

A 2Gb/s/pin CMOS Asymmetric Serial Link*

Kun-Yung Ken Chang, William Ellersick, Shang-Tse Chuang, Stefanos Sidiropoulos, and Mark Horowitz
Computer System Laboratory, Stanford University
Stanford, CA 94305

ABSTRACT

The design of an asymmetric serial link poses a number of trade-offs for the designer. This paper describes measurements from a 0.25 μ m CMOS test chip which show that a properly designed asymmetric link can achieve 2Gb/s using single-ended signalling with a bit-error rate $< 10^{-14}$.

Architecture

Using high-speed serial links in network switches can provide significant performance improvements [1]. However, conventional serial links create hot-spots in large scale switches. For example, employing a conventional serial link at each input of a 32×32 crossbar chip overly increases its power dissipation due to the 32 PLLs required. As shown in Figure 1, asymmetric links can reduce the power consumed in the crossbar chips, by adjusting the timing of the crossbar inbound links at the transmitters located at the peripheral chips, [2].

Figure 2 shows the block diagram of the asymmetric serial link. The *smart end* of the link adjusts the phase of both its transmitter and receiver clocks, while the *dumb end* keeps the phase of both its clocks constant. The timing information, required by the smart end to adjust its transmitter clock, is periodically sent over the dumb-to-smart link.

Design Issues for the Asymmetric Serial Link

To help understand the capabilities of these links, we built a test-chip using Texas Instruments' 0.25 μ m CMOS technology. The chip has a single PLL, and a transmitter-receiver pair that can be configured as either the smart or dumb end of an asymmetric link. The four digitally-controlled precision phase adjusters in the PLL allow us to measure the timing and voltage margins of the system during operation. The test chip helped us answer the following questions:

- (1) Asymmetric links require explicit calibration of the smart-to-dumb link, so the designer needs to trade-off effective link bandwidth against phase adjust bandwidth. What calibration frequency is required for good timing margins?
- (2) Since the link cannot transfer data during calibration, the same data receiver could be used to acquire the timing information. Does using the same receiver improve timing margins, and by how much?
- (3) What is the minimum necessary signal amplitude? Can a single-ended link achieve the same performance as a differential link?

Implementation

The asymmetric link transmits data on both edges of its internal clock, and uses a pair of data receivers to recover alternate data bits. A separate pair of timing receivers is employed at each input to recover timing information, as shown in Figure 3. Noise-tolerant, integrating receivers [5] are used to filter out high-frequency amplitude noise, especially reference voltage noise in the single-ended link. The smart end adjusts both its receiver and transmitter clocks, so that the receivers at both ends of the link integrate the incoming data for the entire bit time. The timing receiver clock (RTClk) is shifted by $1/2$ bit time (90°) from the data receiver clock (RClk), to integrate around the data transition to recover timing information.

The clock generation circuitry at the smart end of the links (Figure 4) includes a dual-loop PLL with digitally-controlled phase adjusters [4]. The dashed line encloses the core loop which comprises a VCO-based third-order PLL with a 250MHz reference clock input. The 1GHz VCO clock is digitally divided by four to generate a 250MHz clock, which is phase locked to the reference clock. The six-

stage VCO uses differential elements with replica-biased symmetric loads [3] to generate twelve clock phases, spaced by 30° . Two adjacent phases are selected by clock multiplexors, which form the first stage of the phase adjusters. A digitally-controlled interpolator mixes the two adjacent phases to generate one of 16 finer phases. As a result, each VCO clock cycle can be subdivided into $12 \times 16 = 192$ phases, which results in a phase resolution of 5 ps for a 1GHz clock.

The timing information from each link is processed by a logic block which controls the phase adjusters and in this way closes the outer loop of the dual-loop PLL. Four phase adjusters respectively generate the digital divider clock, the transmitter clock (TClk), the data receiver clock (RClk) and the timing receiver clock (RTClk). However, in normal operation only the transmitter clock phase and the receiver clock phase are adjusted independently. The phase setting for the digital divider clock is fixed, while the phase setting for the timing receiver clock is offset by $1/2$ bit time (90°) relative to the data receiver clock.

The phase setting for the dumb-to-smart link can be adjusted continuously while the smart-to-dumb link can only be adjusted periodically. The smart transmitter clock is adjusted based on a majority vote of several bit times of dumb receiver timing information to filter out high-frequency noise. Furthermore, to minimize dither jitter, the phase adjust logic moves the clocks only after the results of its last decision have propagated through the outer loop.

The clock generation circuit is carefully tuned so that all the output clock paths have approximately the same delay. As shown in Figure 4, both the receiver clock paths have an extra delay circuit to match the clock-to-output delay of the digital divider in the PLL. The digital divider is sized so that its clock-to-output delay matches that of the transmitter, so no extra delay circuit is needed in the transmitter clock path. As a result, the core PLL loop continuously tracks temperature and supply-induced delay variations within its loop bandwidth. Therefore, the transmitter output and the receiver sample time are kept locked in a fixed relationship to the reference clock.

Measurement Setup and Results

To measure the timing margins at the receiver output accurately, the user can externally control the phase of each of the internal clocks. For example, one can externally control the phase error between the data receiver clock and the timing receiver clock while the link is in lock until bit errors are detected.

Two test chips were used to measure performance, one configured as the dumb end, and the other as the smart end of the link. Figure 5 shows the 2Gb/s measured waveform at the transmitter outputs of both the smart-to-dumb and the dumb-to-smart links using 650mV signalling, with 100mV induced PLL supply noise. The measured bit-error rates for both links are less than 10^{-14} . This confirms that the smart-to-dumb link (periodic phase adjustment) can run as fast as the dumb-to-smart link (continuous phase adjustment), despite the significantly different outer loop bandwidth. Moreover, since the transmitter output and receiver sample time are phase locked by the PLL, the asymmetric link can still run at-speed without periodic calibration, achieving the same performance (less than 10^{-14} bit-error rate) with only an initial calibration process.

Figure 6 shows the measured timing margins of the link at 2Gb/s, with and without PLL supply noise injected at the dumb-end chip. Due to the receiver offset and clock path mismatch between the data and timing receivers, the phase adjustment of both links are slightly off-center. The last row of Figure 6 shows the timing margin of a smart

* Supported by Texas Instruments

receiver experiment with no explicit timing receivers. External phase control was used to shift the data receiver clock by 1/2 bit time (90°) to acquire timing information, reducing the sample time offset from 23.4ps to 5ps.

To measure the minimum required signal amplitudes, we increased the transmitted signal amplitudes until the bit error rates were less than 10^{-12} . Our measurements showed that at 2Gb/s with 100mV induced PLL supply noise, the minimum required differential signal amplitude was 150mV, while the minimum single-ended signal amplitude was 600mV. The former amplitude is smaller because the differential link has twice the effective signal amplitude, and the single-ended link has to overcome reference voltage noise and other coupling noise to the signal. To fully verify robustness, the smart-to-dumb link was reconfigured to use 650mV single-ended signalling with the same PLL supply noise, and achieved less than a 10^{-14} bit-error rate. This shows that with sufficiently large signal swings, single-ended signalling is practical at 2Gb/s.

Table 1 summarizes the performance of the asymmetric link. The PLL and interpolators consume 175mW at 2Gb/s, so 32 of them would consume 5.6W. However, the asymmetric link would consume more power in distributing clocks over the entire crossbar chip because of the additional balanced clock distribution tree. Therefore, we estimate that the asymmetric link reduces the power of a 32×32 crossbar by 4.5W. These experimental results demonstrate that a properly designed asymmetric link can transfer data at 2Gb/s per pin while offering significant reductions in the power and complexity of the crossbar chip.

Acknowledgments

We would like to thank Nick McKeown, Ken Yang, Jeff Hsieh from Stanford and Martin Izzard, Ah-Lyan Yee from Texas Instruments for support and assistance.

References

- [1] Nick McKeown, et al. "Tiny Tera: A Packet Switch Core," IEEE Micro, Jan/Feb. 1997.
- [2] Ken K. -Y. Chang, et al. "A 2Gb/s Asymmetric Serial Link for High-Bandwidth Packet Switches," Hot Interconnects '97, pp171-179.
- [3] J. Maneatis, "Low-Jitter Process-Independent DLL and PLL Based on Self-Biased Techniques," IEEE JSSCC, Nov. 1996.
- [4] S. Sidiropoulos, M. Horowitz, "A Semi-Digital Dual Delay-Locked Loop," IEEE JSSCC, Nov. 1997.
- [5] S. Sidiropoulos, M. Horowitz, "A 700 Mbps/pin CMOS Signalling Interface Using Current Integrating Receivers," IEEE JSSCC, May 1997.




Figure 1 32×32 crossbar and 32 port cards with conventional serial links and asymmetric serial links




Figure 2 Asymmetric Serial Link




Figure 3 Receiver Architecture and Phase Adjust Logic




Figure 4 Smart End Clock Generation

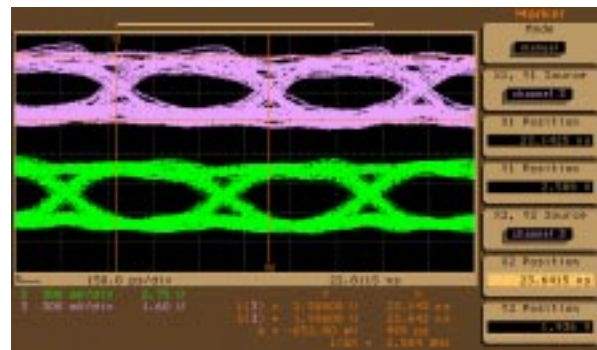


Figure 5 Transmitter Output (Top: Smart Tx, Bottom: Smart Rx)




Figure 6 Receiver Timing Margins (2Gb/s)

Supply Voltage	2.5V
Technology	TI 0.25 μ m CMOS
PLL Range (VCO rate)	200MHz-1.44GHz
Power Dissipation - PLL	175mW @2Gb/s
Jitter (@1 GHz, quiet supply)	17.7ps pk-pk/1.95ps RMS
Supply Sensitivity (pk-pk)	0.75ps/mV
Max Data Rate	2.2Gb/s
Min Input Amplitude (differential signalling)	150mV
Min Input Amplitude (single-ended signalling)	650mV

Table 1: Performance Summary