3.8 A 0.45-to-0.7V 1-to-6Gb/s 0.29-to-0.58pJ/b Source-Synchronous Transceiver Using Automatic Phase Calibration in 65nm CMOS

Woo-Seok Choi[1], Guanghua Shu[1], Mrunmay Talegaonkar[1], Yubo Liu[1], Da Wei[1], Luca Benini[2], Pavan Kumar Hanumolu[1]

[1]University of Illinois, Urbana, IL,
[2]University of Bologna, Bologna, Italy

Supply voltage (V_{DD}) scaling offers a means to greatly reduce power in serial link transceivers. Ideally, power efficiency at a given data rate can be improved by reducing V_{DD} while increasing the number of multiplexed circuits operating in parallel at lower clock frequencies [1]. Though increasing the amount of parallelism is desirable to scale V_{DD}, in practice, it is limited by two main factors. First, increased sensitivity to device variations (threshold voltage/dimension mismatch) at lower V_{DD} makes it extremely challenging to generate equally spaced multi-phase clocks needed in multiplexed transmitter and receiver. Phase calibration methods can correct phase-spacing errors [2,3], but their effectiveness at lower V_{DD} is limited as the calibration circuits themselves become sensitive to device variations. Second, as oscillator output swing reduces with V_{DD}, its phase noise degrades, making low-noise clock generation difficult to implement with low power dissipation. Phase noise can be suppressed by embedding the oscillator in a wide bandwidth analog phase-locked loop (APLL), but conventional charge-pump-based APLLs are difficult to design at low supply voltages (V_{DD} <0.5V). Digital PLLs (DPLLs) can operate at low V_{DD}, but they suffer from conflicting noise bandwidth tradeoffs, which prevent increasing the bandwidth to suppress oscillator phase noise adequately. In this paper, we present a phase-calibration method that enables the operation of a source-synchronous transceiver down to V_{DD} of 0.45V. The energy efficiency and data-rate of the prototype transceiver scale from 0.29 to 0.58pJ/b and 1.3Gb/s to 6Gb/s, respectively, as V_{DD} is varied from 0.45 to 0.7V.

Figure 3.8.1 shows the block diagram of the source-synchronous transceiver. The data transmitter (Tx) is composed of an 8:1 output-multiplexed single-ended voltage-mode (VM) output driver and a self-calibrated PLL-based multi-phase generator. Using 8 equally spaced phases provided by the PLL, the output driver serializes 8b parallel data into a 1b full-rate data stream. In each of the 8 slices of output driver, data bit, D_i, is gated with one unit interval (UI) pulse generated by performing AND operation on two adjacent clock phases, Φ_{i+1}, and Φ_i (i = 0~7). Compared to a conventional VM driver, this implementation is better suited for low voltage operation as it reduces the number of stacked devices in the output driver [4]. The output swing is set to 200mV_{pp} using an on-chip low-dropout regulator implemented using a pseudo-differential common-source error amplifier architecture for low-voltage operation. NMOS transistors in the output driver (M_{1i} and M_{2i}) are sized to have output impedance close to 50Ω. Instead of a full-rate clock, a 1/8th rate clock is forwarded to the receiver using a VM driver to save power.

Also shown in Fig. 3.8.1 is the un-terminated receiver (Rx) consisting of 8 parallel offset-cancelled low-voltage slicers, clocked by 8 phases provided by the PLL. It samples the transmitted serial data and generates 8 de-serialized data streams. Each slicer is implemented using a modified sense-amplifier in which offset cancellation is performed at the output using a digitally tunable capacitor array. De-skewing needed to mitigate path mismatch between the data and clock channels is performed using a single digitally controlled delay line (DCDL) in the forwarded clock path. A limiting amplifier converts low-swing forwarded clock into a rail-to-rail signal and feeds it to the DCDL, whose output is used as reference clock to the PLL. Similar to Tx PLL, the Rx PLL also employs the developed phase-calibration method to compensate for phase spacing errors.

The block diagram of the low-voltage PLL is shown in Fig. 3.8.2. Using analog proportional path and digital integral path, this hybrid architecture helps extend the PLL bandwidth thereby achieving low noise performance with low power consumption [5]. To minimize reference spur caused by offset and non-zero set-up time of the bang-bang phase detector used in the integral path, a symmetric NAND-based decision circuit is used for generating digital E/L outputs. The VCO is implemented using a 4-stage ring oscillator consisting of dual control pseudo-differential delay stages. Current source I_c that is common to all the delay stages is controlled by the PLL, while the output capacitor in each of the delay stages is separately controlled by the calibration loop. In the absence of calibration, Monte Carlo simulations indicate that the standard deviation of the phase mismatch is around 12% at V_{DD} = 0.5V, reducing to 5% at V_{DD} = 0.6V.

The four-step phase calibration technique is illustrated in Fig. 3.8.3. In the first step, the PLL is locked to the reference clock. Therefore, sum of the phase spacings between 8 adjacent phases adds up to exactly 8UI with average phase spacing of 1UI, regardless of phase-spacing errors. In the second step, a digital estimate of the average phase spacing between adjacent phases is obtained by using a digital delay-locked loop (DLL). To this end, two adjacent phases, Φ_{i+1} and Φ_i, are selected and fed to the DLL that forces the DCDL delay to be equal to the phase difference ($\Delta\Phi_i = \Phi_{i+1}-\Phi_i$), which is nominally equal to 1UI. Under this condition, the accumulator output, K[i] represents $\Delta\Phi_i$ in digital domain. By repeating this process for all the 8 adjacent phase combinations K[i], for all i = 0~7, is determined, and the mean of all K[i] values, μ_K, provides an accurate estimate of 1UI phase difference within the accuracy of the DCDL. This estimate, μ_K, is used next to correct the phase spacing errors. In third step, the PLL is disabled to prevent any interaction between the PLL and the calibration loop; the DCDL control is set to μ_K and as before two adjacent phases are selected and fed to an alternate DLL formed around the VCO delay stage. In steady state, the output capacitance of the delay stage is tuned such that phase spacing is equal to the delay of DCDL. By repeating this process for all the 8 adjacent phase combinations, $\Delta\Phi_i = \Delta\Phi_{i+1}$ is achieved for all values of i. In the final step, the PLL is enabled back to normal operation for re-locking the VCO to the reference clock. Because the PLL forces the sum of $\Delta\Phi_i$ to be 8UI, and the calibration loop forces all $\Delta\Phi_i$ to be equal, $\Delta\Phi_i$ = 1UI is achieved as desired. Note that since the same DCDL and phase detector are used for both estimating and calibrating all phases, non-idealities such as mismatch-induced offsets of the loop components that are exacerbated by low-voltage operation do not affect calibration accuracy.

The prototype transceiver, fabricated in a 65nm CMOS process using only regular-V_t devices, occupies an active area of 0.15mm^2 and is packaged in a 48pin QFN package. The measured transmitter output eye diagrams for various data-rates in the range of 1.3 to 6Gb/s are shown in Fig. 3.8.4. At low V_{DD} (<0.6V), timing margin is severely degraded by phase-spacing errors. Phase calibration improves timing margin by more than 30% at V_{DD} = 0.5V and by 50% at V_{DD} = 0.45V. The measured PLL output jitter (F_{OUT} = 250MHz at 0.5V) is 7.7ps$_{rms}$ (<0.2% of a period) and 72.8ps$_{pp}$. Because of the large bandwidth of the hybrid PLL, the achieved jitter performance is superior to a previous low-V_{DD} DPLL [6]. The measured transceiver bathtub curves based on the average bit-error rate (BER) of 8 parallel receiver outputs are shown in Fig. 3.8.5. The receiver is driven by the transmitter with 200mV_{pp} swing PRBS7 data through a channel consisting of package parasitics, 0.5-inch Tx-side FR4 trace, 36-inch SMA cable, and 1-inch Rx-side FR4 trace. Error-free operation (BER <10^{-10} with 0.15UI timing margin) is achieved only when both the sampler offset and phase spacing calibration are exercised. The entire transceiver powered by a single 0.5V supply dissipates 670μW at 2Gb/s, translating to an energy efficiency of 0.33pJ/b. When a lower V_{DD} is used in the output driver, energy efficiency is improved to 0.2pJ/b. The prototype transceiver is also tested at different data rates and supply voltages with the results summarized in Fig. 3.8.5. At 0.7V supply, the receiver achieves BER <10^{-10} with 0.1UI timing margin without phase calibration. Performance of the transceiver is compared with other state-of-the-art energy-efficient off-chip links in Fig. 3.8.6. The die micrograph is shown in Fig. 3.8.7.

Acknowledgement:
This research was in part supported by NSF under CAREER Award EECS-0954969. We thank Berkeley Design Automation for providing Analog Fast Spice (AFS) simulator.

References:
[1] J. Kim, M. Horowitz, "Adaptive supply serial links with sub-1V operation and per-pin clock recovery," *ISSCC Dig. Tech. Papers*, pp. 268-269, Feb. 2002.
[2] L. Lee, *et al*, "A sub-10-ps multiphase sampling system using redundancy, *IEEE J. Solid-State Circuits*, pp. 265-273, Jan. 2006.
[3] A. Agrawal, *et al*, "Area efficient phase calibration of a 1.6 GHz multiphase DLL," *Proc. IEEE CICC*, pp. 1-4, Sept. 2011.
[4] Y.-H. Song, *et al*, "A 0.47-0.66 pJ/bit, 4.8-8 Gb/s I/O transceiver in 65nm CMOS," *IEEE J. Solid-State Circuits*, pp. 1276-1289, May 2013.
[5] W. Yin, *et al*, "A 0.7-to-3.5 GHz 0.6-to-2.8 mW highly digital phase-locked loop with bandwidth tracking," *IEEE J. Solid-State Circuits*, pp. 1870-1880, Aug. 2011.
[6] W. Liu, *et al*, "A PVT tolerant 10 to 500 MHz all-digital phase-locked loop with coupled TDC and DCO," *IEEE J. Solid-State Circuits*, pp. 314-321, Feb. 2010.
[7] G. Balamurugan, *et al*, "A scalable 5-15 Gbps, 14-75 mW low-power I/O transceiver in 65 nm CMOS," *IEEE J. Solid-State Circuits*, pp. 1010-1019, Apr. 2008.

 978-1-4799-6224-2/15/$31.00

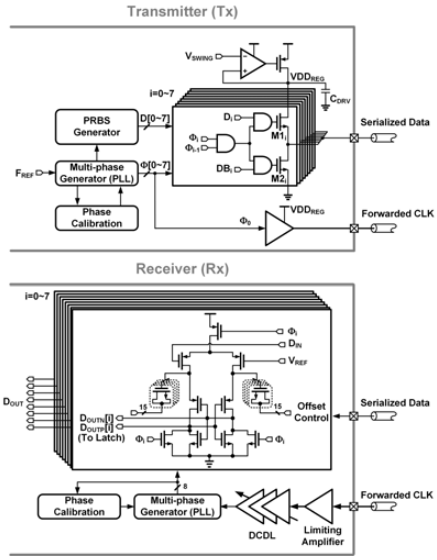

Figure 3.8.1: Overall block diagram of the transceiver.

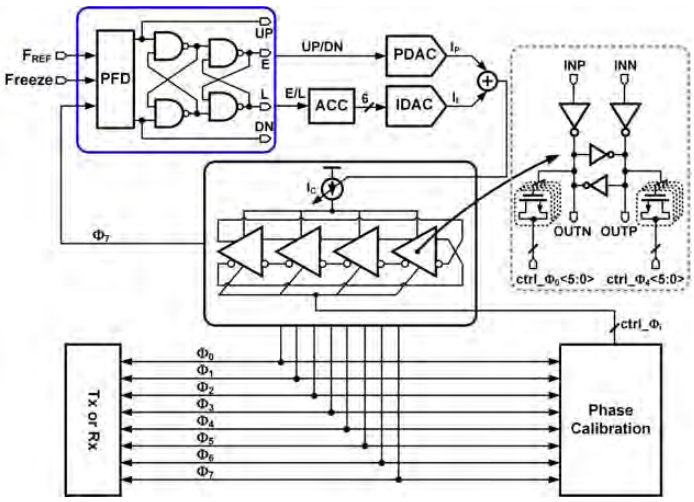

Figure 3.8.2: Low-voltage hybrid PLL.

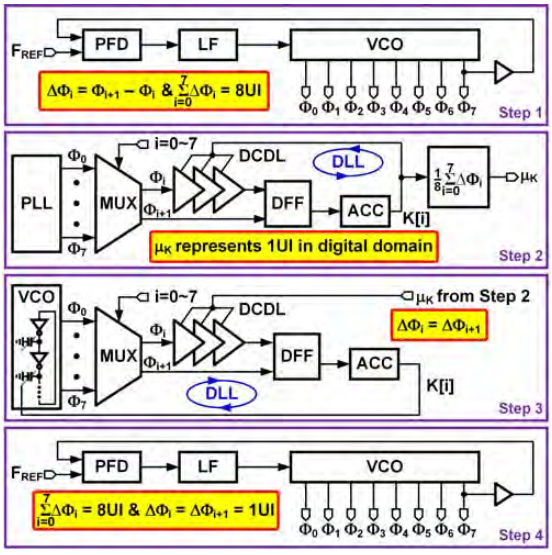

Figure 3.8.3: Phase-calibration scheme.

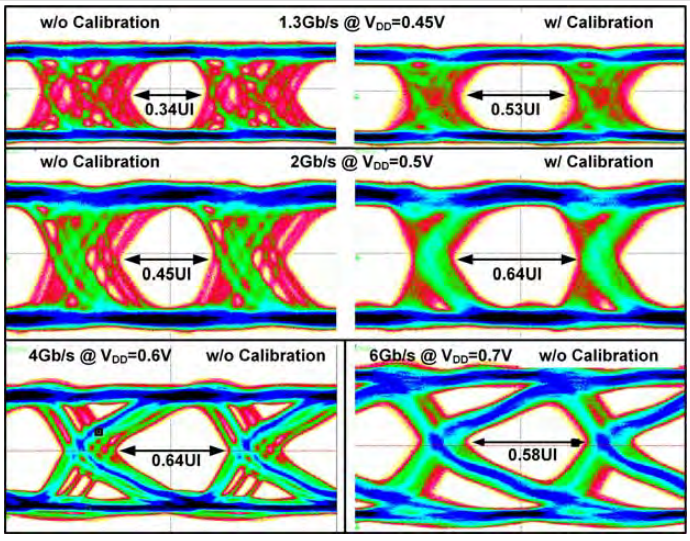

Figure 3.8.4: Measured transmitter output eye diagrams at different data-rates and supply voltages.

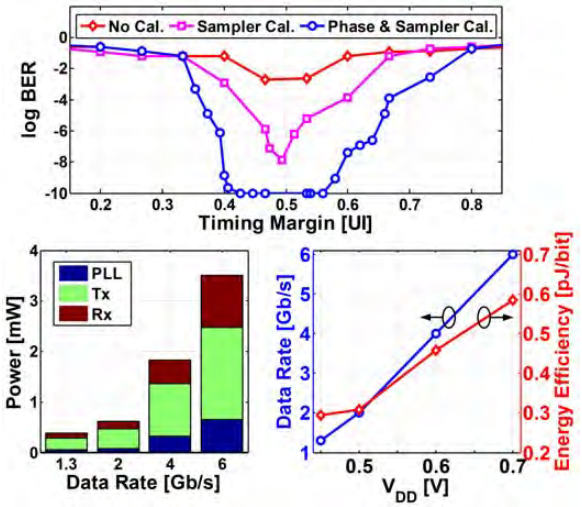

Figure 3.8.5: Measured bathtub curves for the transceiver at 0.5V supply and its data-rate, energy efficiency and power breakdown at different supply voltages.

	This Work					[5] JSSC'13 Y.Song	[7] JSSC'08 G.Balamurugan
Process	65nm CMOS					65nm CMOS	65nm CMOS
Supply Voltage (V)	0.45		0.5	0.6	0.7	0.6 ~ 0.8	0.68
Data Rate (Gb/s)	1	1.3	2	4	6	4.8 ~ 8	5
PLL Energy Efficiency (pJ/bit)	0.04	0.04	0.04	0.08	0.11	No PLL	N/A
Tx Energy Efficiency (pJ/bit)	0.21	0.18	0.20	0.26	0.31	0.3 @ 6.4Gb/s	1.5
Rx Energy Efficiency (pJ/bit)	0.07	0.07	0.07	0.12	0.17	0.17 @ 6.4Gb/s	1.2
Total Energy Efficiency (pJ/bit)	0.32	0.29	0.31	0.46	0.58	0.47 @ 6.4Gb/s	2.7

Figure 3.8.6: Performance summary and comparison with state-of-the-art energy efficient transceivers.

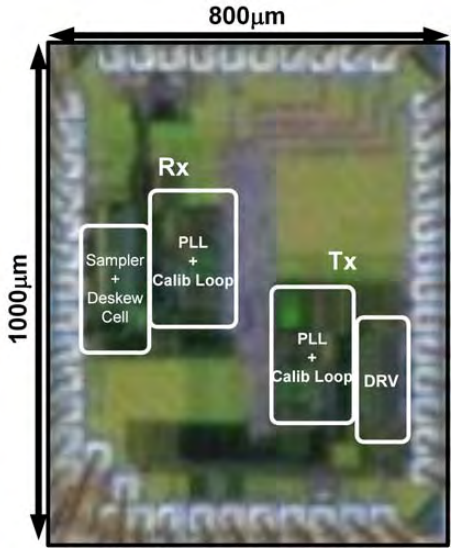

Figure 3.8.7: Transceiver die micrograph.

978-1-4799-6224-2/15/$31.00

Session 4 Overview: *Processors*

HIGH-PERFORMANCE DIGITAL SUBCOMMITTEE

Session Chair: *Atsuki Inoue, Fujitsu Labs., Kanagawa, Japan*

Session Co-Chair: *Jinuk Luke Shin, Oracle Corporation, Redwood Shores, CA*

As compute power is increasingly migrating to large data centers and the cloud, microprocessors face progressively more stringent design constraints. This year's processor session introduces three new "big iron" processors in 22nm/20nm technology providing higher performance and power efficiency. Increasing core counts and cache sizes, with adaptive power management techniques, are applied to optimize performance/W. A high-density microserver computing node is optimized for Big Data and shows its capability in handling a large number of threads. For future intelligent computing, such as deep learning and inference, a specialized processor is demonstrated. Other papers in this session describe adaptive techniques, as well as design optimizations to improve throughput and energy efficiency.

4.1 22nm Next-Generation IBM System z Microprocessor **1:30 PM**
J. Warnock, IBM Systems and Technology, Yorktown Heights, NY

In Paper 4.1, IBM presents their 678mm^2 System z microprocessor chip implemented in IBM's 22nm SOI, high-k metal-gate technology with 17 layers of wiring and functional thru 5GHz. The design comprises 8 dual-threaded cores with private 4MB eDRAM L2 cache and a 64MB unified eDRAM L3 cache, joined by a single global clock mesh and connected to high-bandwidth memory, SMP, and embedded PCI IO links.

4.2 A 20nm 32-Core 64MB L3 Cache SPARC M7 Processor **2:00 PM**
P. Li, Oracle, Redwood Shores, CA

In Paper 4.2, Oracle introduces the SPARC M7 processor with 32 S4 cores and 1.6TB/s bandwidth 64MB L3 cache to deliver more than 3.0× throughput performance over its predecessor. The chip, fabricated in 20nm, features an on-chip network (OCN) with 0.5TB/s data bandwidth, 280 SerDes lanes with 18Gb/s line rate and 1TB/s total bandwidth, and adaptive clocking scheme to deliver 2× supply noise guard-band reduction.

4.3 Fine-Grained Adaptive Power Management of the SPARC M7 Processor **2:30 PM**
V. Krishnaswamy, Oracle, Redwood Shores, CA

In Paper 4.3, Oracle describes the power management system of the SPARC M7 processor. It comprises 48 on-die digital dynamic power meters, 16 temperature sensors, and a configurable hardware proportional feedback controller which directs actuation techniques, including clock cycle skipping and DVFS, to optimize performance by 17% on a power-constrained workload.

 978-1-4799-6224-2/15/$31.00

4.4 Energy-Efficient Microserver Based on a 12-Core 1.8GHz 188K-CoreMark 28nm Bulk CMOS 64b SoC for Big-Data Applications with 159GB/s/L Memory Bandwidth System Density **2:45 PM**

R. Luijten, IBM Research, Rüschlikon, Switzerland

In Paper 4.4, IBM demonstrates a 36W density-optimized, hot-water cooled microserver compute node based on a 1.8GHz CPU with 48GB DDR3 DRAM achieving peak memory bandwidth of 43.2GB/s. 128 of such nodes are combined in a 34.7L 2U rack for a total of 6TB of memory at 159GB/s/L, along with 10Gb/s Ethernet, SATA, and USB.

4.5 The Xeon® Processor E5-2600 v3: A 22nm 18-Core Product Family **3:15 PM**

N. Nassif, Intel, Hudson, MA

In Paper 4.5, Intel's next-generation Xeon processor features 18 dual-threaded 64b Haswell cores, 45MB L3 cache, 4 DDR4-2133MHz memory channels, 40 8GT/s PCIe lanes and 60 9.6GT/s QPI lanes to achieve a 33% performance boost. The processor has 5.56B transistors on a 22nm tri-gate CMOS die measuring 663.5mm^2. Fully integrated voltage regulator (FIVR) technology enables per-core p-states and independent uncore frequency scaling.

4.6 A 1.93TOPS/W Scalable Deep Learning/Inference Processor with Tetra-Parallel MIMD Architecture for Big-Data Applications **3:45 PM**

S. Park, KAIST, Daejeon, Korea

In Paper 4.6, KAIST presents a 213.1mW, battery-powered processor with 4 deep learning, 2 deep inference cores, and a true random number generator. Manufactured in 65nm CMOS with 3.75M transistors and 8 wiring layers on 10mm^2, it achieves 411.3GOPS peak performance at 200MHz, taking ~20ms for object recognition at 640×480 resolution.

4.7 A 409GOPS/W Adaptive and Resilient Domino Register File in 22nm Tri-Gate CMOS Featuring In-Situ Timing Margin and Error Detection for Tolerance to Within-Die Variation, Voltage Droop, Temperature and Aging **4:15 PM**

J. P. Kulkarni, Intel, Hillsboro, OR

In Paper 4.7, Intel presents an adaptive and resilient register file with in-situ timing margin and error-detection techniques to tolerate within-die process variation, voltage droop, temperature and aging. The chip demonstrates 409GOPS/W operation in 22nm tri-gate CMOS to achieve 21% throughput and 67% energy-efficiency improvement with 6.4-to-12.8% area and 0.2-to-0.3% power overheads.

4.8 A 28nm x86 APU Optimized for Power and Area Efficiency **4:45 PM**

K. Wilcox, AMD, Boxborough, MA

In Paper 4.8, AMD presents their next-generation APU, Carrizo, which includes the latest x86 core, Excavator. Implemented in a 28nm HKMG process, the SoC occupies 250.04mm^2 with more than 3.1B transistors. The core realized a 23% area reduction and a 40% power reduction compared to the previous core design. The architectural improvements and power management innovations enable a reduction of typical energy use by more than 50%.