16.8 A 1.17pJ/b 25Gb/s/pin Ground-Referenced Single-Ended Serial Link for Off- and On-Package Communication in 16nm CMOS Using a Process- and Temperature-Adaptive Voltage Regulator

John M. Wilson[1], Walker J. Turner[1], John W. Poulton[1], Brian Zimmer[2], Xi Chen[2], Sudhir S. Kudva[2], Sanquan Song[2], Stephen G. Tell[1], Nikola Nedovic[2], Wenxu Zhao[1a], Sunil R. Sudhakaran[2], C. Thomas Gray[1], William J. Dally[2]

[1]Nvidia, Durham, NC; [2]Nvidia, Santa Clara, CA
[a]now with Broadcom, Irvine, CA

Toward the end of the Moore's-law era, increases in system complexity will rely more heavily on packaging technology. Systems will increasingly comprise multiple chips that must be linked by high-speed data channels carrying a substantial fraction of on-chip bandwidth. To take advantage of inexpensive organic packages and conventional printed circuit (PC) boards, data links that are both energy and pin efficient are needed. A link between neighboring packages is by far the more challenging application due to increased cross-talk, signal attenuation, and reflections from impedance discontinuities. The combination of signal integrity challenges and production margining requires increased amplitude, equalization, ESD protection, and PVT-tolerant circuit design techniques.

We describe a short-reach link that connects chips on the same package or on neighboring packages over conventional PC board channels. These links use single-ended signaling to conserve pins, operate at 25Gb/s/pin and 1.17pJ/b, and use a simple but robust clock-forwarding scheme to cancel jitter. The overall link design is shown in Fig. 16.8.1. Eight data lanes share a common forwarded clock. Data is transmitted on an in-phase clock (Iclk) while the forwarded clock is transmitted using a quadrature clock (Qclk). Data and clock lanes are designed to have closely matched channel delays. At the receiver, the forwarded clock is directly buffered and driven out to data samplers; a programmable delay in each data lane delays the received data by an amount matched to the insertion delay of the clock buffer chain.

The link uses ground-referenced signaling (GRS), described in [1]. GRS avoids many of the usual problems of single-ended signaling. The ground network, typically the lowest impedance network in a system, is used as the signaling voltage reference, eliminating the need for a precise matched reference at the receiver. Return currents flow only on ground, providing a clean line termination. Signals must be driven symmetrically positive and negative with respect to ground; GRS uses an output-multiplexed pair of charge pumps, shown in Fig. 16.8.2 (left), to produce these voltage levels. Since the charge-pump transmitter draws the same current from the supply regardless of data polarity, simultaneous switching noise is largely eliminated. A capacitively coupled auxiliary transmitter injects extra charge into the line on data transitions, providing tunable transmitter equalization across 5 strength settings. The receiver, shown in Fig. 16.8.2 (right), employs a pseudo-differential amplifier that performs four functions: 1) signal gain; 2) level-shift from the about-ground line signal up to CMOS levels; 3) linear EQ using gm-C active inductors in the amp load; and 4) single-ended to differential conversion. First-stage gain is about 6dB, with RX EQ boost set at maximum. A pair of CMOS inverters form the second amp stage and provide further gain of about 10dB. A common-mode feedback system keeps the first stage output centered on the inverter threshold. Termination is comprised of the adjustable Rterm resistor in parallel with $1/g_m$ of the input stage. Input offset is coarsely trimmed by differentially varying the current in the two arms of the amplifier, while Rtrim provides fine offset trim. The input amp consumes about 2.6mA.

Clocks are driven from a central PLL to the individual data and clock lanes using conventional CMOS buffers, each of which has digitally controlled insertion-delay and duty-factor trim. Clocks are distributed via top-level metal with skew < 1ps.

The link employs a novel power supply regulation scheme, shown in Fig. 16.8.3. A phase-locked loop locks a CMOS ring oscillator (RO) at 25GHz to a 1.56GHz reference clock; neither frequency is critical. The full-rate RO output is divided by 2 to provide in-phase and quadrature clocks Iclk/Qclk. Vreg_PLL is the VCO control input, set by a PMOS control element that regulates down from the external Vdd_IO supply. The I/O circuitry operates from a second supply, Vreg, whose digital regulator uses Vreg_PLL as a reference voltage. This voltage is set in the PLL so that an exemplar CMOS circuit (the RO) operates at a fixed rate independent of PVT, thus the I/O circuitry, which operates on a supply voltage that is nominally equal to Vreg_PLL, also operates at a fixed rate independent of PVT. At the expense of regulator losses, this arrangement varies the internal supply voltage to flatten current consumption and circuit speed across PVT (see Fig. 16.8.3), thereby saving the considerable power that would otherwise be needed to provide margin. The flattening of current across PVT also aids in satisfying electro-migration requirements.

An important aspect of power efficiency is energy consumed when link traffic is variable. This link features a "pause" mode that reduces power dissipation by 75%, while providing fast entry and exit times (<5ns). When "pause" is needed, the last data bit is transmitted on a low-to-high transition of Iclk, and the clock is then held high for an extended period. Since Iclk is not toggling, data transmitters are effectively powered down. The clock lane continues to run on Qclk, but sends a stream of 1s. At the receiver, a simple analog detector observes that the clock has not transitioned for several cycles and turns off current consumers (mainly the data input amps). The PLLs continue to run both at the transmitter, to keep the clock forwarding charge pumps running, and at the receiver, to maintain Vreg. At the end of "pause", the clock begins toggling again, 2-5 (programmable) invalid words are transmitted over the data links while the receiver powers up, and then live data can begin transmission.

The link is fabricated in a TSMC 16nm FINFET 12-metal process, and operates from an external power supply voltage of 0.95V. It comprises a transceiver "brick", centered in a 4×5 bump array on a 150µm pitch (Fig. 16.8.7). Overall area is 686µm×565µm, where the I/O brick occupies 81,406µm^2. Figure 16.8.6 contains the circuit area breakdown of the system sub-blocks.

The link is intended to operate over short channels in multi-chip and multi-package environments. A model of the experimental package-to-package channel is shown in Fig. 16.8.4. It is 80mm long and has attenuation of -8.5dB at Nyquist. Individual lanes are matched in length to the clock lane within ±4ps. Crosstalk for the package-to-package system is held below -31.6dB for the sum of all eight aggressors. This is achieved by: 1) using a checkerboard pattern of signal bumps with ground/power bumps; 2) stripline routing in the package and PC board; 3) placing grounded shields between traces in the PC board only; and 4) placing PC board routes in the next-to-bottom PC board layer to minimize any via stub effects. T-coils are used on each lane to improve back-match, with a robust ESD protection device included. Link attenuation is overcome using a combination of transmitter and receiver linear EQ. The data and clock eyes shown in Fig. 16.8.2 were probed at the output of a transmitter. The effects of crosstalk, due to transitioning of data lanes, can be seen as voltage noise at the minimum and maximum values of the clock waveform.

Figure 16.8.5 shows the bathtub curves for the package-to-package link and a 10mm on-package link. The on-chip phase interpolator provides 1.33ps steps, and the 8 data lanes are matched to within a 6-step window, then subsequently trimmed to 3 steps using the on-chip data delay elements. At BER = 10^{-15}, the aggregate eye opening is 0.42UI for the off-package channel and 0.77UI on-package.

Figure 16.8.6 shows the energy per-bit breakdown for the link and a comparison with other work in the general area of short-reach links. The presented link normalizes operation across PVT and passes 110°C 10-year EM requirements to achieve 25Gb/s, 1.17pJ/b, single-ended signaling suitable for both on-package and high-loss off-package channels.

Acknowledgements:
This research was, in part, funded by the U.S. Government under the DARPA CRAFT program. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.
DISTRIBUTION A. Approved for public release: distribution unlimited

References:
[1] J. Poulton, et al., "A 0.54pJ/bit 20Gb/s Ground-Referenced Single-Ended Short-Haul Serial Link in 28nm CMOS for Advanced Packaging Applications", *IEEE JSSC*, vol. 48, no. 12, pp. 3206-3218, Dec. 2013.
[2] A. Shokrollahi, et al., "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," *ISSCC*, pp. 182-183, Feb. 2016.
[3] J. Song, et al., "A 1-V 10-Gb/s/pin Single-Ended Transceiver with Controllable Active-Inductor-Based Driver and Adaptively Calibrated Cascaded-Equalizer for Post-LPDDR4 Interfaces," to appear in *IEEE TCAS-I*, accepted June 2017.

 978-1-5090-4940-0/18/$31.00

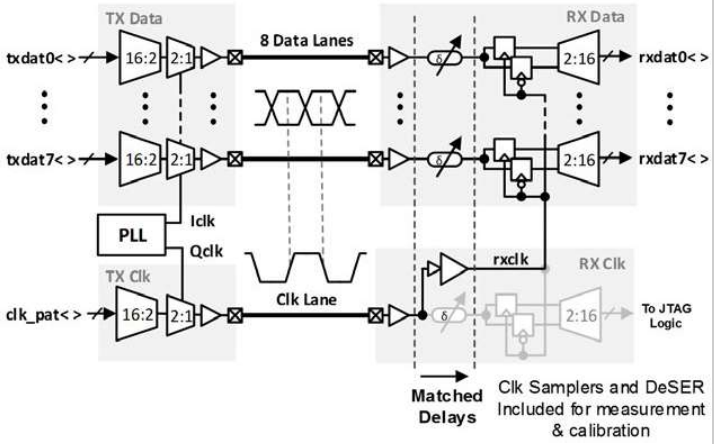

Figure 16.8.1: GRS link overview.

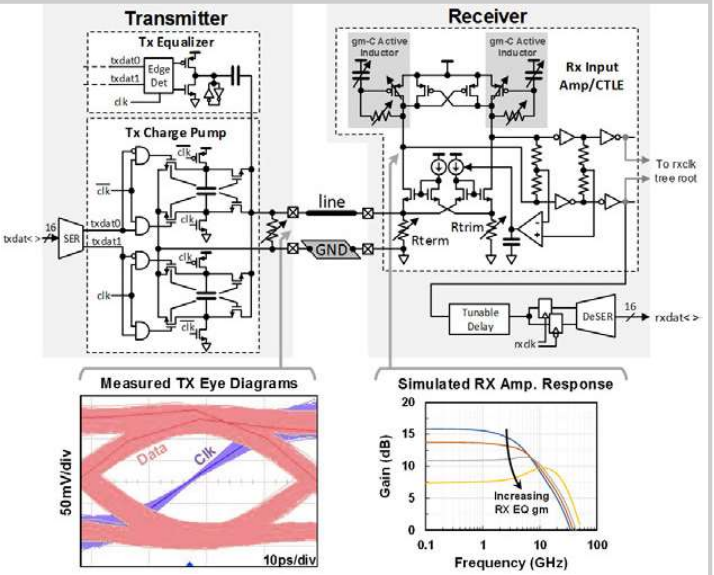

Figure 16.8.2: Transmitter and receiver schematics.

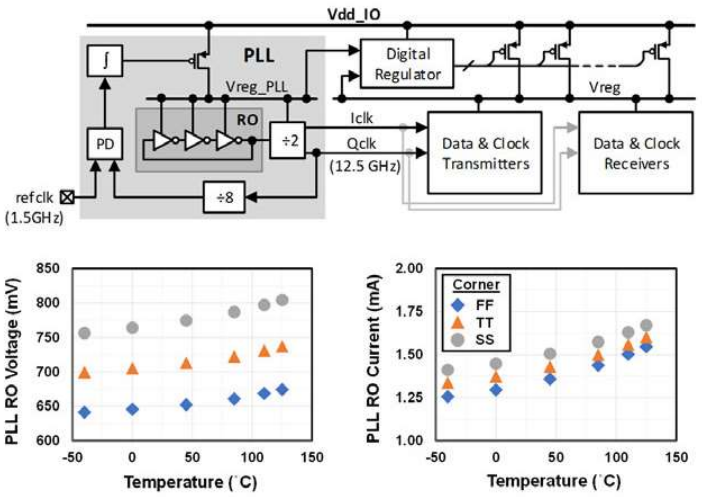

Figure 16.8.3: PVT tolerant design: PLL & digital regulator schematic and ring-oscillator voltage and current across corner cases.

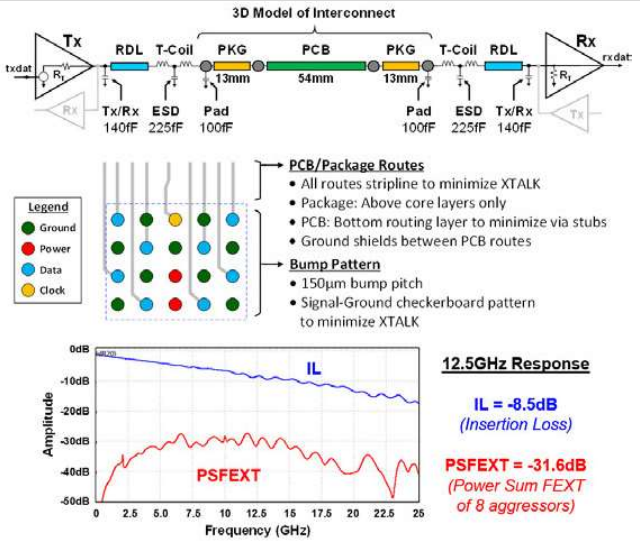

Figure 16.8.4: PCB channel model, bump pattern, and complete channel response.

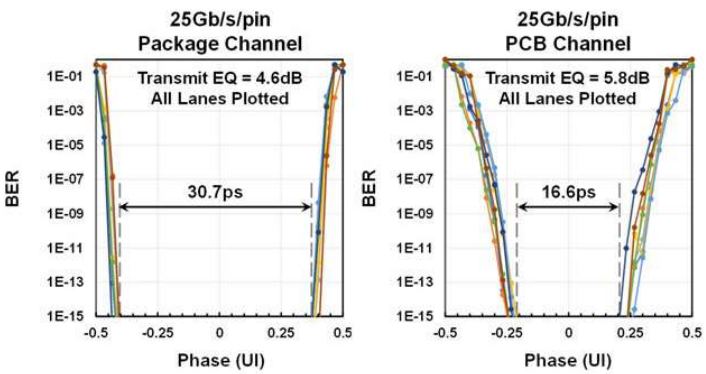

Figure 16.8.5: 25Gb/s PRBS-31 bathtub curves for 10mm package channel & PCB channel.

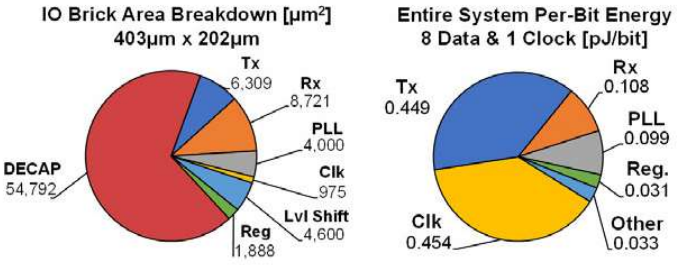

	[3]	[2]	[1]	This Work
Signaling	Single-Ended	CNRZ-5	Single-Ended	Single-Ended
Data Rate / Pin [Gb/s]	10.0	20.83	20.0	25.0
Reach	100mm	12mm	4.5mm	80mm
Type of Packaging	MCM or PCB	MCM	MCM	MCM or PCB
CMOS Technology	65nm	28nm	28nm	16nm
Energy / Bit [pJ/bit]	4.18	0.94	0.54	1.17

Figure 16.8.6: Energy and area breakdown for this work & comparison to prior work.

16

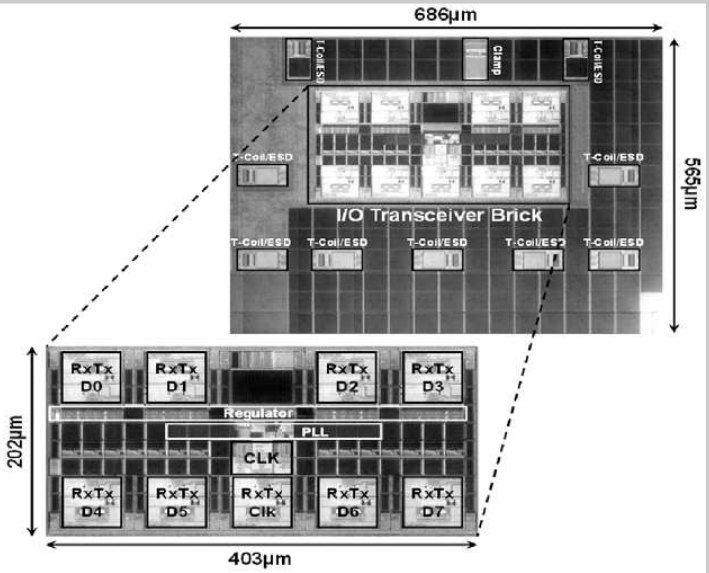

Figure 16.8.7: Die photo.

978-1-5090-4940-0/18/$31.00