10.1 A Pin-Efficient 20.83Gb/s/wire 0.94pJ/bit Forwarded Clock CNRZ-5-Coded SerDes up to 12mm for MCM Packages in 28nm CMOS

Amin Shokrollahi[1], Dario Carnelli[1], John Fox[2], Klaas Hofstra[1], Brian Holden[1], Ali Hormati[1], Peter Hunt[2], Margaret Johnston[1], John Keay[2], Sergio Pesenti[1], Richard Simpson[2], David Stauffer[1], Andrew Stewart[2], Giuseppe Surace[2], Armin Tajalli[1], Omid Talebi Amiri[1], Anton Tschank[2], Roger Ulrich[1], Christoph Walter[1], Fabio Licciardello[1], Yohann Mogentale[1], Anant Singh[1]

[1]Kandou Bus, Lausanne, Switzerland,
[2]Kandou Bus, Northampton, United Kingdom

High-speed signaling over package substrates is key to delivering the promise of 2.5D integration. Applications abound and include high-density memory interfaces, sub-division of large dies to increase yield and lower development time, sub-division of a die to achieve upward or downward scalability, or connecting to an off-chip SerDes or optics engine. Each of these in-package applications typically has high throughput and onerously low power constraints along with a low-loss channel. Several solutions have been proposed. Interposer substrates [1], or Chip-on-Substrate-on-Wafer [2] allow for very high-density wiring and low power using CMOS transceivers. Their high manufacturing and testing cost makes them prohibitive for anything but high-end applications. A different approach using high-speed ground-referenced single-ended signaling is reported in [3], which is intended for shorter channels up to 4.5mm and a BER of 1e-12. An approach using differential signaling on up to 0.75" of Megtron 6 material and a BER of 1e-9 is reported in [4]. A comparison is given in Fig. 10.1.1.

Our approach generalizes differential signaling. Correlated quaternary values are transmitted on a 6-wire bus with a forwarded clock (FCLK) SerDes, which achieves high bandwidth, coupled with low power. The set of transmitted values belongs to a code consisting of 32 codewords called the CNRZ-5-code. 5b are transmitted on the 6 wires in every UI. The codebook consists of judiciously chosen permutations of [+1, +1/3, +1/3, -1/3, -1/3, -1]. At the receiver, the codewords are detected by the 5 self-referencing comparators shown in Fig. 10.1.2. The specific choice of the codewords eliminates transmit common-mode and simultaneous switching output noise. The CNRZ-5 codebook and its comparators have been designed to optimize the ISI-Ratio [5], a figure of merit that measures the susceptibility of a signaling code to inter-symbol Interference (ISI). The smallest possible ISI-ratio is 1 and is achieved by both the differential NRZ and the 6-wire CNRZ-5 codes by design. Traditional quaternary signaling codes like differential PAM-4 have an ISI-ratio of 3, making them inherently much more susceptible to ISI-noise.

The CNRZ-5 SerDes in 28nm CMOS comprises a transmitter, Tx, a receiver, Rx, and a common block, CB. The transmitter consists of 6 data and 2 clock wires transmitting a differential FCLK at a quarter rate, Fig. 10.1.3. The 6 wires of the output driver are terminated to a common point, V_{cm}, with a nominal value at $V_{dda}/2$. A digital encoder transforms five streams of 32b each of incoming data (total 160b) into 12 streams of coded signals for the 6 wires. Custom logic is used to convert 2b of encoded data per wire into 3 control signals for creating the wire levels in the output driver. A single CNRZ-5 encoder operating at 1.5625 GHz has 23 gates and a maximum logic depth of 2 and is implemented with standard cells. A sliced architecture source-series-terminated driver is used, where the output voltage swing, common mode and termination resistors are programmable. A nominal termination impedance of 75Ω is chosen to reduce power, while maintaining signal integrity with 50Ω MCM channels together with series inductors (layers M8-M9). An FCLK architecture provides a simple receiver clock data alignment (CDA) process, thus reducing power consumption and system complexity. The differential FCLK has a frequency of 3.125GHz (baudrate/8) and is rate but not skew matched to the 6 data wires. A local Tx mini-PLL provides the Tx clocks.

Figure 10.1.4 shows the architecture of the CNRZ-5 receiver. The Rx 6 data wires and differential FCLK are DC coupled to the line; a level shifter is used to bring the input CM voltage to an optimum value for the CTLE and samplers. The receiver is a quadrature-rate architecture scheme with 5 parallel quarter rate strong-arm samplers (one per sub-channel of the CNRZ-5 code) plus one sampler dedicated to the FCLK CDA, which is sampling data crossings. The self-referencing comparators provide immunity to common-mode noise across the 6 input data wires. T-coils (layers M8–M9) are used to compensate the parasitic input capacitance and achieve the desired input return loss. They are placed under the bumps to save area. The samplers have wide-range 7b low-speed offset-correction DACs with a resolution of 2.5mV. They can be used to reconstruct a receive eye at the output of the CTLE. The 6 Rx data wires and receive signal path circuitry are closely skew matched. A local Rx mini-PLL generates a local quarter-rate clock from the received FCLK; the Rx mini-PLL generates four clock phases (0, 90, 180, 270) using a complementary 4-stage fed forward Ring Oscillator (RO), which also provides an additional 315 degree phase for the CDA algorithm. The Rx mini-PLL maximizes bandwidth, and minimizes jitter transfer peaking.

A central low phase noise RO-based main-PLL generates a 3.125GHz clock from common RefClk frequencies. The main-PLL is a three-stage pseudo-differential RO supplied by a regulator from a 1.5V/1.8V supply. It has a programmable charge pump/loop filter combination with additional elements for fast start; digital lock detector and auto-restart circuits are included. The output of the main-PLL is multiplied by 2 in a Tx mini-PLL for the local four phase 6.25GHz transmitter clocks; the Tx mini-PLL is a type-II RO PLL. It rejects supply noise injected deterministic jitter through its wide bandwidth (250MHz). To properly track the jitter on the received signal, the Rx mini-PLL has a higher bandwidth (500MHz). The main-PLL is part of the common-block layout that also provides an analog test bus and temperature sensors. The Rx and Tx mini-PLLs are part of the Rx and Tx blocks, respectively, and include lock detectors.

The CNRZ-5 FCLK SerDes link is fabricated in a 28nm 1P9M process. Its layout is shown in Fig. 10.1.5 and the testchip, with a single instance of the SerDes, is shown in Fig. 10.1.7. The SerDes size is 1.503×0.4185mm^2 (Si) and has been laid out in a modular fashion such that additional 6 wire chords may be added to multiple chord SerDes utilizing one main-PLL. The testchip achieves 166Gb/s per mm die-edge (Rx plus Tx). A quad chord SerDes will offer 444Gb/s per mm (Rx plus Tx) with standard 150µm pitch bump technology. The FCLK SerDes supports data rates from 12.5 to 25Gbaud, internal Tx to Rx loop-back, PRBS31 and user pattern generation and verification, and a receiver eyescope. The test setup is a MCM with four testchip dies arranged to support 3 MCM channels of lengths 5mm, 12mm and 24mm; two of the four die provide Tx and Rx 8-wire breakouts respectively. Figure 10.1.6 shows a horizontal bathtub plot with an opening of 20ps at BER=1e-15, and an example of an eye diagram using the on-die eyescope at 25Gbaud over 12mm of MCM trace with 1.3dB loss; residual untracked jitter is 500fs (<0.015UI). Long runs show a vertical eye opening of 65mV at BER<1e-15 at nominal conditions (V_{dda}=V_{ddd}=1V, V_{ddh}=1.5V) and a power efficiency of 1.1pJ/b. The link achieves a power efficiency of 0.94pJ/b at 25Gbaud with BER<1e-15 on the same channel at V_{dda}=0.925V, V_{ddd}=0.8V, and V_{ddh}=1.4V.

Acknowledgements:
The authors would like to thank Graham Dring for supporting the layout of the chip, and EPFL, Upen Reddy and Jose Torres-Fretts at Cisco Systems Inc., and Gerard Mas at STMicroelectronics for their support with silicon testing.

References:
[1] B. Kim et al., "A 10 Gb/s compact low-power serial I/O with DFE-IIR equalization in 65 nm CMOS", *IEEE J. Solid State Circuits*, vol. 44, no. 12, pp. 3526-3538, Dec. 2009.
[2] M.-S. Lin et al., "An Extra Low Power 1 Tbit/s Bandwidth PLL/DLL-less eDRAM PHY using 0.3V Low Swing IO for 2.5D CoWoS Application," *IEEE Symp. VLSI Circuits*, pp. C16-C17, 2013.
[3] J. W. Poulton et al., "A 0.54 pJ/b 20 Gb/s Ground-Referenced Single-Ended Short-Haul Serial Link in 28nm CMOS for Advanced Packaging Applications," *IEEE J. Solid State Circuits*, vol. 48, no. 12, pp. 3207-3218, Dec. 2013.
[4] T. O. Dickson et al., "A 1.4 pJ/bit, power scalable 16x12 Gb/s source-synchronous I/O with DFE receiver in 32 nm SOI CMOS technology," *IEEE Custom Integrated Circuits Conf.,* pp. 5-10, Sep. 2014.
[5] A. Hormati et al., "Method and Apparatus for Low-Power Chip-to-Chip Communications with Constrained ISI-Ratio", U.S. Patent 9,100,232.

 978-1-4673-9467-3/16/$31.00

Reference	[1]	[2]	[3]	[4]	This work
Year	2009	2013	2012	2014	2015
pJ/bit	1.9	0.11	0.54	1.4	0.94
BW/pin (Gb/s)	8.9	1.1	20	6	20.83
Technology	45 nm	40 nm	28nm	32 nm	28nm
Signaling	D	SE	GRSE	D	CNRZ-5
Channel loss	≤ 20 dB	NA	≤ 1 dB	≤ 3 dB	≤ 3 dB
Substrate	SiCa	SiCa	MCM	Meg 6	MCM
Reach	≤ 40mm	≤ 1mm	≤ 4.5 mm	≤ 19mm	≤ 12mm
BER	1e-9	NA	1e-12	1e-9	1e-15

This work		
Block	Power (mW)	Area (mm²)
Tx	63.70	0.252
Rx	43.92	0.252
Common	9.96	0.126
Total	**117.58**	**0.629**
Ebit (pJ/bit)	**0.94**	

Figure 10.1.1: Comparison to prior work (SiCa=Silicon carrier, SE = single ended, D = differential, GRSE=ground referenced single-ended, MCM = MCM organic substrate, Meg 6 = Megtron 6) and power/area breakdown.

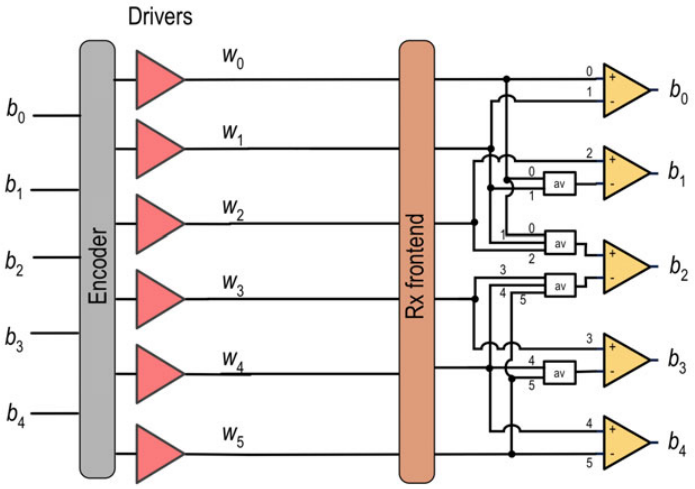

Figure 10.1.2: Functional diagram of the CNRZ-5: encoder, wires, multi-input comparators.

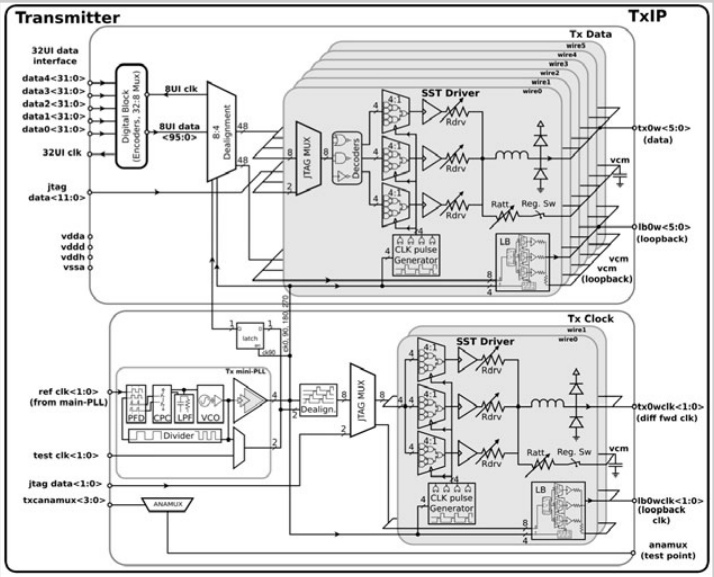

Figure 10.1.3: CNRZ-5 transmitter block/circuit diagram.

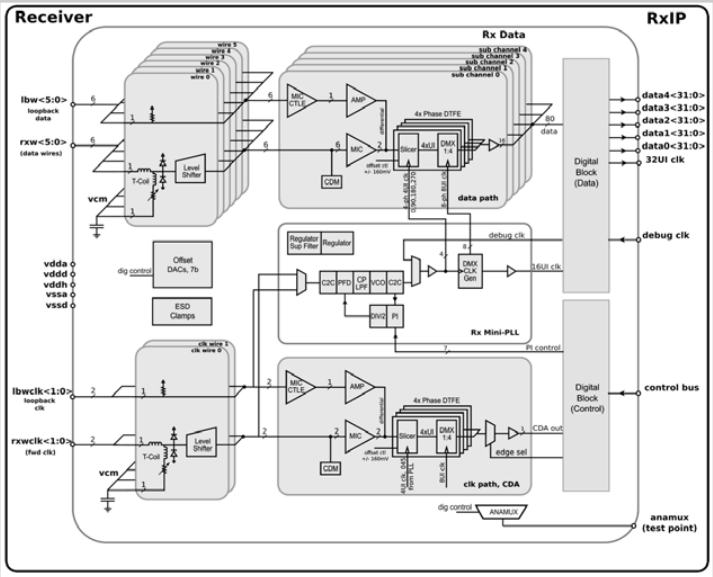

Figure 10.1.4: CNRZ-5 receiver block/circuit diagram.

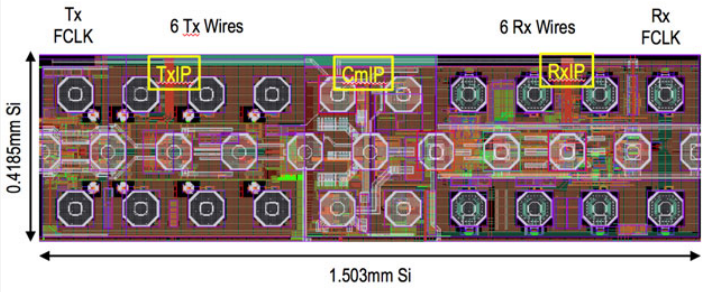

Figure 10.1.5: CNRZ-5 IP layout (CmIP = Common IP).

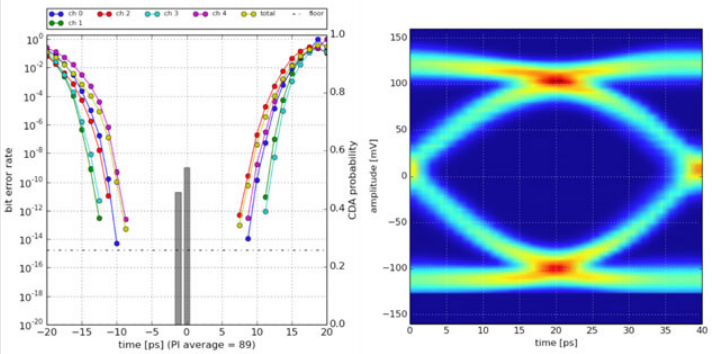

Figure 10.1.6: CNRZ-5 horizontal BER bathtub curve and example eye plot at the Rx for a 12mm MCM channel at 25GBd, PRBS31 pattern.

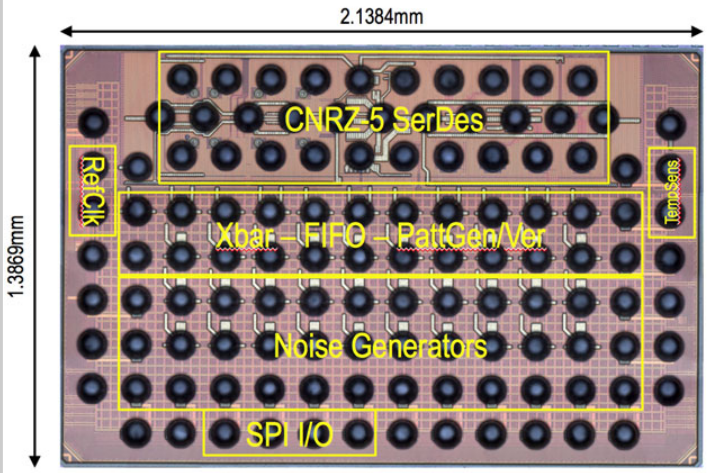

Figure 10.1.7: CNRZ-5 testchip.

978-1-4673-9467-3/16/$31.00