

A 20Gb/s 0.13um CMOS Serial Link Transmitter Using an LC-PLL to Directly Drive the Output Multiplexer

Patrick Chiang¹, William J. Dally^{1,2}, Ming-Ju Edward Lee³, Ramesh Senthinathan³, Yangjin Oh¹ and Mark Horowitz¹

¹Stanford University
353 Serra Mall #212
Stanford, CA 94305, USA

²Velio Communications
2249 Zanker Rd
San Jose, CA 95131, USA

³ATI Technologies, Inc.
4555 Great America Pkwy, #501
Santa Clara, CA 95054, USA

Tel: (650) 725-8086, E-mail: pchiang@leland.stanford.edu

Abstract

A 20Gb/s transmitter is implemented in 0.13um CMOS technology. Eight 2.5Gb/s data streams are 4:1 multiplexed, sampled, and retimed into two 10Gb/s data streams. A final 20Gb/s 2:1 output multiplexer, clocked by complementary phases of an LC-VCO (voltage controlled oscillator) in a phase-locked loop, creates 20Gb/s data. The VCO is integrated with the output multiplexer, resonating the load and eliminating the need for clock buffers. Power, active die area, and jitter (RMS/pk-pk) are 165mW, 650um x 350um, and 2.37ps/15ps, respectively.

Introduction

New developments in high speed serial links have been crucial in scaling off-chip system bandwidth with CMOS on-chip bandwidth. Bandwidth/pin must continue to increase, as CMOS scaling steadily increases the bandwidth/mm on-die. As these data rates increase, the importance of timing precision in link performance becomes the dominant factor in the ability to scale off-chip with on-chip bandwidth.

Previously, many serial link designs have used multi-phase phase locked loops, using multi-phase multiplexing, to achieve fast bandwidth/pin performance. [1,2] As data rates increase, and timing uncertainty becomes the critical bottleneck in link performance, these links suffer reduced timing margin for a few reasons. One problem is the difficulty in maintaining phase symmetry between the multiple phases. For example, threshold mismatches and capacitive layout mismatches in the timing vernier may cause static phase errors and unequal eye openings. A second problem is that conventional serial links use a series of post-PLL clock buffers, in order to increase the clock fanout for the transmitter multiplexing. As these buffers lie outside of the phase locked loop, their jitter is not reduced by the PLL feedback loop, resulting in a significant source of high frequency clock jitter, such as through a noisy supply voltage.

As a result, some current CMOS serial link transmitters of 10Gb/s retime the data at the full rate of 10GHz. [3,4] While this mitigates the phase symmetry and jitter issues, full rate architectures also increase power, area, and circuit complexity, as the on-die circuitry bandwidth is the same as the off-chip bandwidth.

In this paper, we describe a 20Gb/s transmitter implemented in 0.13um CMOS, where the final 2:1 output multiplexer/driver capacitance is subsumed directly into

the complementary nodes of the 10GHz LC-VCO, alleviating many of the aforementioned problems. First, the complementary nodes of the high Q LC resonator obey low static phase offset, resulting in symmetric eye openings. Second, by driving the final 2:1 output driver directly by the complementary sinusoidal phases of a 10GHz LC-VCO, the output data jitter depends solely on the LC-PLL jitter, as the post-PLL buffers are no longer necessary. Finally, as only the final 2:1 output multiplexer is running at the 20Gb/s line rate, this architecture achieves low area, power, and circuit complexity.

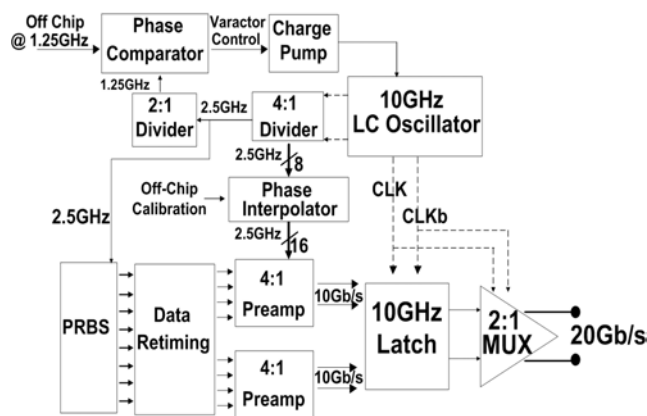


Figure 1. Transmitter architecture.

System Architecture

Figure 1 shows the complete transmitter architecture. A 20 bit, pseudo random bit sequence (PRBS) generator, creates an eight-bit wide data stream @ 2.5GHz. These eight data streams are retimed and sent to two sets of 4:1 multiplexers, creating two 10Gb/s data streams offset in phase by 50ps. The two 10Gb/s data streams are retimed to the 10GHz clock domain by two analog sample-and-hold circuits. The final 2:1 output buffer multiplexes the two 10GHz data streams to 20Gb/s. This architecture allows for the final transmitter jitter generation to depend solely on the jitter of the complementary 10GHz clock (CLK and CLKb).

All transmitter timing is derived from a 10GHz phase locked loop that employs a spiral-inductor voltage-controlled oscillator, as shown in Figure 2. The 0.5nH differential spiral is fabricated using top level metal (metal eight), and has a simulated Q of 11.8 and 10.4, with and

without ground shield, respectively. This small difference in Q is due to the inherent low doped substrate, as well as the large distance (5.6 μ m) of separation between metal eight and the substrate.

Using a LC-based voltage controlled oscillator has several advantages. First, a 10GHz oscillation frequency would be difficult to achieve, with any significant fanout, using a ring oscillator. Secondly, subsuming the output multiplexer capacitance directly into the resonator removes any post-PLL clock buffers, which are a significant source of high frequency jitter. Third, the LC voltage controlled oscillator gives lower jitter than a ring oscillator due to the smaller VCO tuning range and large Q . Finally, static phase offset of the resonator nodes is minimized, as explained below.




Figure 2. Spiral inductor voltage-controlled oscillator.

There are two potential problems that can affect the static phase offset of both resonator outputs. First, transistor mismatch such as threshold mismatch causes 2ND order distortion that distorts the sinusoids and creates asymmetry in the rise and fall time. This has minimal effect on the zero crossings of the clock signals but affects the sinusoidal symmetry, which can cause unequal eye openings. This asymmetry is attenuated by the large Q of the resonator, as the larger Q increases signal amplitude, minimizing the effect of threshold mismatch. Simulation with one PMOS having a 50mV threshold mismatch exhibits less than 1ps of static phase error.

Second, capacitance mismatch of both resonator nodes also has an affect on waveform symmetry. A capacitance difference between complementary outputs acts as a capacitive divider between the two nodes. As well, a rise and fall time asymmetry will occur, due to the asymmetric current charging/discharging of different capacitances. Both the voltage difference and rise/fall timing asymmetry cause static phase offset and uneven zero crossings. By increasing the drive strength of the $-g_m$ inverters, the capacitive dependent voltage is mitigated, as the complementary nodes become voltage limited by the supply. As well, improving the resonator Q also minimizes the effect of capacitance mismatch as the voltage mismatch is mitigated by larger voltage swings. With 50mV transistor mismatch and 5% capacitance mismatch, voltage amplitude mismatch is $< 4\%$, and static phase offset $< 2ps$.

The 10GHz clock from the PLL is divided down to 2.5GHz by a 1:4 source-coupled frequency divider, and this signal is further divided to 1.25GHz by a 1:2 dynamic-CMOS frequency divider. A high voltage (3.3V) PMOS source follower buffers the voltage between the charge

pump output and the varactor control. This buffer provides isolation of the charge pump, in an attempt to minimize any nonlinearity and coupling between the charge pump and the VCO/varactors.

The transmitter multiplexing starts with a 20 bit pseudo random bit generator (PRBS), creating eight parallel 2.5Gb/s data streams. To operate at this frequency, dynamic logic is employed along with careful design of metal interconnect and transistor sizing. After being retimed, the eight data streams are multiplexed to two 10Gb/s data streams using eight 2.5GHz phases from the phase locked loop divider. The generation of the 10Gb/s data stream is illustrated in Figure 3(a). The data is multiplexed using a current summing technique, forming a 600mV voltage swing. Each rising and falling clock of the 4:1 multiplexer is controlled by an interpolator increasing timing margin, and allows both 10Gb/s data streams to be moved 100ps in phase. Using a three NMOS pulldown stack decreases the bandwidth, but allows more precise control of the timing margin of the 10Gb/s data streams, as compared to a NAND two stack multiplexer. [5]

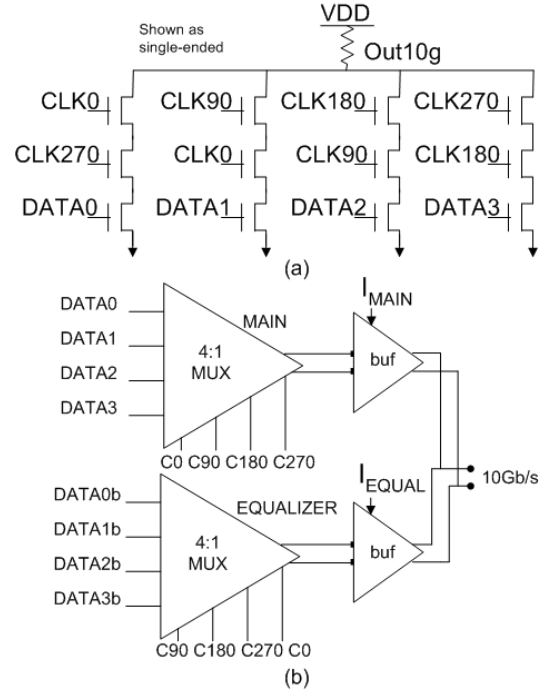


Figure 3. (a) 4:1 10Gb/s multiplexer. (b) 2-tap current summing equalizer for on-chip 10Gb/s signals.

After the data is multiplexed to 10Gb/s, it is buffered and current summed again with an equalizing path. (Figure 3(b)) The equalizer has two purposes. First it helps improve the bandwidth of the three stack NMOS in the 4:1 preamp. Secondly, it compensates for the limited bandwidth/hysteresis in the proceeding 10GHz analog sample-and-hold, described now.

The two 10Gb/s data streams are retimed by a set of analog sample-and-holds, shown in Figure 4. On the falling edge of the 10GHz clock, data is sampled onto the internal nodes MID and MIDb. Note that the analog sampled voltage is not reset to a common mode voltage, since only true and complement phases are provided. Therefore, significant hysteresis exists, and the previous

stage equalizer becomes a necessity. The timing aperture of this sampling gate depends on the RC time constant of the internal node, and the edge rate of the sampling clock. Simulation shows a 24 ps aperture time for this latch.




Figure 4. 10Gb/s analog sample-and-hold.

In order to ensure that the two 10GHz sampling clocks sample the two 10Gb/s data streams when the data is valid, phase interpolators are placed between the PLL output and the 4:1 multiplexers. This essentially shifts the zero crossings of the 10Gb/s data streams in relation to the sampling 10GHz CLK and CLKB. Each phase interpolator exhibits two methods of phase control. Tri-state current summing inverters provide 2 bits of coarse control, while 2 bits of fine capacitive trimming control another 2 bits. A maximum phase step of 7ps, across 50ps of range, is observed in simulation.

Finally, the retimed 10Gb/s data streams are multiplexed by the final 2:1 output driver, shown in Figure 5. The 2:1 output driver is implemented using source coupled pairs, current summing on both differential pairs through 50 Ohm termination resistors.




Figure 5. (a) 10GHz analog sample-and-hold circuit. (b) Final 2:1 output driver/multiplexer.

One possible concern is increased jitter due to data dependent kickback from the 2:1 output driver back into the LC resonator. This occurs through data dependent modulation through C_{GD} into the resonator. This effect is mitigated, as the total capacitance of the tail NMOS nodes is only 38fF of the total 1pF resonance capacitance, of which only 13fF is parasitic drain-gate capacitance. Also, the differential pair provides essentially a virtual ground at

the drain of the tail NMOS, reducing any data dependent kickback into the resonant tank. Simulation shows less than 1ps pk-pk added jitter due to this transmitter architecture.




Figure 6. Die photo micrograph.

Experimental Results

The 20Gb/s transmitter was fabricated in a standard CMOS 0.13um, 1.2V process. Cascade Microtech on-die wafer probes are used to achieve 50 Ohm termination for the 20Gb/s signal. The off-chip 1.25GHz clock is derived from an Agilent 8133A, with RMS and pk-pk jitter of 1.25ps and 8.9ps, respectively. All jitter/eye diagram measurements are done using an HP54754A. The active area of the transmitter is 650um x 350um, with 130um x 130um used for the inductor. The PLL and the total transmitter power(including PLL) are 33mW and 165mW, respectively. The tuning range of the transmitter is from 16.32Gb/s to 19.2Gb/s, and was observed in parasitic extraction simulation late in tapeout. Our target 20Gb/s data rate was not achieved due to un-optimized layout of the phase locked loop and the transmitter muxes in relation to the LC resonator. This caused excessive wiring capacitance between these blocks. For example, 200um of wide top layer metal wiring separates the transmitter muxes and the frequency divider. More careful attention to placement of the initial 4:1 frequency divider would have mitigated this frequency offset.

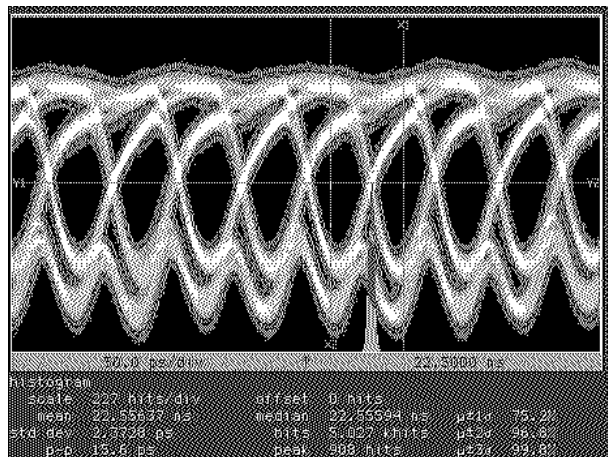


Figure 7. Transmitter eye diagram.

Figure 7 shows an eye diagram at the 19.2Gb/s data rate. (both outputs superimposed on plot) The RMS and pk-pk jitter are 2.37ps and 15.6ps. The eye opening amplitude is approximately 105mV. Measured static phase offset between eight consecutive bits is less than 2ps. There is a 60mV voltage ripple seen at the bottom of the eye, also predicted from simulation. This phenomenon occurs during the transition period of the sinusoidal complementary clocks, when neither tail NMOS is fully on, and the differential pair no longer acts with a constant bias current. To observe the effect of data dependent kickback into the LC resonator causing increased jitter, measurements of the transmitter PLL's 2.5GHz clock were observed in various transmitter configurations. With only the PLL on and the transmitter off, the jitter(RMS,pk-pk) is 1.2ps and 10ps, respectively. When the transmitter is completely on (sending PRBS data), the jitter is 2.25ps and 15.6ps. Leaving the PRBS and digital logic on but turning off the analog multiplexing stages gives a jitter of 2.3ps and 16.7ps. This shows that the increase in PLL jitter is due to digital logic causing power supply noise, and the effect of data dependent kickback is not a dominant factor. Note that the 2.5GHz clock jitter with transmitter off is comparable to the 8133A clock jitter. Therefore, the true PLL performance measurements are limited by the input clock noise.

Conclusion

A 20Gb/s serial link was designed in a 0.13um CMOS process. The final transmitter output multiplexer is clocked directly by the two complementary phases of the LC resonator in the 10GHz phase locked loop. This allows most of the transmitter to run at the half rate of 10Gb/s, decreasing the area, power consumption, and complexity in respect to full-rate architectures. In addition, subsuming the output mux capacitance into the resonator removes PLL clock buffering, a significant source of high frequency jitter. Finally, due to the inherent symmetry of a LC voltage controlled oscillator, both complementary phases exhibit little static phase offset, resulting in symmetric eye openings. The 20Gb/s transmitter dissipates 165mW in 650um x 350um.

Acknowledgements

The authors would like to thank John Poulton, Ken Mai, and Jaeha Kim for helpful discussions, and Mark Kellam, Robert Palmer, and Hiok-Tiaq Ng for fabrication support, Professor Tom Lee for testing equipment, and Pauline Prather for bond wire expertise.

TABLE I
TEST CHIP PERFORMANCE SUMMARY

Active area	Transmitter: 0.23mm ² PLL: 0.073mm ² Inductor: 0.017mm ²
Power (105mV diff. swing)	Transmitter: 165mW PLL: 33mW LC-VCO: 6.4mW
Transmitter tuning range	16.32 - 19.2 GHz
Transmitter output jitter	2.37ps RMS, 15.2ps p-p
Static phase offset	< 2ps

References

- [1] C.-K. Yang and M. Horowitz. "A 0.8um CMOS 2.5Gb/s oversampling receiver and transmitter for serial links." *IEEE Journal of Solid-State Circuits*, December 1996, Vol. 31, No. 12, pp. 2015-2023.
- [2] M.-J. E. Lee, W. J. Dally, P. Chiang. "Low-Power Area-Efficient High-Speed I/O Circuit Techniques." *IEEE Journal of Solid-State Circuits*, November 2000, Vol. 35, No. 11, pp. 1591-1599.
- [3] J. Cao, et. al. "OC-192 Transmitter and Receiver in standard 0.18-um CMOS." *IEEE Journal of Solid-State Circuits*, December 2002, Vol. 37, No. 12, pp. 1768-1780.
- [4] L. Henrickson, et. al. "Low-Power Fully Integrated 10-Gb/s SONET/SDH Transceiver in 0.13-um CMOS." *IEEE Journal of Solid-State Circuits*, October 2003, Vol. 38, No. 10, pp. 1595-1601.
- [5] W. Ellersick, et. al. "A Serial-Link Transceiver Based on 8 Gsample/s A/D and D/A Converters in 0.25um CMOS", *JSSCC Dig. Of Tech. Papers*, pp. 58-59, Feb. 2003.