

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

Patrick Chiang, *Student Member, IEEE*, William J. Dally, *Fellow, IEEE*, Ming-Ju Edward Lee, *Member, IEEE*, Ramesh Senthinathan, *Senior Member, IEEE*, Yangjin Oh, and Mark A. Horowitz, *Fellow, IEEE*

Abstract—A 20-Gb/s transmitter is implemented in 0.13- μm CMOS technology. An on-die 10-GHz LC oscillator phase-locked loop (PLL) creates two sinusoidal 10-GHz complementary clock phases as well as eight 2.5-GHz interleaved feedback divider clock phases. After a $2^{20} - 1$ pseudorandom bit sequence generator (PRBS) creates eight 2.5-Gb/s data streams, the eight 2.5-GHz interleaved clocks 4:1 multiplex the eight 2.5-Gb/s data streams to two 10-Gb/s data streams. 10-GHz analog sample-and-hold circuits retune the two 10-Gb/s data streams to be in phase with the 10-GHz complementary clocks. Two-tap equalization of the 10-Gb/s data streams compensate for bandwidth rolloff of the 10-Gb/s data outputs at the 10-GHz analog latches. A final 20-Gb/s 2:1 output multiplexer, clocked by the complementary 10-GHz clock phases, creates 20-Gb/s data from the two retimed 10-Gb/s data streams. The LC-VCO is integrated with the output multiplexer and analog latches, resonating the load and eliminating the need for clock buffers, reducing power supply induced jitter and static phase mismatch. Power, active die area, and jitter (rms/pk-pk) are 165 mW, 650 $\mu\text{m} \times 350 \mu\text{m}$, and 2.37 ps/15 ps, respectively.

Index Terms—High-speed I/O, LC oscillators, multiplexing, OC-192, phase-locked loops, sample and hold circuits, transmitters.

I. 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 one of the dominant factors in the ability to scale off-chip with on-chip bandwidth.

Previously, many serial link designs have used multiphase phase-locked loops (PLLs), using multiphase 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. Various techniques have been proposed to alleviate this static phase offset problem [3]–[5], but they suffer from large area overhead as well as residual static phase offset from quantization error.

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 PLL, 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. Various techniques have been proposed to alleviate buffer sensitivity [2], [6], but the need for high fanout due to large transmitter multiplexing (muxing) capacitances results in large power dissipation for these techniques.

As a result, some current CMOS serial link transmitters of 10 Gb/s retune the data at the full rate of 10 GHz [7], [8]. 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 20-Gb/s transmitter implemented in 0.13- μm CMOS, where the final 2:1 output multiplexer/driver capacitance is subsumed directly into the complementary nodes of the 10-GHz 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 10-GHz 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 20-Gb/s line rate, this architecture achieves low area, power, and circuit complexity.

The following is the outline for the rest of this paper. Section II briefly describes the overall transmitter architecture. Section III discusses in detail the performance of the LC-VCO, delving into the power supply sensitivity and residual static phase offset. Section IV discusses the PLL. Section V describes the transmitter muxing circuitry to achieve 20-Gb/s data rate. Section VI shows the experimental results.

Manuscript received September 3, 2004; revised December 3, 2004. This work was supported by grants from Intel Corporation, MARCO, and Cray.

P. Chiang, Y. Oh, and M. A. Horowitz are with the Computer Systems Laboratory, Stanford University, Stanford, CA 94305 USA (e-mail: pchiang@stanford.edu).

W. J. Dally is with Stream Processors, Inc., Sunnyvale, CA 94085 USA, and also with the Computer Systems Laboratory, Stanford University, Stanford, CA 94305 USA.

M.-J. E. Lee is with ATI Technologies Systems Corporation, Inc., ATI Research Silicon Valley Inc., Santa Clara, CA 95054 USA.

R. Senthinathan is with ATI Technologies, Inc., Markham, ON L3T 7X6, Canada.

Digital Object Identifier 10.1109/JSSC.2004.842841

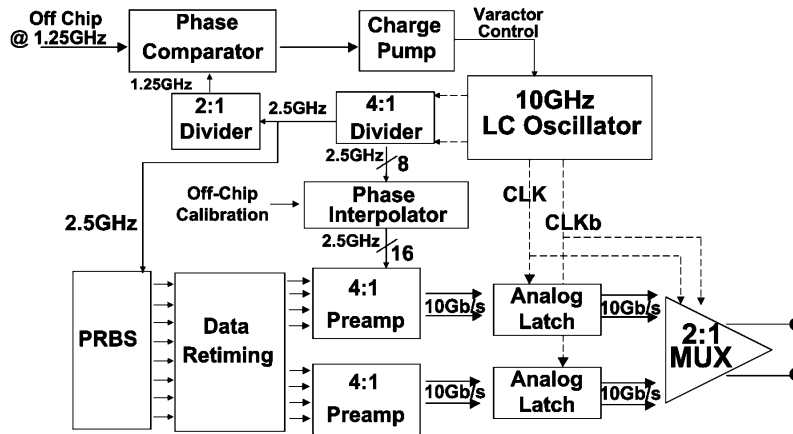


Fig. 1. Transmitter architecture.

II. SYSTEM ARCHITECTURE

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

Notice that a significant amount of delay exists between the 16 2.5-GHz clock phases and the complementary 10-GHz phases, making it difficult to ensure valid setup/hold time of the two 10-Gb/s data streams in relation to the 10-GHz complementary clocks. Each of the clock phases suffers from large delay variance, as the phases pass through low-to-high level conversion, interpolator stages, and fanout buffering before arriving at the 4:1 10-Gb/s multiplexer. Hence, a 5-bit DAC phase interpolator is needed to adjust each of the sixteen phases. Coarse interpolation is achieved using tri-state buffer current summing, while fine interpolation is done through capacitive trimming, resulting in a minimum interpolation step of 8 ps and a maximum step of 50 ps. Calibration logic to adjust the 16 2.5-GHz phases in relation to the 10-GHz complementary clocks was not implemented on this chip, and therefore, timing margin is achieved off-line by adjusting the interpolator steps through a scan chain.

III. LC-VCO

The transmitter output jitter/skew is limited by the intrinsic jitter/skew of the clocks generated by the phase-locked loop. Any additional buffer stages between the PLL (point of generation) and the transmitter muxing (point of use) add additional jitter/skew to the nominal jitter/skew of the PLL. Qualitatively, this can be seen in Fig. 2. As the number of buffer stages in the clock buffer chain increases, static phase offset and power supply induced jitter increase. Static phase offset increases as more devices exist in the clock chain, increasing




Fig. 2. Relationship between clock buffering and increased static phase offset/power supply induced jitter.

the probability for more threshold voltage mismatch and capacitance layout mismatch, resulting in timing skew between multiphase clocks. Likewise, each additional buffer adds additional power supply sensitivity, resulting in larger timing jitter as a result of power supply noise. At the limit, using the architecture implemented in this paper, no buffers are used between the voltage-controlled oscillator (VCO) (point of generation) and the transmitter muxing (point of use), leaving only the residual power supply sensitivity/static phase skew of the intrinsic VCO itself. The next section will describe the implemented VCO and the residual power supply sensitivity/static phase skew of the VCO.

Fig. 3(a) shows the implemented LC-VCO, consisting of a spiral inductor placed across the complementary nodes of cross-coupled inverters. The 0.5-nH differential spiral is fabricated using top level metal (metal eight), and has a simulated Q (using ASITIC [9]) 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 μm) of separation between metal eight and the substrate.

Using an LC-based VCO has several advantages. First, a 10-GHz oscillation frequency would be difficult to achieve, with any significant fanout, using a ring oscillator, or would require a prohibitively large amount of power. For example, a three-stage 10-GHz current-mode logic (CML) ring oscillator was simulated in HSPICE, burning 40 mW, versus 5 mW for a $Q = 10$ LC-VCO.

Second, subsuming the output multiplexer capacitance directly into the resonator removes any post-PLL clock buffers,

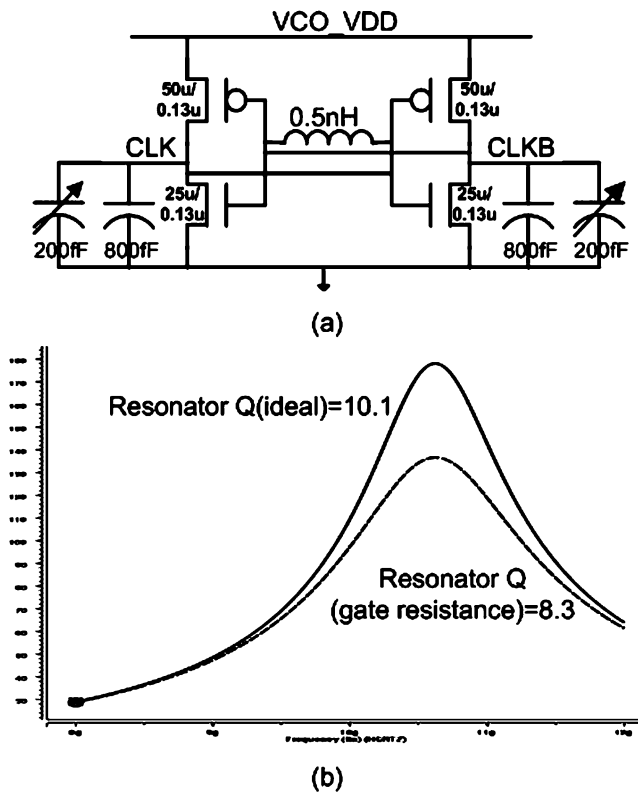


Fig. 3. (a) LC-VCO schematic. (b) Resonator Q with/without poly gate resistance.

which are a significant source of high-frequency jitter. The measured supply sensitivity (fraction delay change per fraction of supply change [2], [7]) of the LC-VCO was measured by applying a 10-ps rise time, 10% noise voltage on the supply, resulting in 0.11 (transient phase step) and 0.01 (steady-state phase error). These numbers are comparable to a supply regulated delay line [2] and power supply compensated buffer [7], but only dissipates 5 mW while driving a large fanout of 1 pF, at a cost of larger inductor area. Higher rejection of power supply noise can be achieved with smaller devices and higher oscillator Q , while incurring the penalty of larger inductor area.

Third, subsuming the load capacitance directly into the LC tank and resonating the capacitance significantly decreases the power dissipated. For example, simulation of two stages of inverters driving the clock loading capacitance burns an additional 10 mW.

One possible problem might concern resonator degradation due to the many sources of capacitances that sum to the total resonator capacitance. Half of the tank capacitance is dominated by high- Q varactor capacitance (460 ff), while passive wiring capacitance makes up another quarter of total capacitance, leaving approximately 270 ff of NMOS gate capacitance. Careful layout to reduce any parasitic resistances, such as S/D series resistance and gate resistance, is crucial towards minimizing resonator loss. For example, the varactor uses 32 parallel fingers of $6 \mu\text{m}/0.36 \mu\text{m}$, strapped on both sides of the transistor. Even with special care, the series gate poly resistance dominates resonator loss; series gate resistance is 1.5Ω , causing a Q degradation of 17%, seen in Fig. 3(b).

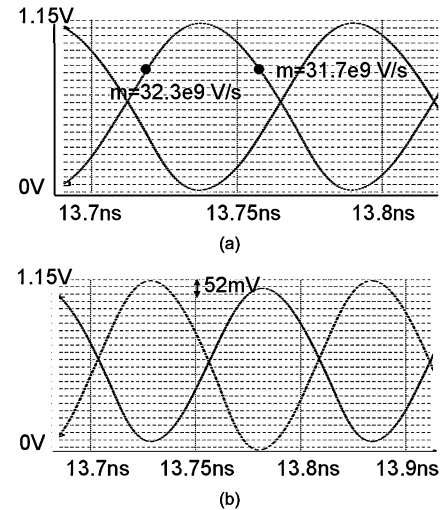


Fig. 4. (a) Rise/fall asymmetry due to -50-mV PMOS threshold voltage mismatch. (b) Voltage mismatch/asymmetry due to 5% capacitance mismatch between resonator nodes.

Another possible problem might arise from long wiring traces between the VCO and various blocks. One particular route, from the LC-VCO to the frequency divider, is unusually long ($300 \mu\text{m}$), with 20 ff of the frequency divider gate capacitance at the end. ASITIC parasitic extraction showed that this M8 thick metal ($300 \mu\text{m} \times 4 \mu\text{m}$) at 10 GHz is observed as a 0.1-nH $0.2\text{-}\Omega$ 20-ff transmission line. Transient simulation shows a phase lag at the divider of <1 ps, and a Q degradation of 1% from this long wiring trace. Next, the residual static phase offset of the VCO is described.

There are two potential problems that can affect the static phase offset of both resonator outputs. First, transistor mismatch such as threshold mismatch causes second-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. For example, as seen in Fig. 4(a), when one PMOS of the cross coupled pair has a -50-mV threshold mismatch, the rise time slope is larger than the fall time, due to the improved g_m of the PMOS device with the smaller threshold voltage. However, the resulting static phase error is still <1 ps.

Second, capacitance mismatch of both resonator nodes also has an effect on waveform symmetry and possibly static phase offset. 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 threshold voltage mismatch is mitigated by larger voltage swings. Fig. 4(b) illustrates a simulation with 50-mV transistor

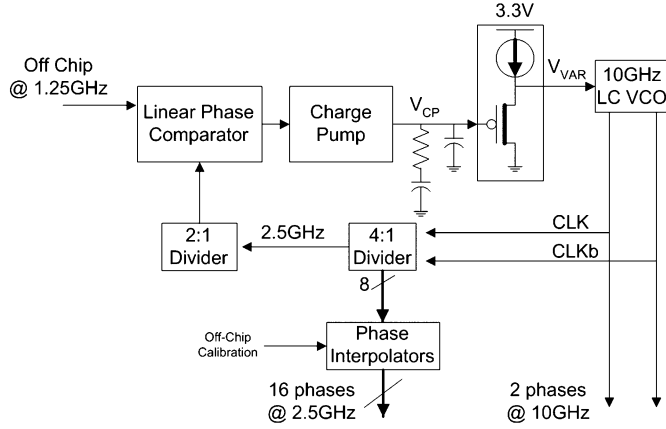


Fig. 5. 10-GHz PLL.

mismatch and 5% capacitance mismatch, exhibiting voltage amplitude mismatch $<4\%$, and static phase offset <2 ps.

IV. PHASE-LOCKED LOOP (PLL)

The block diagram for the PLL is seen in Fig. 5. A 1.25-GHz off-chip clock is sent to a linear phase comparator, a conventional charge pump, and two-pole loop filter. An unsilicided poly resistor is used for the loop filter zero, and the two filter capacitors use NMOS gate capacitance to achieve high capacitive density. The LC-VCO generates 10-GHz complementary clocks, which are divided by a 4:1 CML divider. The output is further divided by a digital 2:1 divider in order to run the phase comparator and charge pump at a lower frequency of 1.25 GHz.

A PMOS source follower, using the thick oxide transistors, is used as a level shifter between the charge pump loop filter (V_{CP}) and the actual VCO varactor control (V_{VAR}). This has two advantages. One is that the loop filter voltage is more isolated from the VCO, reducing phase noise enhancement from nonlinear effects, such as large-signal AM modulation. For example, varactor control voltage ripple is reduced $3\times$ from 230 mV to 75 mV. Second, gain reduction in the source follower increases the varactor voltage dynamic range. Since the source follower converts the 1.2-V charge pump voltage into the 3.3-V voltage domain, it allows for larger varactor bias dynamic range and therefore, increased tuning range.

V. TRANSMITTER MULTIPLEXING

The transmitter 8:1 muxing starts with a 20-bit ($2^{20} - 1$) PRBS, creating eight parallel 2.5-Gb/s data streams as in Fig. 6(a). These eight data streams are 4:1 multiplexed into two 2-tap equalized 10-Gb/s data streams, using the 16 interpolated clocks from the PLL. The two 10-Gb/s data streams are retimed to the clock domain of the 10-GHz LC-VCO clock by the use of 10-GHz analog latches. The two 10-Gb/s retimed data streams are then sent to the final 2:1 output multiplexer, achieving 20-Gb/s data rate at the differential output.

A. 4:1 10-Gb/s Multiplexer

Sixteen 2.5-GHz clock phases create two 10-Gb/s data streams using a 4:1 output multiplexing scheme. Fig. 6(b) shows one of the 10-Gb/s multiplexing paths, using four




Fig. 6. (a) Transmitter muxing block diagram. (b) 4:1 10-Gb/s multiplexer. (c) 4:1 10-Gb/s multiplexer with 2-tap equalization.

2.5-GHz phases. Notice that both the rising and falling edges of a single clock are separated into two clock phases; for example, clk0_{UP} and clk0_{DOWN} are two separate clocks. It is possible to use only four clock phases for each 4:1 multiplexer, but because a phase interpolator adjusts the phases for the 4:1 multiplexing stage, additional design complexity exists when the rising/falling edges of an interpolated clock need to be adjusted independently of each other.

Alternatively, the use of a two-stage NMOS pulldown multiplexer can be achieved using a dynamic AND structure, which merges the clock phase with the data [10], as opposed to the implemented three-stack NMOS multiplexer. However, such a scheme adds two stages of transistors between the clock signal and the multiplexer, allowing for possible static timing mismatch at the multiplexer that cannot be accounted for by the preceding phase interpolator. Using a three-stage NMOS pulldown multiplexer ensures that if the eight interpolated multiphases are aligned with the two 10-GHz complementary clock phases, timing closure can be insured.

While the implemented three-stage NMOS pulldown multiplexer minimizes the possible static phase offset problem, gain/bandwidth of the multiplexer becomes an issue. Also, aperture time/bandwidth limitations of the preceding 10-GHz analog latch increases intersymbol interference (ISI), translating to ISI at the final 20-Gb/s output eye. Two-tap equalization is employed to relax the bandwidth constraints, by equalizing the

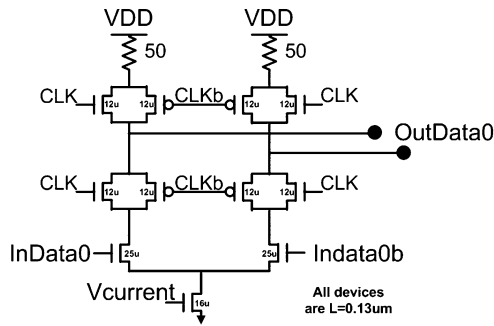


Fig. 7. 10-GHz analog latch.

low-pass channel characteristics. Fig. 6(c) shows the current summing at the output of the 10-Gb/s data streams, using a delayed version of the data to achieve 10-Gb/s 2-tap equalization. The benefits of this equalization will be described in the next section.

B. 10-GHz Analog Latch

Fig. 7 shows the schematic of the 10-GHz analog latch used. As the complementary 10-GHz clocks CLKb and CLK fall, the differential signal InData0 is sampled onto the intermediate node OutData0. Full pass gates are used to mitigate charge injection loss onto the sampled nodes. Other types of high-speed analog latches were deemed less effective. For example, a pass-gate analog sampler is ineffective, as a large (400–600 mV) voltage swing is sampled onto node OutData0. A CML latch is another alternative, but it has reduced aperture time for large differential voltage swings, as charge sharing through the differential pair occurs when the clock is sampled and the current is not immediately switched off.

While the aperture time is very fast at the output nodes for the implemented latch, significant hysteresis exists. Since only two-phase 10-GHz clocking exists, one phase of clock is used for sampling and the other for holding; no clock phase exists to reset the sampled nodes and remove the hysteresis. However, hysteresis can be observed as bandwidth reduction, or ISI, and compensated for using the preceding 10-Gb/s 2-tap equalization. Fig. 8 illustrates the sampled differential analog latch voltage (OutData0) with and without pre-emphasis. When pre-emphasis of OutData0 is turned off, significant ISI is observed for three post-cursor bit periods when the “010” data stream is transmitted. When pre-emphasis is enabled, the latched differential voltage remains constant over that same transmitted bit sequence. Fig. 9 also shows the utility of pre-emphasis of the analog latch. Note that the dotted line is for voltage waveforms with equalization and the solid line for waveforms without equalization. Fig. 9 (Panel 4) shows the analog latch input with and without pre-emphasis. The output of the latch in Fig. 9 (Panel 7) shows the latched analog 10-Gb/s data streams, with significant ISI at fast edge transitions. Fig. 9 (Panel 12) shows the 20-Gb/s transmitter output with and without equalization of the analog latches. A significant amount of eye closure (15%) is observed in the 20-Gb/s output eye for the nonequalized latch (solid line) as compared to the equalized case (dotted line).




Fig. 8. Simulated 10-GHz analog latch outputs without/with equalization.

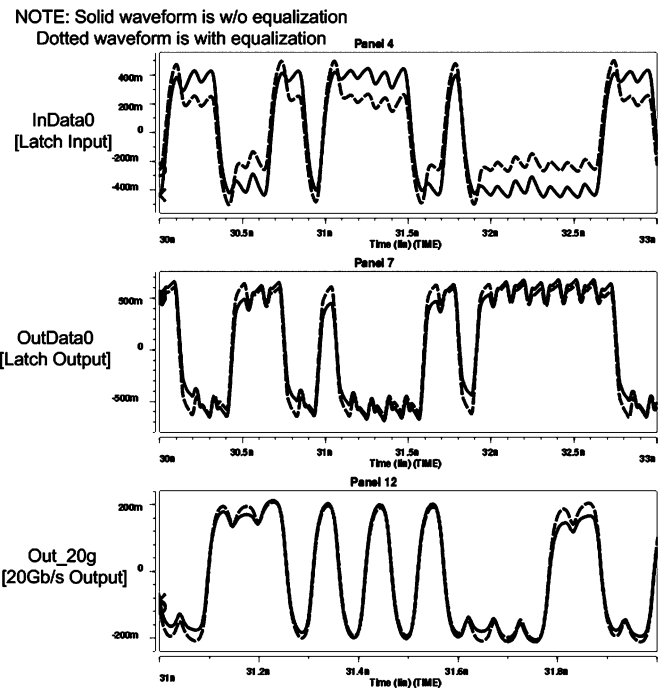


Fig. 9. Panel 4) Simulated inputs to 10-GHz analog latch with/without equalization. Panel 7) Simulated output of 10-GHz analog latch with/without equalization. Panel 12) Simulated output of 20-Gb/s 2:1 output multiplexer.

C. Final 20-Gb/s 2:1 Output Multiplexer

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

One issue of concern with this merged VCO/output driver design is the possibility of data-dependent voltage noise/kickback, degrading VCO phase noise. This problem can arise from data-dependent modulation of the differential pair tail voltage, causing charge variation/kickback through C_{GD} , ultimately affecting resonator tank capacitance and oscillation frequency. This data-dependent kickback into the LC-VCO was simulated in HSPICE by changing one of the inputs of the 2:1 multiplexer dynamically, causing the 20-Gb/s output data to toggle. The period of the VCO is measured before and after the dynamic data change, with the observed VCO period changing less than 0.2 ps. This small VCO perturbation can be attributed to the ability of the differential pair tail node to maintain a virtual




Fig. 10. Final 20-Gb/s 2:1 output multiplexer.




Fig. 11. Die photograph.


ground. Simulation shows only a 40-mV tail node voltage modulation associated with the output voltage change, resulting in an insignificant amount of charge to perturb the VCO resonator.

VI. EXPERIMENTAL RESULTS


Fig. 11 shows the die photograph. Including the bonding pads, the die area is 700 $\mu\text{m} \times 1100 \mu\text{m}$. The two arrays of inner pads along the top and bottom of the die allow for on-die, high bandwidth 50- Ω termination impedance using Cascade Microtech probes. The 20-Gb/s differential data stream travels through the on-die high bandwidth probes, through two sets of dc blocking connectors, one meter of coax, and into the oscilloscope input. Total insertion loss is 3.6 dB at 10 GHz, exhibiting some amount of frequency-dependent attenuation.

Fig. 12(a) shows the spectrum analyzer output of a 2.5-GHz divide-by-4 clock from the standalone PLL, locked at 9.6 GHz. The measured -3-dB bandwidth of the locked PLL is 13 MHz. There is bandwidth peaking of around 3.6 dB in the PLL. Due to an inappropriately sized loop filter resistor, the damping δ is less than 0.5. The phase noise at 1-MHz offset is -99.7 dBc/Hz observed in the presence of a large noise floor, and -120 dBc/Hz beyond 100 MHz.

The standalone PLL jitter was measured using an HP54754A oscilloscope, and an Agilent 8133A waveform generator. The jitter of a 1.25-GHz clock source from the 8133A was 1.13 ps (RMS), 8.9 ps (pk-pk). Fig. 12(b) shows that the jitter histogram of the 10-GHz clock output is 0.97 ps (rms), 8 ps (pk-pk). Since the PLL bandwidth of 13 MHz is relatively high, much of the



(a)



(b)

Fig. 12. (a) Power spectrum of 2.5 Hz divided clock from the 10-GHz PLL. (b) 10-GHz clock jitter histogram.

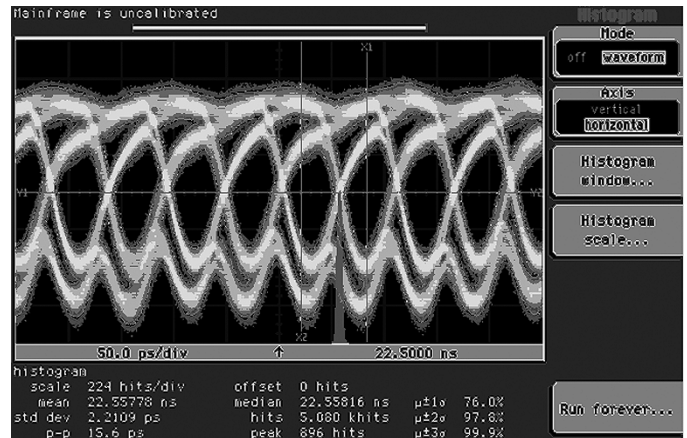


Fig. 13. 20-Gb/s eye diagram.

PLL output jitter is passed directly from the 8133A input reference. This suggests that a quieter input reference is needed to ascertain the true performance of the 10-GHz PLL. The measured integrated jitter (50 kHz-80 MHz) was 4 mUI (rms).

Fig. 13 shows the eye diagram, at a data rate of 19.2 Gb/s. The outputs are superimposed on each other. The eye opening amplitude is approximately 105 mV at the oscilloscope input. The measured jitter is 2.37 ps (rms) and 15.6 ps (pk-pk). Measured static phase offset between eight consecutive bits is less

TABLE I
MEASURED JITTER OF 2.5-GHz PLL CLOCK WITH VARIOUS CONFIGURATIONS

Various Configurations	Jitter(RMS)	Jitter(pk-pk)
Transmitter Off, PRBS Off	1.25ps	10ps
Transmitter On, PRBS Off (data:00000000)	1.25ps	8.9ps
Transmitter On, PRBS Off (data:00001000)	1.21ps	10ps
Transmitter Off, PRBS On	2.3ps	16.7ps
Transmitter On, PRBS On	2.25ps	15.6ps

than 2 ps. Notice that there is a 60-mV 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. Instead, the current drops during the transitions, causing the output nodes to pull up to the supply during the zero crossings of the clocks.

The transmitter works from 16.32–19.2 Gb/s, slower than our initial specification of 20 Gb/s. This was actually observed in parasitic extraction simulation, showing the transmitter working from 16.8–19.6 Gb/s. The cause of this is excessive capacitive loading of the complementary resonator clock nodes, with wide top thick metal ($300 \mu\text{m} \times 4 \mu\text{m}$) connecting the frequency divider and the final 2:1 output multiplexer. More careful floorplanning of the layout would have placed the divider closer to the output multiplexer, reducing this clock loading. Increasing PLL tuning range, such as using switch capacitor banks for the tank resonator, can also help alleviate this potential problem.

We attempted to determine the effect of data-dependent kickback on PLL jitter accumulation by measuring the transmitter in various configurations. This was done by measuring the 2.5-GHz divided clock coming from the PLL itself. If poor isolation exists between the transmitter VCO and the transmitter latches/muxes, we would expect the 2.5-GHz PLL clock to have the largest jitter when PRBS data is transmitted. Table I illustrates the results of our measurements. When only the PLL is enabled, with the PRBS and the transmitter muxing disabled, the measured jitter is 1.25 ps, 10 ps (rms, pk-pk). With the PLL and transmitter muxing on (sending all zeroes) and the PRBS off, the jitter is 1.25 ps, 8.9 ps (rms, pk-pk), illustrating no change in jitter. With the same configuration, while sending a lone one, the jitter again does not change, with the measured jitter at 1.21 ps, 10 ps (rms, pk-pk). The jitter of the 2.5-GHz clock increases substantially (2.3 ps, 16.7 ps) when the transmitter is turned off, but the PLL and PRBS are both on. This illustrates that the dominant source for the increased jitter is the power supply noise, either/both in the divide-by-4 or the buffer chain to the output pin. Finally, when the PLL, PRBS, and the transmitter are on, sending random 20-Gb/s data at the output, the measured jitter is 2.25 ps, 15.6 ps. Since this measured jitter is roughly the same as the situation with the PRBS on and transmitter off, we can infer that the increase in measured jitter is dominated by digital power supply jitter

caused by the PRBS. Measured results using the a spectrum analyzer do not uncover any other conclusions. While these results do not imply that data-dependent jitter does not increase PLL jitter accumulation, the results clearly suggest that such data-dependent accumulation is less significant than digital power supply noise.

VII. CONCLUSION

A 20-Gb/s serial link was designed in a $0.13\text{-}\mu\text{m}$ CMOS process. The final transmitter output multiplexer is clocked directly by the two complementary phases of the LC resonator in the 10-GHz PLL. This allows most of the transmitter to run at the half rate of 10 Gb/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 20-Gb/s transmitter dissipates 165 mW in $650 \mu\text{m} \times 350 \mu\text{m}$. As conventional link channels are unlikely to achieve flat frequency response through 10-GHz bandwidth, future work includes increasing output drive strength and implementing transmitter pre-emphasis to combat intersymbol interference.

ACKNOWLEDGMENT

The authors would like to thank J. Poulton, K. Mai, and J. Kim for helpful discussions, and M. Kellam, R. Palmer, and H.-T. Ng for fabrication support, Prof. T. Lee for testing equipment, and P. Prather for bond wire expertise.

REFERENCES

- [1] C.-K. Yang and M. Horowitz, "A $0.8 \mu\text{m}$ CMOS 2.5 Gb/s oversampling receiver and transmitter for serial links," *IEEE J. Solid-State Circuits*, vol. 31, no. 12, pp. 2015–2023, Dec. 1996.
- [2] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1591–1599, Nov. 2000.
- [3] L. Wu and W. C. Black Jr., "A low-jitter skew-calibrated multiphase clock generator for time-interleaved applications," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2001, pp. 396–397.
- [4] C.-H. Park, O. Kim, and B. Kim, "A 1.8-GHz self-calibrated phase-locked loop with precise I/Q matching," *IEEE J. Solid-State Circuits*, vol. 36, no. 5, pp. 777–783, May 2001.
- [5] R. Farjad-Rad *et al.*, "0.622–8.0 Gb ps 150 mW Serial IO macrocell with fully flexible preemphasis and equalization," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 2003, pp. 63–66.
- [6] M. Mansuri and C.-K. Ken Yang, "A low-power low-jitter adaptive-bandwidth PLL and clock buffer," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, vol. 1, Feb. 2003, pp. 430–505.
- [7] J. Cao *et al.*, "OC-192 transmitter and receiver in standard $0.18\text{-}\mu\text{m}$ CMOS," *IEEE J. Solid-State Circuits*, vol. 37, no. 12, pp. 1768–1780, Dec. 2002.
- [8] L. Henrickson *et al.*, "Low-power fully integrated 10-Gb/s SONET/SDH transceiver in $0.13\text{-}\mu\text{m}$ CMOS," *IEEE J. Solid-State Circuits*, vol. 38, no. 10, pp. 1595–1601, Oct. 2003.
- [9] A. M. Niknejad, "Modeling of passive elements with ASITIC," in *Proc. IEEE Radio Frequency Integrated Circuits (RFIC) Symp.*, June 2002, pp. 303–306.
- [10] W. Ellersick *et al.*, "A serial-link transceiver based on 8 Gsample/s A/D and D/A converters in $0.25 \mu\text{m}$ CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2003, pp. 58–59.



Patrick Chiang (S'99) received the B.S. degree in electrical engineering and computer sciences from the University of California at Berkeley in 1997, and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 2001. He is currently working toward the Ph.D. degree in the Computer Systems Laboratory, Stanford University.

In 1998, he was with Datapath Systems (now LSI Logic), working on analog front-ends for DSL chipsets. In 2004, he was a consultant at Telegent Systems, Sunnyvale, CA, working on various mixed-signal RF circuits. His interests are in ultra-wideband RF architectures, high-speed serial links, and circuit design for biological systems.




William J. Dally (M'80–SM'01–F'02) received the B.S. degree in electrical engineering from the Virginia Polytechnic Institute, Blacksburg, the M.S. degree in electrical engineering from Stanford University, Stanford, CA, and the Ph.D. degree in computer science from the California Institute of Technology (Caltech), Pasadena.

He is the Willard R. and Inez Kerr Bell Professor of Engineering and the Chair of the Department of Computer Science at Stanford University. He and his group have developed system architecture,

network architecture, signaling, routing, and synchronization technology that can be found in most large parallel computers today. While at Bell Telephone Laboratories, he contributed to the design of the BELLMAC32 microprocessor and designed the MARS hardware accelerator. At Caltech, he designed the MOSSIM Simulation Engine and the Torus Routing Chip which pioneered wormhole routing and virtual-channel flow control. While a Professor of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology, his group built the J-Machine and the M-Machine, experimental parallel computer systems that pioneered the separation of mechanisms from programming models and demonstrated very low overhead synchronization and communication mechanisms. At Stanford University, his group has developed the Imagine processor, which introduced the concepts of stream processing and partitioned register organizations. He has worked with Cray Research and Intel to incorporate many of these innovations in commercial parallel computers, with Avici Systems to incorporate this technology into Internet routers, co-founded Velio Communications to commercialize high-speed signaling technology, and co-founded Stream Processors, Inc. to commercialize stream processor technology. He currently leads projects on high-speed signaling, computer architecture, and network architecture. He has published over 150 papers in these areas and is an author of the textbooks *Digital Systems Engineering* (Cambridge, U.K.: Cambridge University Press, 1998) and *Principles and Practices of Interconnection Networks* (San Mateo, CA: Morgan Kaufmann, 2003).

Dr. Dally is a Fellow of the Association for Computing Machinery (ACM) and has received numerous honors including the IEEE Seymour Cray Award and the ACM Maurice Wilkes award.



Ming-Ju Edward Lee (S'98–M'01) received the B.S. degree in electrical engineering and computer sciences from the University of California at Berkeley in 1997, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2000 and 2001, respectively.

He is currently an Engineering Manager of circuit design with ATI Research Silicon Valley, Santa Clara, CA. From 2000 to 2003, he was with Velio Communications Inc., where he led a 6.25-Gb/s serial I/O development.



Ramesh Senthinathan (S'82–M'86–SM'04) received the B.S. degree in computer engineering from the State University of New York at Buffalo in 1984, and the M.S. and Ph.D. degrees in electrical engineering from the University of Arizona, Tucson, in 1986 and 1992, respectively.

He is currently a Director of Engineering responsible for I/O, analog, and technology development with ATI Technologies, Santa Clara, CA. He was a Director of Engineering with Velio Communications, Milpitas, CA, responsible for both technology and product development from 2001 to 2003. From 1995 to 2001, he was with Intel Corporation as a Director and Distinguished Engineer for Communication and Microprocessor groups in Sacramento and Folsom, CA. He was responsible for all aspects of circuit design of Pentium Pro (Desktop and Server) and Pentium III microprocessor design. With Level One acquisition, he moved to the Intel Communication group, and managed Analog Front End group for Intel's DSL division. From 1993 to 1995, he was with Motorola, Inc. as a Principal Engineer responsible for 16/24 bit DSP analog and I/O circuit designs. He was a Staff Engineer with IBM Research from 1992 to 1993. From 1986 to 1989, he was a Design Engineer with the ASIC and Microcontroller groups at Intel Corporation, Chandler, AZ. He has published more than 50 refereed papers in these areas, and is a coauthor of the book *Simultaneous Switching Noise of CMOS Devices and Systems* (New York: Springer, 1993).



Yangjin Oh received the B.S. degree from Seoul National University, Seoul, Korea, in 2000 and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 2002. He is currently working toward the Ph.D. degree at Stanford University.

He interned as a Circuit Design Engineer at National Semiconductor, Santa Clara, CA, during the summer of 2002. His research interests include signal processing and mixed-signal circuit design in wireless communication.



Mark A. Horowitz (S'77–M'78–SM'95–F'00) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978, and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is the Yahoo Founder's Professor of Electrical Engineering and Computer Science at Stanford University. His research area is in digital system design, and he has led a number of processor designs including MIPS-X, one of the first processors to include an on-chip instruction cache, TORCH, a statically scheduled, superscalar processor that supported speculative execution, and FLASH, a flexible DSM machine. He has also worked in a number of other chip design areas including high-speed and low-power memory design, high-bandwidth interfaces, and fast floating point. In 1990 he took leave from Stanford to help start Rambus Inc., Los Altos, CA, a company designing high-bandwidth chip interface technology. His current research includes multiprocessor design, low power circuits, memory design, and high-speed links.

Dr. Horowitz received the Presidential Young Investigator Award and an IBM Faculty development award in 1985. In 1993, he received the Best Paper Award at the IEEE International Solid-State Circuits Conference.