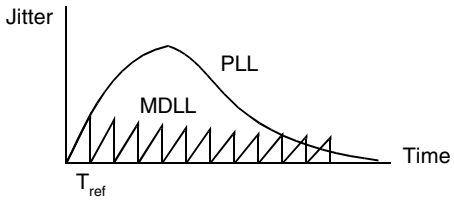

Figure 6: Jitter Response of a PLL and a MDLL.

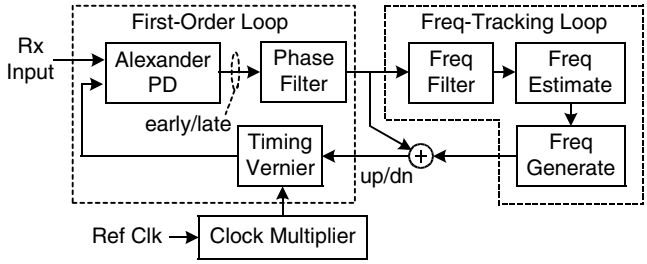

Figure 7: A digital clock recovery architecture.

grammable (e.g., to maximize jitter filtering or minimize lock time). Furthermore, the digital control to the timing vernier can be easily bypassed to allow flexible positioning of the sampling clocks for testing purposes.

2.5. Equalization

Skin effect, dielectric absorption and discontinuities cause a channel to exhibit *frequency-dependent loss* (FDL). A pulse representing a bit not only gets *attenuated* by the channel but is *spread out* in time, causing ISI. Figure 8 shows the frequency response and the 6.25 Gb/s pulse response of a backplane channel. A significant amount of ISI is present at the adjacent sample points in the pulse response (the vertical grid lines are spaced at the sample points, 160 ps apart).

A filter, or equalizer, with an inverse channel response can be used to counteract FDL. A commonly used filter is a discrete-time symbol-spaced FIR filter. Oftentimes it is implemented at the transmitter (transmitter pre-emphasis) with direct current summing of different taps at the output [3] [11]. Figure 8 shows the effect of a 4-tap filter (1 main tap and 3 post-cursor taps) in the frequency domain and time domain on the same backplane channel. Since a portion of the available transmitter current is assigned to the equalization taps, in effect transmitter pre-emphasis *attenuates* the low-frequency component to achieve a flat spectrum overall. With pre-emphasis, the amount of ISI is significantly

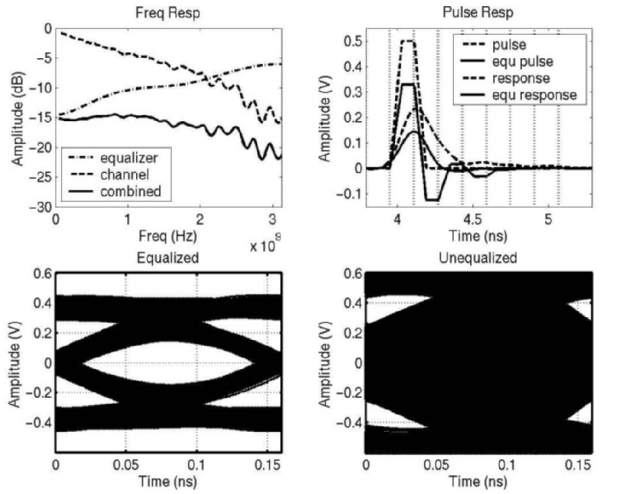

Figure 8: Response of a backplane channel.

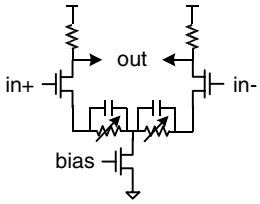

Figure 9: Active RC-based receive equalization.

reduced. As shown in Figure 8, it opens up a completely closed eye (PRBS 23 pattern).

Sometimes it is beneficial to place the equalizer at the receiver. Although a discrete-time FIR approach can be used, it is significantly more complicated than transmitter pre-emphasis since high-speed sampling, multiplication, and addition of analog values are required. An alternative is an active high-pass filter, shown in Figure 9 [6]. The gain of this circuit goes up with frequency as the capacitor decreases the amount of source degeneration. The equalization gain can be adjusted through the variable resistors.

3. Future Challenges

As gate lengths are scaled by α (at a rate of about 20% per year), gate delay also scales as α and transistor ω_T scales as $1/\alpha$. Signaling bandwidth can also scale as $1/\alpha$ if the timing uncertainties, dominated by clock jitter and channel interference can be made to scale at the same rate. This section investigates the scalability of clock jitter and dis-

cusses how channel interference can be improved through circuit level and system level techniques. With careful circuit and system design, we expect the bandwidth of electrical signals on boards, over backplanes, and over cables to scale to at least 40 Gb/s. I/O energy per bit is expected to scale as α to α^2 in the near future, but will eventually be limited by α. In contrast, the switching energy per function for digital logic scales as α^3. As a result, the fraction of I/O power in a system will increase for the foreseeable future.

3.1. Scalability of Clock Jitter

Analysis of a CMOS inverter ring oscillator suggests that clock jitter can be made to scale with α if higher reference clock frequencies are used and if an increasing percentage of I/O area and power is devoted to clock generation. We investigate the effects of the three most important noise sources: power supply noise, substrate noise, and thermal noise.

3.1.1. Power Supply Noise A $k\%$ change in supply voltage results in a $k\%$ change in the period of a CMOS ring oscillator. Assuming that supply noise remains a constant fraction of the supply, if the reference clock frequency remains constant, the p-p jitter will remain constant since both the rate and the duration of jitter accumulation are fixed. In other words, jitter as a percentage of the bit time increases. To ameliorate this problem, we can increase the supply noise rejection and/or increase the reference clock frequency.

Local supply regulation, shown in Figure 10, is commonly used to isolate critical circuits [11]. On-chip digital switching often generates significant supply noise. To first-order approximation, the amount of noise rejection by this type of regulator is proportional to C_1/C_2. Therefore, supply rejection can always be improved with area. It also improves with process scaling as long as the *area* of C_1 scales slower than α^2. For multiphase oscillators, however, the area of the delay element often needs to remain constant to keep phase mismatch a fixed fraction of the bit time. In this case, the area of C_1 must *increase* with process scaling to improve the supply rejection. A bit-rate oscillator is advantageous in this regard since it does not rely on matching of the delay stages to produce precise clock phases.

The frequency of the crystal reference is limited by its thickness and cannot be expected to scale as aggressively as the semiconductor technology. Since on-chip LC oscillators exhibit a much better jitter performance, it is advantageous to multiply the reference clock to an intermediate frequency with a global on-chip LC oscillator and use local ring oscillators to generate the final high-frequency clocks whenever integration or tunability is a concern. Fortunately,

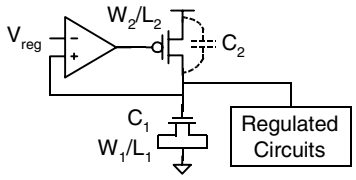

Figure 10: A local supply regulator.

the Q of on-chip inductors is improving with the availability of more metal layers in advanced CMOS processes.

With a combination of higher reference clock frequencies and better supply noise rejection, jitter induced by power supply noise should continue to scale with the bit time in the foreseeable future.

3.1.2. Substrate Noise, also caused mostly by digital switching, is a major concern in highly integrated applications. Fortunately, process remedies are now readily available to reduce its effect. For example, many processes now offer deep NWELL to isolate a subcircuit from the rest of the chip. Recent work has demonstrated better than 50 dB attenuation of substrate noise with only $200\mu m$ of separation in an epi process [8]. Judicious use of this structure should keep substrate noise a negligible effect on sensitive circuits such as clock generators.

3.1.3. Thermal Noise Unlike supply and substrate noise whose magnitude can be attenuated externally, thermal noise is inherent in the device[4]. The rms jitter of an N-stage CMOS ring oscillator when placed in a PLL or a MDLL is [10] [15]

$$\sigma_{jit} = \frac{\Gamma_{rms}}{2\pi f_0 C V_{sw}} \sqrt{\frac{N}{2} \frac{\overline{i_n^2}}{\Delta f} \tau_L} \qquad (1)$$

where f_0 is the oscillator frequency. For a PLL, τ_L is $1/2\pi f_L$, where f_L is the loop bandwidth. For a MDLL, τ_L is $1/f_{ref}$, where f_{ref} is the reference clock frequency. Γ is the impulse sensitivity function (ISF) and determines the sensitivity of the oscillator to a noise impulse [10]. For example, noise occurring at the edge of the clock produces more jitter than that at the peaks. It can be shown that Γ_{rms} scales as $\alpha^{1.5}$ due to sharper edges at higher frequencies. CV_{sw} is the maximum charge swing and determines how easily the oscillator nodes can be moved. It scales as α^2. $\overline{i_n^2}/\Delta f$ is the amount of thermal noise on one node and remains approximately the same with scaling[5]. This analysis

4 Although MOSFETs suffer from poor $1/f$ noise, its effect on jitter can be reduced significantly by balancing the rise and fall time of the clock. Furthermore, low-frequency jitter such as $1/f$ noise is attenuated when placed in a high-bandwidth loop.

indicates that while the clock period scales as α, the rms jitter scales as $\sqrt{\alpha}$ for a fixed reference clock frequency and as α if the reference clock frequency scales at the same time. Furthermore, increasing the width of the delay element improves jitter in a square root fashion due to a higher charge swing.

It is instructive to compare the magnitude of jitter induced by thermal noise and that induced by supply noise. Recent measurement of a $0.25\mu m$ 1.33 GHz CMOS ring oscillator showed a thermal-noise-induced phase noise of -111.5 dBc/Hz at a 1 MHz offset from the carrier [10]. For a MDLL with a multiplication factor of 10, this roughly translates into a rms jitter of 0.173 ps at the end of a reference clock period. The p-p jitter for $< 10^{-15}$ probability is 2.77 ps. In contrast, a 5% supply noise with a 20 dB power supply rejection results in roughly 37 ps p-p jitter.

In summary, by increasing the reference clock frequency and increasing the oscillator width, thermal-noise-induced jitter should scale well with the bit time. In addition, in highly integrated applications, thermal noise will likely remain a negligible effect for the foreseeable future.

3.2. Channel

High-speed I/Os are typically used between chips on a printed circuit board, across a connectorized backplane, and across short distance cables (tens of meters). The FDL (in dB) scales linearly with bandwidth for typical circuit boards, where dielectric absorption dominates, and as the square-root of bandwidth for cables, where skin effect dominates. In addition, discontinuities can cause significant FDL beyond these fundamental loss mechanisms. While equalization can flatten the spectrum of these channels, total attenuation as well as external interferences will ultimately limit the achievable bit rate. In this section we focus on backplane channels because they are the most challenging in terms of attenuation and cross-talk.

For many systems (e.g., switches and routers), bandwidth is upgraded through gradual replacement of cards in an existing backplane. These *legacy* backplanes, suitable for the speed requirement at the time they are designed, often exhibit very high attenuation and very low signal-to-interference ratio (SIR) as bit rate increases. Figure 11 shows the cumulative distribution functions (CDFs) of ISI, 8 cross-talk aggressors, and the total (ISI plus cross-talk) for one such backplane (same channel as Figure 8) running at 6.25 Gb/s. The right side of the plot stops at the amplitude of the received pulse. Therefore, the probability of a bit error due to a particular interference is the intersection of the

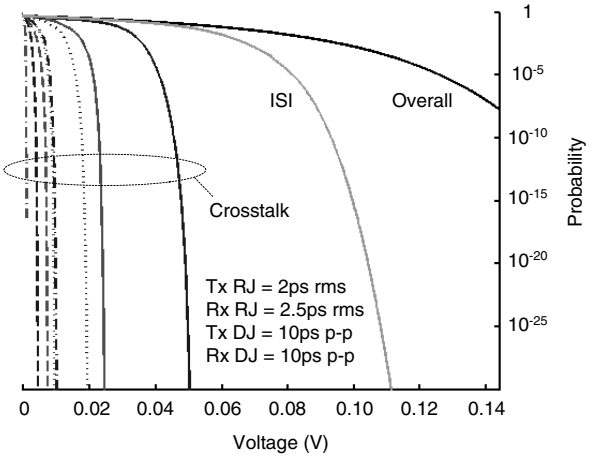

Figure 11: CDFs of channel interferences for a backplane, including the effects of clock random jitter (RJ) and deterministic jitter (DJ).

curve with the y-axis. Although a 4-tap equalizer is used, the bit-error-rate (BER) is still unacceptably high at 10^{-7}.

The decreasing signal-to-inteference ratio is best managed through a combination of circuit level and system level improvements. Currently, most high-speed I/Os use a 2-tap linear filter that is manually adjusted through either trial-and-error or channel analysis. As longer filters are required to further remove the ISI, *adaptive equalization*, in which the tap coefficients are optimized by hardware, becomes an a critical requirement [20]. It not only obviates the need for user intervention that is often time-consuming but also improves the effectiveness of equalization by including the effects of package and termination non-idealities that are lost in s-parameter or eye measurements. A non-linear filter, such as a decision-feedback equalizer (DFE), can further improve the margin by equalizing the signal without amplifying the cross-talk. In contrast, a high-pass linear filter commonly used to equalize the channel amplifies the high-pass cross-talk significantly.

Because the channel response attenuates while the cross-talk amplifies at high-frequencies, sending more bits per unit bandwidth through multi-level signaling is an attractive way to manage this problem [7]. Figure 12 compares binary and 4-level eye diagrams for the same *symbol rate*. The horizontal eye opening of 4-level signaling is less than binary signaling due to limited slew rate. Furthermore, its vertical eye opening is *less* than $1/3$ that of binary signaling due to the voltage noise at the intermediate level (V_n). Multi-level signaling often requires additional overhead bandwidth to ensure enough useful transitions[6] are present for clock re-

5 It has been observed that thermal noise increases in deep-submicron technologies due to high-field carrier heating. Fortunately, voltage scaling helps reduce this effect.

covery. The exact benefit of multi-level signaling needs to be simulated on a per-channel basis, performing an analysis similar to that shown in Figure 11. However, a useful rule-of-thumb is that the SIR must increase by at least 12 dB in the octave from $1/4$ to $1/2$ the bit rate for 4-level signaling to be advantageous.

Careful system design is needed in addition to circuit level innovations to sustain continuous bandwidth scaling. For example, via stubs often cause FDL to be much worse than expected from skin effect and dielectric loss due to quarter wavelength resonance. *Back-drilling*, in which the unused portion of the via is removed, provides a very cost-effective way to push out this resonance [18]. Without back-drilling, a 180 mil thick FR4 backplane via stub creates a resonance at about 5 GHz for typical via sizes.

The primary source of cross-talk in most systems is the backplane connector. New connectors are being introduced with ground shields completely surrounding each signal pair to reduce cross-talk. Signals flowing in opposite directions are isolated from each other to avoid near-end cross-talk, which is much more detrimental than far-end cross-talk since the interference is not attenuated by the full length of the channel along with the signal. Cross-talk coupling less than -50 dB has been demonstrated on a typical backplane with these improvements [18] [17].

With a 50 mV receiver sensitivity now available in commercial high-speed I/Os, 26 dB of FDL at 1/2 bit rate can be tolerated for a typical 1 V p-p input. Using the techniques mentioned above, along with low-loss laminates, $<$ 20 dB of FDL up to 10 GHz has been demonstrated on fully connectorized backplane channels up to 70cm. 10 Gb/s data transmission without any equalization has been demonstrated, and 20 Gb/s data transmission with simple 2-tap pre-emphasis is now possible [17] [16]. With further investment, it appears that achieving $<$ 30 dB FDL up to 20 GHz and one meter is not out of reach. This, combined with further process and circuit improvements on receiver sensitivity, jitter, and equalization, should enable a 40 Gb/s transceiver over backplanes in the future. Of course, these benefits cannot be fully realized unless the whole system, including the backplane, is completely upgraded.

3.3. Current State-of-the-Art and Future Trend

Figure 13 shows that the bandwidth of production backplane channels has been doubling every two years since 1999. 3.125 Gb/s channels are now commonplace and 6.25 Gb/s and 10 Gb/s channels have been demonstrated [6] [21]. It is clear that this bandwidth growth trend is not sustainable since device speed is only doubling every 3–4 years. Techniques such as multi-level signaling only provide a one-time bandwidth increase. Since 1999, I/O technology has been *catching up* to the semiconductor technology, making the *super-Moore's-Law* bandwidth trend possible. A practical limit of the symbol time for high-speed I/Os is about 2 FO4 (fanout of 4 inverter delay). In $0.13\mu m$ CMOS technologies, this limit is about 7 Gb/s (or 12 Gb/s for 4-level signaling). It is expected that the per-channel backplane bandwidth growth will be limited by semiconductor scaling beyond 10 Gb/s and at least up to 40 Gb/s when the channel imperfections become the critical bottleneck.

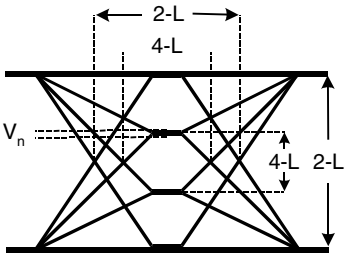

Figure 12: A comparison of binary and 4-level eye diagrams.

High-speed I/O energy per bit will ultimately be limited by the transmitter output drive, which requires at least a constant current to overcome fixed noise and higher loss in the channel. As a result, transmit energy per bit scales as α. For a CMOS inverter based multi-phase clock multiplier, energy per bit also scales as α since gate capacitance must increase to scale transistor mismatch. For a bit-rate oscillator, where matching is less of an issue, energy per bit scales as α^3. The rest of the circuits, including the transceiver data paths and the digital clock recovery unit, are digital logic and hence scale as α^3. As a result, the energy per bit for one high-speed I/O is expected to scale as α to α^2 in the near future, but will eventually be limited by α. In comparison, the switching energy for a digital logic function scales as α^3. This α^2 difference in scaling is partly offset by increased integration. As $G = 1/\alpha^3$ more core logic bandwidth is integrated on a chip (holding total core power constant), Rent's rule suggests that only $I = G^{2/3} = 1/\alpha^2$ more I/O bandwidth will be required, consuming $1/\alpha$ times as much I/O power. The ratio of I/O power to core power on a chip will hence increase by $1/\alpha$ with technology scaling.

4. Conclusion

High-speed I/O circuits enable signaling rates to scale with device technology, largely independent of transmission

6 Only transitions from outer levels to outer levels and inner levels to inner levels are useful for clock recovery.

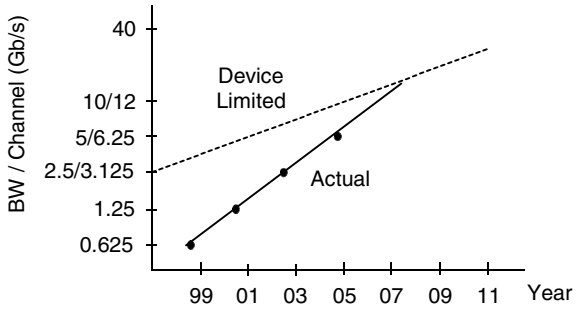

Figure 13: Bandwidth of production backplane channels using CMOS technologies.

length. However, for this scaling to continue, care must be taken to ensure that timing uncertainty, due to clock jitter and channel interference, scales at the same rate. Our analysis shows that by using higher frequency reference clocks and devoting a larger fraction of area to clock circuits, jitter due to power supply noise, substrate noise, and thermal noise will scale appropriately. Channel interference can be managed at the circuit level and the system level simultaneously to push the backplane bandwidth to at least 40 Gb/s in the future.

Although the bandwidth of production backplane channels has been doubling every two years for the past three generations, this trend is not expected to be sustainable since transistor speed is only doubling every 3–4 years. It is expected that the per-channel backplane bandwidth growth will be limited by semiconductor scaling beyond 10 Gb/s.

Finally, I/O will continue to dominate the energy consumption of the system. The ratio of I/O power to the logic power on a chip will increase as $1/\alpha$ with technology scaling as the I/O bandwidth is scaled with the logic bandwidth according to Rent's rule.

Acknowledgment

The authors would like to thank H. Yazdanmehr for his support and R. Rathi, A. Nguyen, J. Tran, F. Heaton, T. Greer, T. Stone, S. Gowni, Y. Tan and E. Dinis-Mcveigh for their contributions to the high-speed I/O work at Velio.

References

[1] J. D. H. Alexander. Clock recovery from random binary data. *Electronics Letters*, 11:541–542, Oct. 1975.

[2] J. Cao et al. OC-192 transmitter and receiver in standard $0.18\mu m$ CMOS. *IEEE J. Solid-State Circuits*, 37:1768–1779, Dec. 2002.

[3] W. J. Dally and J. Poulton. Transmitter equalization for 4Gb/s signaling. In *Proc. Hot Interconnect*, pages 29–39, 1996.

[4] W. J. Dally and J. W. Poulton. *Digital Systems Engineering*. Cambridge University Press, 1998.

[5] R. Farjad-Rad et al. A low-power multiplying DLL for low-jitter multigigahertz clock generation in highly integrated digital chips. *IEEE J. Solid-State Circuits*, 37:1804–1812, Dec. 2002.

[6] R. Farjad-Rad et al. 0.622-8Gpbs 150mW serial IO macrocell with fully flexible preemphasis and equalization. In *Symposium on VLSI Circuits Dig. Tech. Papers*, pages 63–66, 2003.

[7] R. Farjad-Rad, C.-K. K. Yang, and M. Horowitz. A $0.3\mu m$ CMOS 8-Gb/s 4-PAM serial link transceiver. *IEEE J. Solid-State Circuits*, 35:757–764, May 2000.

[8] L. M. Franca-Neto et al. Enabling high-performance mixed-signal system-on-a-chip (SoC) in high performance logic CMOS technology. In *Symposium on VLSI Circuits Dig. Tech. Papers*, pages 164–167, 2002.

[9] F. M. Gardner. Charge-pump phase-locked loops. *IEEE Trans. Comm.*, COM-28:1849–1858, Nov. 1980.

[10] A. Hajimiri and T. H. Lee. *The Design of Low Noise Oscillators*. Kluwer Academic Publishers, 1999.

[11] M.-J. E. Lee, W. J. Dally, and P. Chiang. Low-power, area-efficient, high-speed serial I/O circuit techniques. *IEEE J. Solid-State Circuits*, 35:1591–1599, Nov. 2000.

[12] M.-J. E. Lee et al. Jitter transfer characteristics of delay-locked loops – theories and design techniques. *IEEE J. Solid-State Circuits*, 38:614–621, Apr. 2003.

[13] M.-J. E. Lee et al. A second-order semi-digital clock recovery circuit based on injection locking. In *ISSCC Dig. Tech. Papers*, pages 74–75, 2003.

[14] T. H. Lee. *The Design of CMOS Radio-Frequency Integrated Circuits*. Cambridge University Press, 1998.

[15] J. A. McNeill. Jitter in ring oscillators. *IEEE J. Solid-State Circuits*, 32:870–879, June 1997.

[16] H. Merkelo and T. Hochberg. Late breaking news and advances in new limits for ultrahigh data rates on copper. In *Proc. DesignCon*, pages 190–202, 2003.

[17] R. Moedinger, J. Lappoehn, and H. Merkelo. New backplane connector technology for 10 Gb/s applications. In *Proc. DesignCon*, pages 226–238, 2003.

[18] M. Shafer, B. Das, and G. Patel. Connector and chip vendors unite to produce a high-performance 10 Gb/s NRZ-capable serial backplane. In *Proc. DesignCon*, pages 253–271, 2003.

[19] S. Sidiropoulos and M. Horowitz. A semidigital dual delay-locked loop. *IEEE J. Solid-State Circuits*, 32:1683–1692, Nov. 1997.

[20] J. T. Stonick et al. An adaptive PAM-4 5-Gb/s backplane transceiver in $0.25\mu m$ CMOS. *IEEE J. Solid-State Circuits*, 38:436–443, Mar. 2003.

[21] J. Zerbe et al. Equalization and clock recovery for a 2.5-10Gb/s 2-PAM/4-PAM backplane transceiver cell. In *ISSCC Dig. Tech. Papers*, pages 80–81, 2003.