

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

John W. Poulton, *Fellow, IEEE*, John M. Wilson^{ID}, *Senior Member, IEEE*, Walker J. Turner, *Member, IEEE*, Brian Zimmer, *Member, IEEE*, Xi Chen, *Member, IEEE*, Sudhir S. Kudva, Sanquan Song, *Member, IEEE*, Stephen G. Tell, *Member, IEEE*, Nikola Nedovic, *Member, IEEE*, Wenxu Zhao^{ID}, Sunil R. Sudhakaran, C. Thomas Gray, *Senior Member, IEEE*, and William J. Dally, *Fellow, IEEE*

Abstract—This paper describes a short-reach serial link to connect chips mounted on the same package or on neighboring packages on a printed circuit board (PCB). The link employs an energy-efficient, single-ended ground-referenced signaling scheme. Implemented in 16-nm FinFET CMOS technology, the link operates at a data rate of 25 Gb/s/pin with 1.17-pJ/bit energy efficiency and uses a simple but robust matched-delay clock forwarding scheme that cancels most sources of jitter. The modest frequency-dependent attenuation of short-reach links is compensated using an analog equalizer in the transmitter. The receiver includes active-inductor peaking in the input amplifier to improve overall receiver frequency response. The link employs a novel power supply regulation scheme at both ends that uses a PLL ring-oscillator supply voltage as a reference to flatten circuit speed and reduce power consumption variation across PVT. The link can be calibrated once at an arbitrary voltage and temperature, then track VT variation without the need for periodic re-calibration. The link operates over a 10-mm-long on-package channel with -4 dB of attenuation with 0.77-UI eye opening at bit-error rate (BER) of 10^{-15} . A package-to-package link with 54 mm of PCB and 26 mm of on-package trace with -8.5 dB of loss at Nyquist operates with 0.42 UI of eye opening at BER of 10^{-15} . Overall link die area is $686\ \mu\text{m} \times 565\ \mu\text{m}$ with the transceiver circuitry taking up 20% of the area. The transceiver's on-chip regulator is supplied from an off-chip 950-mV supply, while the support logic operates on a separate 850-mV supply.

Index Terms—Clock forwarding, dynamic voltage scaling, ground-referenced signaling (GRS), multi-chip modules (MCM), single-ended (SE) signaling.

Manuscript received April 24, 2018; revised July 13, 2018 and August 28, 2018; accepted September 23, 2018. Date of publication November 9, 2018; date of current version January 14, 2019. This paper was approved by Guest Editor Yohan Frans. This work was supported by the DARPA CRAFT Program, U.S. Government. Distribution Statement A. Approved for Public Release, Distribution Unlimited. (*Corresponding author: John M. Wilson.*)

J. W. Poulton, J. M. Wilson, W. J. Turner, S. G. Tell, and C. T. Gray are with NVIDIA, Inc., Durham, NC 27712 USA (e-mail: dr.john.wilson@gmail.com).

B. Zimmer, X. Chen, S. S. Kudva, S. Song, N. Nedovic, S. R. Sudhakaran, and W. J. Dally are with NVIDIA, Inc., Santa Clara, CA 94305 USA.

W. Zhao is with Broadcom Inc., Irvine, CA 92618 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSSC.2018.2875092




Fig. 1. Very-short-reach (on-package) and short-reach (package-to-package) link examples.

I. INTRODUCTION

AS THE end of the Moore's Law era approaches, increases in system complexity can no longer rely on the consistent scaling of integrated-circuit feature sizes that the industry has depended on for the last half century. Systems will need to rely more heavily on advanced packaging techniques such as multi-chip modules (MCM), where systems too large for a single die will consist of multiple semiconductor dice that are interconnected, as shown in Fig. 1, using on- and off-package high-speed interconnects. These short-reach links are a key enabler of these scalable systems, which require interconnects capable of supporting data rates approaching on-chip bisection bandwidth while consuming a small fraction of total device power. Similar chip-to-chip links are also a crucial component for optical long-haul communications, where processors and networking switches are interconnected to special purpose electrical/optical (E/O) interface chips. Even in traditional memory interconnects, energy-efficient short-reach links are continually developing to meet future market demands.

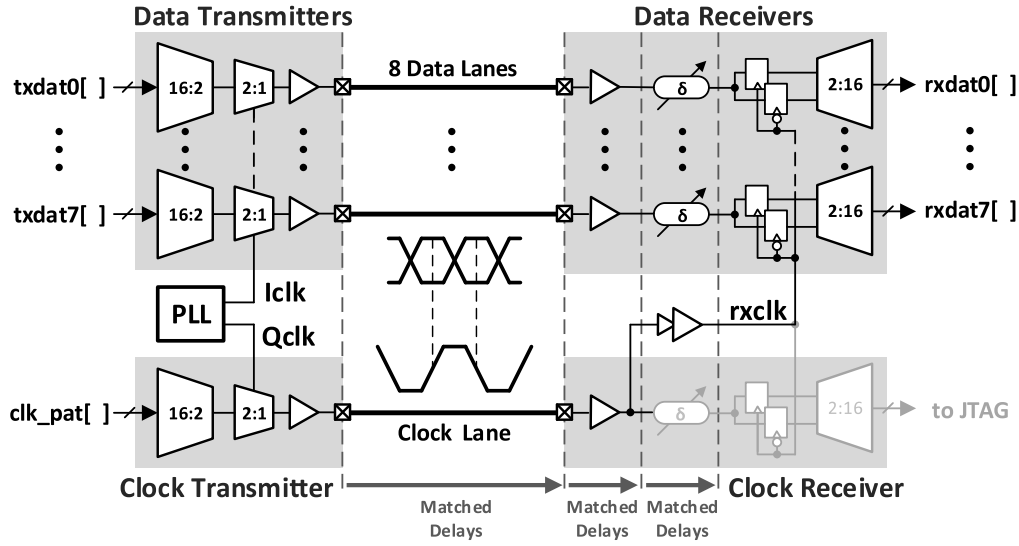


Fig. 2. Overview of the GRS link architecture.

Current state-of-the-art research on short links is mainly focused on E/O interfaces, which operate at 56 Gb/s to meet the OIF-CEI-04.0 standards [2]–[7]. These interconnects are differential and typically use PAM-4 signaling to meet the high-bandwidth requirements. These links operate with 4–6 pJ/bit energy efficiencies and can deal with up to about –20 dB of signal attenuation. Recent advances in memory interconnects include the introduction of High Bandwidth Memory [8], which offers good energy efficiencies of 1–2 pJ/bit but poor pin efficiencies since data rates are limited to 2–4 Gb/s/pin. A single-ended (SE) memory link was also implemented to support 10 Gb/s at about 5 pJ/bit [1]. Progress on very-short-reach links could also be applied to short package-to-package channels. A 16-Gb/s interconnect was developed in [9] to signal across a –17-dB channel while consuming 4.5 pJ/bit (not including the clocking system). A 56-Gb/s interconnect in [10] operates over a –8-dB channel while consuming 2.25 pJ/bit (including clocking). A commercial short-reach link described in [11] uses “chord” signaling, where 5 bits are transmitted over six wires at 25 Gbaud (20.83-Gb/s/pin effective bandwidth) and achieves about 0.93-pJ/bit efficiency for –1.2-dB low-loss on-package channels. None of these systems simultaneously meet the bandwidth, energy-efficiency, and pin-efficiency requirements for our target multi-GPU MCM application.

This paper describes a high-speed link intended for scalable MCM systems that can also support short-reach package-to-package communication. To conserve signal pins, the link uses SE ground-referenced signaling (GRS), a technique introduced in [12] and further described in [13] and [14]. The link operates at 25 Gb/s/pin with a 1.17 pJ/bit energy efficiency, including all clock, regulator, and SerDes logic power. It uses a simple but robust clock-forwarding scheme to cancel jitter without the need for a clock recovery system at the receiver. A co-design of the link circuitry and channel takes full advantage of the ability to match the insertion delays of multiple data/clock lanes within the short-reach interconnect. The low attenuation that can be achieved in a co-designed channel is leveraged to achieve a simple,

robust, and low-power design. Section II describes the overall link design, including the clock-forwarding scheme and a novel supply regulator that compensates for PVT variation. Section III describes the circuit design details of the GRS transmitter and receiver. Section IV describes the experimental results of a production-quality GRS link that supports both on-package and package-to-package communication. Section V concludes the paper.

II. GRS LINK DESIGN

A. Clock Forwarding Architecture

The overall GRS link architecture is shown in Fig. 2, where a forwarded clock is transmitted in quadrature with eight bundled data lanes. Though omitted from Fig. 2 for clarity, both ends of the test-chip link are organized as transceivers, where the transmitters and receivers are identical for all clock and data lanes. The forwarded clock receiver retains its de-serializer to enable the phase and duty-cycle calibration.

The transmitters consist of 16:2 serializers followed by 2:1 multiplexing stages and line drivers. A ring-oscillator (RO) phase-locked loop (PLL), described below, produces 12.5-GHz in-phase (**Iclk**) and quadrature (**Qclk**) clocks at the transmit end of the link. The last 2:1 multiplexing stage of the data transmitters is clocked by **Iclk** while the forwarded clock multiplexer is timed by **Qclk** to achieve quadrature data-to-clock timing relationship. The **Iclk** and **Qclk** distribution buffers are designed to have identical propagation delays and power supply sensitivities.

To preserve the quadrature clock-to-data timing at the receiver input, all data and clock channels are routed with near identical lengths. This ensures the channels are delay-matched to a tolerance much smaller than 1 UI.

At the receive end of the link, the data and clock signals are amplified to CMOS levels by identical amplifiers. The clock receiver amplifier drives the root of the receiver clock (**rxclk**) buffer to provide sampling clocks at each data receiver lane. In the data lanes, adjustable CMOS delay elements are inserted between receiver amplifiers and samplers with an insertion delay (δ) designed to match that of the **rxclk** buffer at




Fig. 3. Overview of GRS.

the mid-point of the adjustment range. The tunable delay is adjusted by trimming the internal fan-out of a series of CMOS stages, so that the power supply sensitivity of δ tracks that of the **rxclk** buffer. Duty factor is trimmed by adjusting the relative P-to-N transistor sizes in the delay stages.

The use of near-identical circuitry for transmitting and receiving the forwarded clock and data signals, in addition to delay-length matching all channels, ensures the data, and clock insertion delays are matched across the entire link and track across variations in PVT. This simple mechanism cancels nearly all sources of jitter, including power-supply-induced jitter, over a frequency range approaching the link bandwidth. Additionally, the jitter requirements of the PLL are greatly relaxed, and any timing variations are exported within the relative timing of the clock and data signals.

B. Ground-Referenced Signaling

The link employs GRS to avoid many of the difficulties with conventional SE signaling. Referring to Fig. 3, GRS operates by transmitting signals that toggle symmetrically around the ground potential (0 V), as described in [12]–[14]. A pair of switched-capacitor charge pumps (*Pump0* and *Pump1*) functions as the GRS transmitter, which operates at bit-rate to implement the symmetric signal voltage generator, 2:1 data multiplexer, and line driver. GRS signaling circuits are described in more detail in the next section.

This arrangement avoids three of the principal difficulties with SE signaling.

- 1) *Simultaneous Switching Noise*: The single largest noise source in conventional SE systems. A GRS charge-pump transmitter consumes nearly constant current from the power supply on each cycle, regardless of data polarity, so simultaneous switching noise is nearly eliminated.
- 2) *Reference Voltage Generation*: SE systems require a reference voltage that is nominally halfway between the HI and LO signal levels at the receiver. Since GRS uses GND as the signal reference (typically the lowest impedance and most robust network in a system), the receiver requires no internal reference, thereby eliminating a major noise source.
- 3) *Frequency-Dependent Termination*: The return currents in conventional SE signaling systems flow through both GND and power supply networks, including the bypass capacitors between them, which introduces frequency-dependent return paths and degrades the line termination quality. GRS simplifies the return loop by ensuring

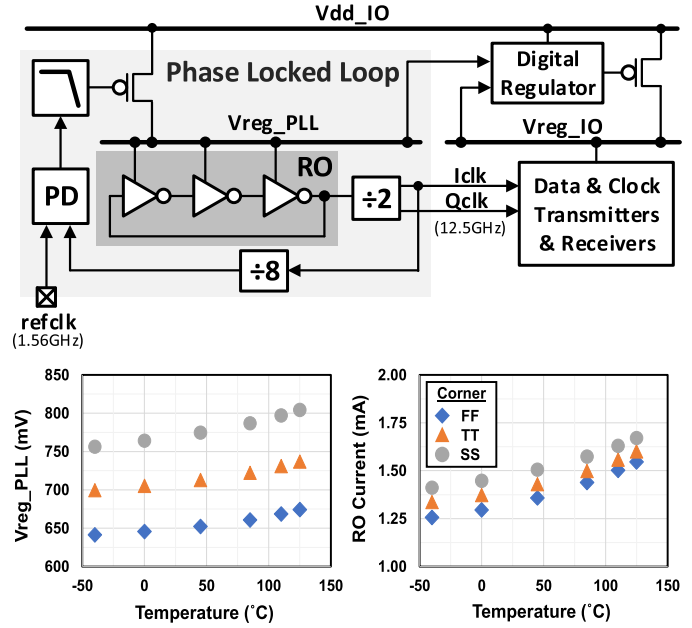


Fig. 4. Dynamic voltage scaling power regulation scheme.

the return current flows only in the ground network, allowing for a simple high-quality termination.

C. Dynamic Voltage Scaling Regulation Scheme

The link employs the power supply regulation scheme shown in Fig. 4, which uses a novel digital regulator described in detail in [15]. A PLL locks a CMOS RO at 25-GHz to a 1.56-GHz reference clock; neither frequency is critical. The full-rate RO output is divided by 2 to provide in-phase and quadrature-phase clocks (**Iclk/Qclk**). The VCO control input (**Vreg_PLL**) is set by a PMOS pass element that regulates the RO supply voltage down from the external supply (**Vdd_IO**). The I/O circuitry operates from a second regulated supply (**Vreg_IO**) whose 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. Therefore, the I/O circuitry, which operates on a supply voltage that is nominally equal to **Vreg_PLL**, also operates at fixed rate independent of PVT. At the expense of regulator losses, this arrangement varies the internal I/O supply voltage to flatten the circuit speed across PVT and reduce current consumption variation across process corners (see Fig. 4 insets), thereby saving considerable power that would otherwise be dissipated after providing sufficient margin across corner cases. The reduction in the current consumption variation across process corners also aids in satisfying electromigration requirements. This regulation scheme is employed at both ends of the link, where the PLL at the receive end is used solely to set the I/O supply voltage, since the generated **Iclk** and **Qclk** signals are not used in the receiver.

D. Fast Entry/Exit Pause Mode

While energy-efficient operation is a requirement for a low-power link, it is not sufficient in cases where the traffic carried by the link is bursty. To deal with this problem, the link enters

a “pause” mode whenever transmission of data is not needed. In “pause” the circuitry consumes very low power while providing very fast entry and exit times from “pause” to “active”. To enter the “pause” mode, the forwarded clock is held HI after transmission of the last data bit and remains in this state for the duration of the pause event. The internal **Iclk** signal stops toggling and is held HI, which drops the clock distribution and data transmitter power to zero, except for the forwarded clock circuitry, which must continue to drive its line HI. Because the forwarded clock is held HI, the recovered clock signal at the receive end of the link (**rxclk**) stops toggling, so clock distribution power at the receiver also drops to zero. A simple analog CMOS circuit with asymmetrically skewed pull-up and pull-down strengths, detects that **rxclk** has paused and signals all current-consuming circuits in the receivers to shut down. The remaining active circuitry includes the PLLs at both ends of the link to keep the Tx **Qclk** toggling and maintain the internal **Vreg_IO** voltages, as well as the forwarded-clock transmitter to hold the forwarded clock signal HI. To exit the “pause” mode, the forwarded clock begins toggling, while **Iclk** is re-started simultaneously. The toggling of the forwarded clock causes **rxclk** to resume at the receive end of the link, and the CMOS pause-mode detector powers up the receiver circuitry. After a small and programmable number of parallel clock times, “live” data transmission can resume. The pause mechanism reduces power to about 25% of active power, while entry and exit times are held to less than 5 ns.

E. Link Calibration

An important design objective was eliminating the need for periodic re-calibration. To make this work correctly, it must be possible to perform calibration at the beginning of operation, including signal offset cancellation, clock delay matching, and duty factor cancellation. Subsequently, the operating parameters must track across supply voltage and temperature variation with sufficient accuracy to maintain link margin. The PLL-referenced power supply regulation scheme described earlier, along with many circuit design techniques, support this “calibrate-once” policy as demonstrated with the experimental results outlined in Section IV (see Fig. 15).

Most calibration steps make the use of a special operational mode in the receiver data samplers. In this mode, the receiver sampling clock **rxclk** is swapped with a locally generated random clock that is uncorrelated with the data. A data pattern with an established ratio of 1’s and 0’s is transmitted to the receiver, sampled by the random clock, and accumulated in counters. Finally, the ratio of 1’s to 0’s is measured, and the results are used to drive the calibration loop. To remove the receiver input offset, for example, a continuous duty-factor-free 11001100 pattern is transmitted, the randomly sampled data accumulated, and the calibration loop drives the receiver offset tuning mechanism to obtain a 1:1 ratio from 1’s to 0’s. Next, a repeating 1010 pattern is transmitted, and the calibration loop drives the duty-factor offset tuning for the received clock signal. A similar technique could be used to tune equalizer settings, though in very short-range links with small channel attenuation, the EQ adjustment can be determined and set at design time.

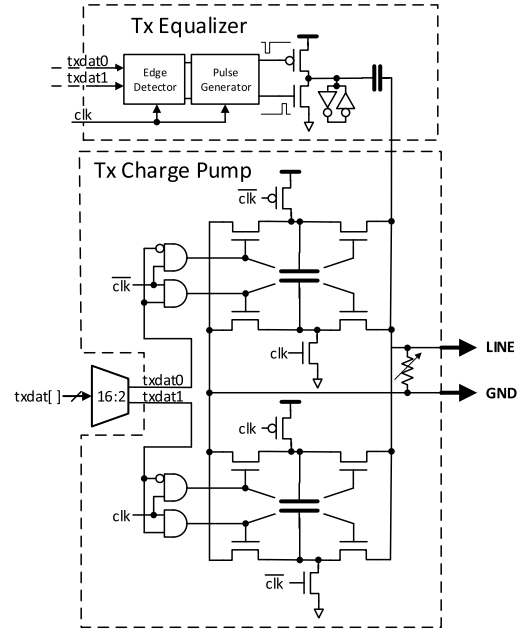


Fig. 5. GRS transmitter schematic.




Fig. 6. Simplified transmitter circuit and approximate line voltage during the drive phase.

III. GRS CIRCUIT DESIGN

A. GRS Transmitter

Fig. 5 shows a schematic of the GRS transmitter, comprising the main charge-pump-based line driver and an auxiliary equalizing transmitter. 16 bits of parallel data are serialized to an even and odd half-bit-rate data pair (**txdat0** and **txdat1**), with final 2:1 multiplexing performed in the charge-pump and equalizer stages. The line driver has two identical charge-pump circuits that operate on opposite phases of the half-bit-rate clock. The upper charge pump pre-charges a storage capacitor to the supply rail when the clock is HI; when clock goes LO, the charge stored in the capacitor is driven into the line, developing a positive line voltage when even data **txdat0** = HI and a negative voltage when **txdat0** = LO. The lower charge pump performs the same operation but on opposite clock phases, driving the odd data (**txdat1**) onto the line.

The equalizer performs additive edge boosting through an ac-coupled voltage mode driver that drives a current impulse into the line during edge transitions to enhance the speed of the output signal transition. An edge detector inside the equalizer drives a 1-UI pulse generator whenever the current data bit

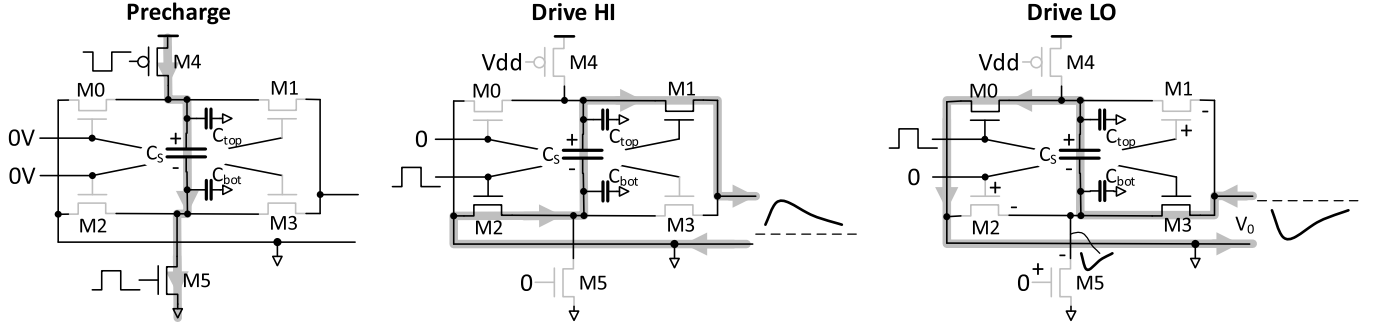


Fig. 7. Details of current flow during pre-charge, drive HI, and drive LO.

differs from the previous one. This is performed by an XOR of even and odd data bits [14], thereby ensuring that the output is only driven during edge transitions and is left in a high-Z state during idle periods to prevent modulation of the transmitter return impedance. Cross-coupled inverters are added to the inboard node of the coupling capacitor to operate as high-impedance keeper cells that hold the node voltage at either supply rail during long periods of inactivity. The equalizer is divided into a number of digitally enabled segments, so that the amount of overdrive can be adjusted at run time to compensate for channel attenuation. Since the charge-pump transmitter and transmit equalizer are comprised of CMOS switches, it is easy to match the delays of their data paths at design time. This approach, combined with the adaptive power supply that flattens circuit delay across PVT, ensures delay match across PVT. The transmitter output is coupled to the line through a T-coil (not shown) that compensates for the parasitic capacitance of the transceiver circuits, ESD clamp, I/O pad, and redistribution layer (RDL) routes.

The behavior of the charge-pump transmitter can be analytically approximated with the help of the simplified circuit of Fig. 6, consisting of the storage capacitor (C_S), pre-charged to an initial voltage (v_{ini}), in series with a switch with switch resistance (R_S). The shunt capacitance at the transmitter output terminal and the line impedance are modeled as C_O and R_O , respectively. The line voltage (V_O) can be shown to be

$$V_O(s) = \frac{\beta v_{ini}}{s^2 + (\alpha + \beta + \gamma)s + \alpha\gamma} \quad (1)$$

where $\alpha = 1/R_S C_S$, $\beta = 1/R_S C_O$, and $\gamma = 1/R_O C_O$. The system has two poles that govern the time evolution of the system such that

$$\pm p_{1,2} = \frac{\alpha + \beta + \gamma}{2} \left(1 \pm \sqrt{1 - \frac{4\alpha\gamma}{(\alpha + \beta + \gamma)^2}} \right). \quad (2)$$

Assuming $p_2 > p_1$

$$v_O(t) = \frac{\beta v_{ini}}{(p_2 - p_1)} (e^{-p_1 t} - e^{-p_2 t}). \quad (3)$$

For the 25-Gb/s transmitter reported here, $R_S \approx 2R_O = 80\Omega$ and $C_S \approx 2C_O = 400$ fF, so $\alpha = 1/2\beta = 1/4\gamma$. Time t_{max} at which the output pulse is at maximum voltage, as well as

the maximum voltage (V_{max}), can be calculated as

$$t_{max} = \frac{1}{(p_2 - p_1)} \ln \left(\frac{p_2}{p_1} \right) \quad (4)$$

$$V_{max} = \frac{\beta v_{ini}}{p_2} \left(\frac{p_2}{p_1} \right)^{\frac{-p_1}{p_2 - p_1}}. \quad (5)$$

For a 750-mV I/O supply voltage (v_{ini}) and the assumptions above for the time constants, $V_{max} = 0.243 \times 750$ mV = 183 mV and $t_{max} = 12.9$ ps are very close to the results of detailed circuit simulations. When the link is sending a stream of 1's, the line voltage is the same at the beginning and end of each UI (V_H) while the peak voltage V_{max} is somewhat higher. The formalism above can be used to calculate these voltages. With the same assumptions as above, these voltages can be shown to be $V_H = 0.082v_{ini}$ (62 mV), $V_{max} = 0.199v_{ini}$ (150 mV), and $t_{max} = 0.238$ UI.

The calculated values are close to the observed and simulated voltages from the GRS transmitter. The roughly 90 mV of ripple on the transmitted signal that occurs when sending a stream of 1's or 0's may appear to be a concern, but this ripple is at twice the Nyquist frequency, and heavily filtered by the channel attenuation. With just a few dB of attenuation at Nyquist, the signal eye at the receiver input shows only a small residual ripple, and its maximum occurs very near the eye center (see Fig. 12).

The analytical treatment developed above is only an approximation since the storage capacitor and switches are non-ideal in several ways. Fig. 7 shows a more detailed schematic of the charge-pump line driver during pre-charge and when driving the line HI and LO.

During pre-charge, the H-bridge transistors M0–M3 are all off while M4 and M5 turn on to charge the storage capacitor C_S to V_{dd} . When driving HI, the pre-charge transistors are turned off and M1 and M2 are enabled. Current flows from the capacitor into the line to generate a positive voltage pulse. When driving LO, transistors M0 and M3 are turned on so C_S is connected so that current is drawn from the line, creating a negative voltage pulse.

There are several asymmetries in these phases that tend to make positive and negative drive voltages differ. First, the parasitic capacitors C_{TOP} and C_{BOT} aid during the positive drive but oppose during the negative drive. This effect is typically less than 10% of the drive voltage and can be compensated by choosing the sizing of M0–M3 appropriately. Second, driving

LO not only drives the line output below ground, but also the bottom terminal of the storage capacitor C_S . The gates of M1, M2, and M5 are at 0 V while the sources of these transistors see a negative voltage excursion, causing the transistors to turn on. For the 100–200-mV amplitudes that can practically be achieved in a GRS transmitter of this type, transistors M1, M2, and M5 enter the weak inversion region of operation where leakage is large enough to introduce drive asymmetry. This leakage diverts current away from the line during the negative drive, further exacerbating drive asymmetry. The leakage problem ultimately limits the drive amplitude of a GRS transmitter, but for short-reach links, a few 100 mV of drive amplitude is sufficient. The problem can be mitigated by choosing the appropriate MOSFET device types for the charge-pump transistors; the best choice is standard-threshold devices for M0–M3 (so the line drive transistors track across global corners), a low-threshold PMOS for M4 (to speed up pre-charge), and a high-threshold device for M5 (since this device is the largest and leakiest NMOS in the pump).

Finally, non-overlap between pre-charge and drive phases must be guaranteed. For the negative drive, M3 is bootstrapped by its negative-going source voltage, so the time it takes to fully turn on is slightly delayed and ensures non-overlap between phases. For positive drive, the gate-drive voltages on M1 and M2 must be slightly delayed to ensure that pre-charge is completely extinguished before initiating line drive. Otherwise, the trailing edge of the pre-charge current aids in driving a positive current into the line, creating further asymmetry. Referring to Fig. 5, the AND gates that logically combine **clk** and even/odd data (**txdat0** and **txdat1**) are modified so that the “enable” clock is delayed by one inverter delay, while “disable” is immediate; this mechanism is sufficient to overcome the overlap problem.

Asymmetric drive leads to offset in the transmitter output signal, which is indistinguishable from the receiver input offset in the calibration scheme we are using. While both forms of offset can be compensated by tuning the receiver’s input-referred offset voltage, if the transmitter output offset varies with temperature and supply voltage, it will be impossible to track output offset without re-calibration. Therefore, it is critical to remove transmitter output offset at design time to the highest degree possible. While the PLL-referenced voltage regulation scheme removes much of the PVT-dependent output offset, the use of thin-oxide MOS capacitors as the charge-pump storage capacitors introduces unacceptable temperature variation in the transmitter output offset. In the design described here, most of the storage capacitor (C_S) is implemented as metal–oxide–metal (MOM) capacitors, which eliminates much of the temperature variation in transmitter output offset.

Charge-pump style transmitters are quite general purpose. It is easy, for example, to perform higher multiplexing ratios than 2:1 to attain higher speeds. Transmitter elements can be combined to perform equalization for de-emphasis, or to generate PAM-4 modulated outputs. When driving the line, a charge-pump transmitter is “floating” (i.e., not referred to either power supply), and this feature may be useful in situations where it is desirable to allow the receiver to set the

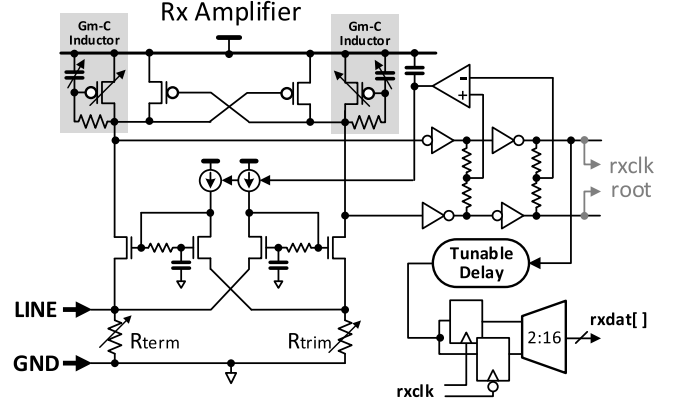


Fig. 8. GRS receiver schematic.




Fig. 9. Input amplifier frequency response versus G_m – C adjustments.

common-mode (CM) voltage in a differential signaling system. The downside of charge-pump transmitters is that the range of operating speed is somewhat limited, since lower-bit-rate pumps require larger values of C_S storage capacitors to generate the same line voltage.

B. Termination Scheme

The return impedance of the charge-pump transmitter is the sum of the switch resistance and the effective resistance of the switched-capacitor circuit in the pump, given by $1/C_S f$. With $C_S = 400$ fF and $f = 25$ GHz (there are two pumps in parallel), this impedance is about 100 Ω . Including switches, the total return impedance is about 140 Ω . A detailed circuit simulation of the transmitter gives an average equivalent impedance very close to this estimate. While the transmitter’s average impedance is well controlled over PVT, thanks to the supply regulation scheme, the switch resistance varies considerably during a drive event, as the transistors pass through saturation and triode modes of operation. In any case, both the average and instantaneous charge-pump impedances are significantly higher than typical channel impedance, so the transmitter can be viewed as a current source.

In cases where the channel response is smooth with no significant impedance discontinuities near the transmitter, no back-match is needed, and the transmitter can drive the line from its own source impedance to provide the largest possible signal into the channel. For cases where the line must be back-matched to absorb reflections from discontinuities in a more complicated channel, an adjustable parallel termination resistor is included to improve return loss. This same termination resistor is repurposed as the incident wave termination at the receiver input.




Fig. 10. Die photograph with Tx/Rx signal (S), clock (C), ground (G), and power (P) ball locations. All lanes are transceivers.

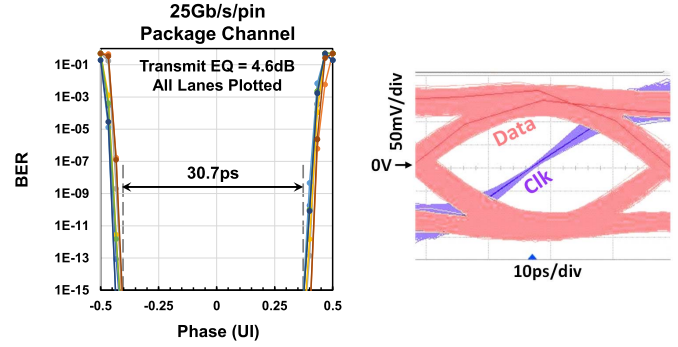


Fig. 12. Measured bathtub curve including all eight data lanes (left) and measured data/clock eye diagram (right), for the on-package link.

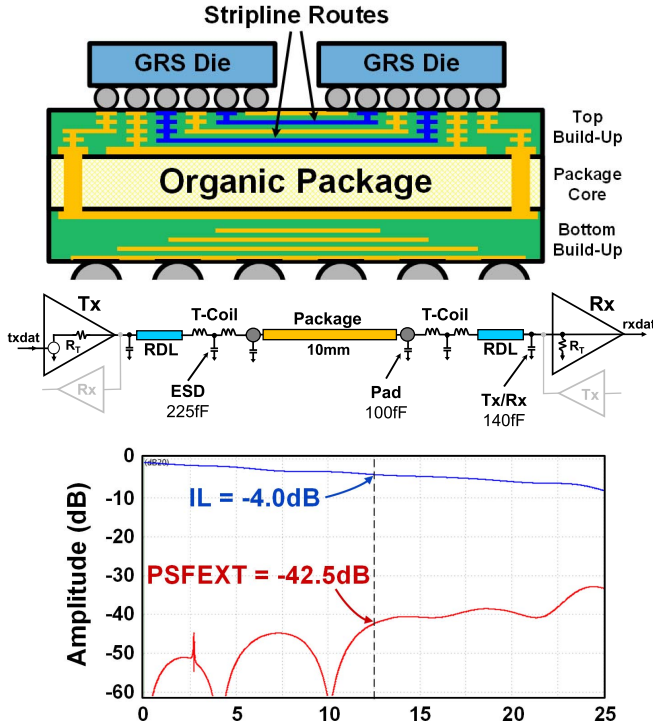


Fig. 11. Channel cross section (top), electrical model (middle), and frequency response (bottom) of the 10-mm on-package channel.

By way of example, consider back-matching a 40- Ω line. The transmitter drives the line through a T-coil and an RDL trace, totaling about 4.6 Ω of series resistance. Setting the termination resistor to 47 Ω (in parallel with the transmitter's 140 Ω) provides the required match. Since most of the termination conductance is provided by the adjustable terminator, variations in the charge-pump's output impedance cause only small changes in the overall termination impedance. In effect, the adjustable termination resistor's large conductance "linearizes" the transmitter's impedance match to the channel.

C. GRS Receiver

Fig. 8 shows the schematic for the GRS receiver, which consists of an input amplifier followed by two subsequent CMOS inverter stages. The receiver input amplifier performs several functions:

- 1) level-converts the input signal from symmetric-about-ground to symmetric-about-the-inverter switching threshold, nominally $V_{dd}/2$;
- 2) provides modest voltage gain to reduce the input-referred offset and noise of the receiver;
- 3) includes high-frequency signal peaking to overcome its own output pole and help equalize the channel;
- 4) converts the SE input signal into a pseudodifferential output.

The subsequent CMOS inverter stage provides most of the signal gain, about 14 dB, while the last inverter stage acts as a limiting amplifier by providing sufficient gain to amplify the signal to CMOS levels. The full-swing differential output signals are used in the forwarded clock lane to drive the root of the receiver clock network (**rxclk**). In the data lanes, only one of these outputs drives the SE data path to reduce power consumption, which would otherwise require differential signal paths within the tunable matched delay line. The tunable delay line that precedes the samplers provides ± 10 ps of delay adjustment, in 1.5-ps steps, to each lane. This can be used to compensate for ± 1.5 -mm trace length mismatch between each data lane and the forwarded clock lane. The delay range and step size do not vary significantly with PVT, thanks to the adaptive regulation scheme.

Several iterations of the GRS input amplifier have been described in [12], [14], and [17]. These are all based on the common-gate (CG) topology, since this configuration does not suffer from bandwidth limitations associated with the Miller effect. In the link reported here, the amplifier of [16] is altered to synthesize a complementary signal at the output of the right-hand branch to realize a pseudodifferential signal. The input stage consists of two identical CG amplifier branches that are biased using reference current mirrors with cross-coupled source connections. The branch on the left behaves as a CG amplifier while the branch on the right operates as a common-source amplifier, providing the necessary signal inversion for pseudodifferential output. To provide signal peaking, the amplifier load uses adjustable G_m -C active inductors [17], [18], which can be tuned by changing the form factor of the G_m devices in addition to the RC feedback networks. A small cross-coupled PMOS pair in parallel with the active inductors improves the effective quality factor. RC feedback networks at the gates of the diode-connected current mirrors form G_m -C active inductors that equalize the current mirror

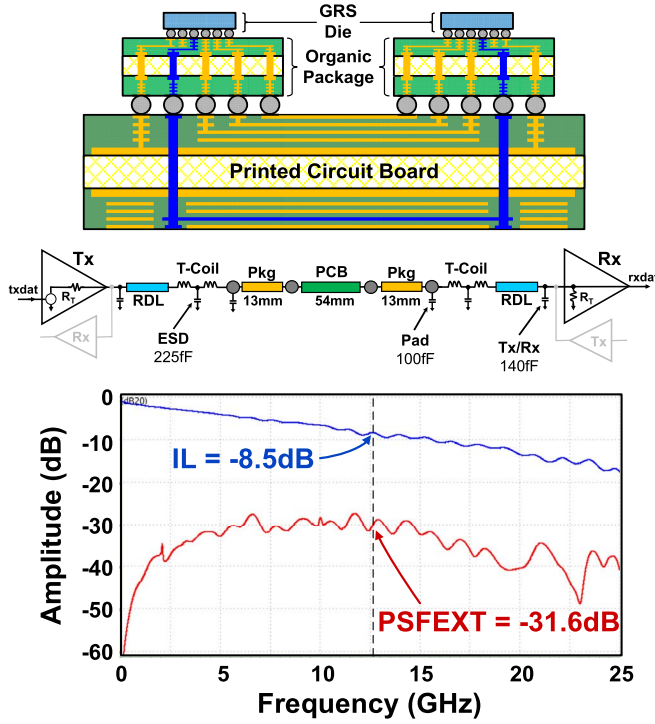


Fig. 13. Channel cross section (top), electrical model (middle), and frequency response (bottom) of the short-reach PCB channel.

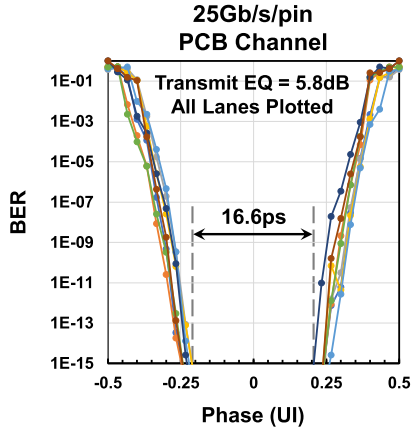


Fig. 14. Measured BER bathtub curves for eight data lanes operating over the short-reach PCB link.

nodes to minimize the insertion delay to the output of the right-hand amplifier branch. To ensure the maximum voltage gain in the second-stage inverters, a CM feedback loop serves the input stage bias currents so the differential signals at the inverter inputs are centered on the inverter switching threshold.

Fig. 9 plots the transfer function of the input amplifier across G_m -C adjustments.

In practice, the high-frequency peaking in the input amplifier is needed to overcome the output pole of the second-stage inverter amplifiers, so most of the channel attenuation is compensated in the transmitter equalizer. Additional receiver equalization could easily be provided at the expense of increased power; the sizing and power consumption of the receiver in this design were chosen to reduce overall link energy.

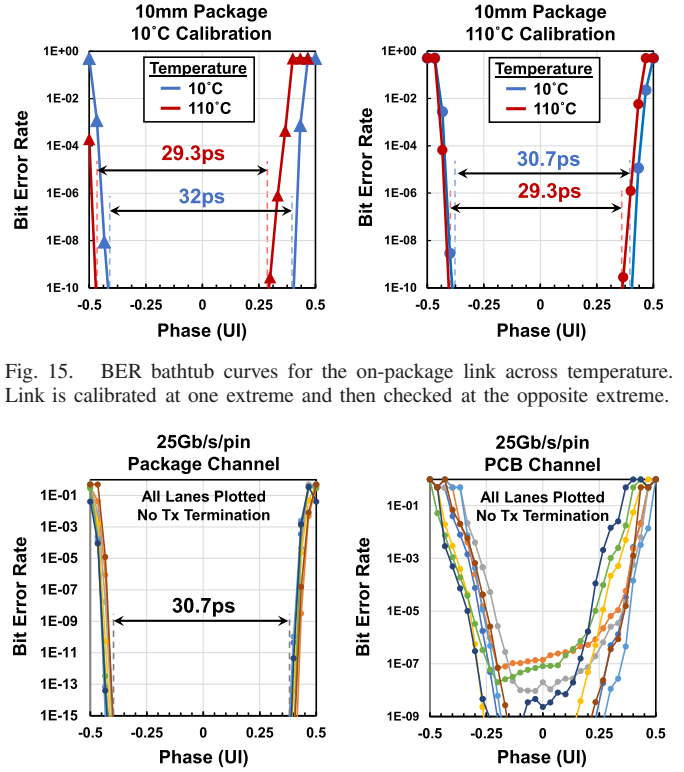


Fig. 15. BER bathtub curves for the on-package link across temperature. Link is calibrated at one extreme and then checked at the opposite extreme.

Fig. 16. BER bathtub curves measured without additional parallel Tx termination for both the on-package link (left) and PCB link (right). The on-package link is unaffected by Tx termination. However, the PCB link is more sensitive to the impedance match between the Tx and the channel.

The receiver input amplifier has gone through a progression of designs as described in [14], from a simple resistor-loaded CG amplifier, to a complementary CG amplifier, and finally the pseudodifferential amplifier described here. The objective of this development process was the reduction of temperature variation in receiver input offset. The amplifier in Fig. 8 provides two tunable offset cancellation mechanisms: 1) The R_{trim} resistor that approximates the parallel combination of line impedance and terminator resistor, and 2) the differential current source that drives the main CG amplifier stage. The CM setting of the mirrored current sources is controlled by a CM feedback loop, as described above. Temperature drift is nearly eliminated by a synergistic combination of the PLL-referenced supply regulation, pseudodifferential input stage design, and continuous CM feedback in the amplifier. See Fig. 15 for an experimental verification of offset cancellation and link margin across temperature.

IV. EXPERIMENTAL RESULTS

The GRS link circuitry was fabricated in a 16-nm FinFET CMOS technology. Fig. 10 shows a die photograph of the link. Overall area including the 5×4 array of C4 bumps is $686 \mu\text{m} \times 565 \mu\text{m}$. The “brick” containing the I/O active circuitry is shown in the dotted outline and occupies $403 \mu\text{m} \times 202 \mu\text{m}$ (area = $81,406 \mu\text{m}^2$). The clock (C) and data signal (S) bumps are arranged in a checkerboard pattern with the ground (G) and power (P) bumps, where supply bumps act as physical shields between adjacent lanes to reduce crosstalk

TABLE I
COMPARISON OF THIS WORK TO OTHER PUBLISHED LINKS

	[1]	[10]	[11]	[12]	This Work
Signaling	Single-Ended	Differential	CNRZ-5	Single-Ended	Single-Ended
Data Rate/pin	10.0 Gb/s	56 Gb/s	20.83 Gb/s	20 Gb/s	25 Gb/s
Reach	100 mm	80 mm	12 mm	4.5 mm	80 mm
Application	MCM or PCB	PCB	MCM	MCM	MCM or PCB
CMOS Technology	65 nm	16 nm	28 nm	28 nm	16 nm
Energy/bit	4.18 pJ/bit	2.25 pJ/bit	0.94 pJ/bit	0.54 pJ/bit	1.17 pJ/bit
Channel Attenuation	-14.5 dB	-11 dB	-1.3 dB	-1.5 dB	-8.5 dB

from the vertical transitions. All lanes of the test chip are transceivers.

A compact floor plan is preferred since all lanes use a common timing reference (the forwarded clock). This helps to minimize clock distribution power [19] and to hold skew below 500 fs. Length-matched routing (270 μ m) on the RDL adds 60 fF of capacitance to the I/O of each lane. Ultimately, the total area used depends on the minimum C4 bump pitch, which is set by the thermomechanical limits of a reticle limited die attached to a particular packaging technology, as well as the number of ground return bumps needed to mitigate crosstalk. Any remaining area is used as power supply capacitance.

The GRS link reported here was used in two applications: first, a very-short-reach link that operates over high-density interconnect (HDI) between chips mounted on the same organic substrate, and second, a short-reach link between packaged chips operating over a short printed circuit board (PCB) channel.

The physical arrangement of the on-package link is shown in Fig. 11 (top), where the chips are mounted on a conventional organic package with a 4–2–4 metal stack-up. The on-package signal routes are restricted to the top-side build-up HDI layers so blind vias could be used for all vertical connections, thereby minimizing stub lengths of the via structures. The channels are routed as strip-lines in the second layer of HDI, with ground planes above and below the signal traces to reduce crosstalk between adjacent lanes. The trace lengths of all lanes are matched to within approximately 1 ps.

Fig. 11 also shows the electrical model (center) and frequency characteristics (bottom) of the 10-mm channel, where the frequency response of the entire channel was modeled using a 3-D field solver. Production-level ESD clamps are included at both ends of the link to ensure robust operation. The extracted capacitances at either end of the link are shown in Fig. 11, which includes the contributions of the I/O pads, ESD structures, and circuit parasitics. T-coils are used to compensate for these capacitance contributions at each end, thereby improving back-match and minimizing inter-symbol interference. The channel has -4 dB of insertion loss (IL)

at Nyquist, while the power sum far-end crosstalk (PSFEXT) is kept 38.5 dB below IL at Nyquist.

Fig. 12 shows the measured bit-error-rate (BER) bathtub curve for all lanes on the 10-mm package channel (left) and a measured eye diagram of data and clock signals (right). The data pattern is PRBS-31 with different seeds for each lane, measured at 25 Gb/s. An on-chip phase interpolator (included only for characterization purposes) provides 1.33-ps steps for the BER bathtub plot. At a BER of 10^{-15} , the measured aggregate eye opening is 30.7 ps (0.77 UI) when using +4.6 dB of transmitter EQ.

Fig. 13 (top) shows the physical arrangement of the package-to-package link, where two chips are mounted on conventional organic packages to signal across a conventional low-cost PCB. The signal routes in the PCB were restricted to the bottom-most strip-line routing layer to minimize stub lengths associated with the plated through hole vias required for vertical transitions within the PCB. To further reduce crosstalk, grounded shields were inserted between adjacent signal routes on the PCB. As shown in Fig. 13 (middle), the link comprises 13 mm of on-package HDI traces on both ends of the link, vertical paths through the package and PCB vias, and 54-mm strip-line routes on the bottom of the PCB.

In this channel, the delays of the clock lane and data lanes are matched to within approximately ± 0.6 mm in length (± 4 ps). The entire interconnect, including package traces, bump and ball connections to the package and PCB, vertical via structures, and PCB traces, were modeled using a 3-D field solver to produce the channel frequency response in Fig. 13 (bottom). Crosstalk is shown as the power-sum of the FEXT signals from all eight aggressors. The channel has -8.5 dB of IL at Nyquist, while the PSFEXT is kept 23.1 dB below IL at Nyquist.

Fig. 14 shows the measured BER bathtub curves for all eight data channels operating at 25 Gb/s. The data pattern is PRBS-31 with different seeds for each data lane. The measured aggregate eye opening at BER = 10^{-15} is approximately 16.6 ps (0.42 UI) when using +5.8 dB of transmit EQ.

Link margining is measured across temperature variation to verify that one-time calibration under arbitrary conditions of




Fig. 17. Area breakdown (left) and power breakdown (right) for the entire GRS link, comprising eight data lanes and one clock lane.

supply voltage and temperature is sufficient to guarantee subsequent link timing margin under VT variation. Fig. 15 shows the bathtub curves for the 10-mm on-package link. The link is first calibrated at one temperature extreme, then, timing margin is measured at the initial temperature and opposite extreme with no intermediate adjustments. Data patterns were PRBS-31 with different seeds on each lane, and upon demonstrating low temperature sensitivity, the measurements were concluded at $\text{BER} = 10^{-10}$ for brevity.

Fig. 16 shows the measured BER bathtub curves for the on-package link (left) and the PCB link (right) without the additional parallel Tx termination used to match the channel impedance. For the case of the on-package link, there was no measurable change in the timing margin across all eight data lanes to a $\text{BER} = 10^{-15}$. This is due to the exceptional signal integrity of the on-package channel, which has no discontinuities. However, when the additional parallel Tx termination is not used with the PCB link, there is an increase in reflections and crosstalk that degrades the timing margin of all lanes. Four of the eight PCB lanes are more sensitive to the absence of a good back-match, and they fail before a $\text{BER} = 10^{-9}$ can be achieved. Both the on-package and PCB link data patterns were PRBS-31 with different seeds on each lane.

Fig. 17 summarizes the area breakdown for the components of the I/O “brick,” and power breakdown for the complete GRS link, which includes eight data lanes and one clock lane. Table I shows a comparison of the GRS link to other contemporary general-purpose short-reach links.

V. CONCLUSION

This paper describes a single-ended, bundled-data clock-forwarding GRS link that supports high signaling bandwidth (25 Gb/s/wire) at low energy per bit (1.17 pJ/bit). The link has been demonstrated to operate at very low BER of 10^{-15} in both on-package and package-to-package applications, with only one-time calibration. The link can be applied to high-speed interconnects within MCMs, electrical/optical interface chips, and memory interconnects, where bandwidth per pin achieved by this link is comparable to recent CEI/OIF 56 Gb/s experimental links but at roughly one-half the energy per bit.

ACKNOWLEDGMENT

The authors would like to acknowledge the assistance and design expertise of the signal integrity, package design, and printed circuit board (PCB) design teams at NVIDIA.

REFERENCES


- [1] J. Song, S. Hwang, H.-W. Lee, and C. Kim, “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,” *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 65, no. 1, pp. 331–342, Jan. 2018.
- [2] OIF-CEI-040-Common Electrical I/O (CEI)-Electrical and Jitter Interoperability Agreements for 6G+ bps, 11G+ bps, 25G+ bps I/O and 56G+ bps, Optical Internetworking Forum (OIF), Dec. 2017.
- [3] P. Upadhyaya *et al.*, “A fully adaptive 19-to-56Gb/s PAM-4 wireline transceiver with a configurable ADC in 16 nm FinFET,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 108–109, Paper 6.4.
- [4] L. Wang, Y. Fu, M. LaCroix, E. Chong, and A. C. Carusone, “A 64 Gb/s PAM-4 transceiver utilizing an adaptive threshold ADC in 16 nm FinFET,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 110–111, Paper 6.5.
- [5] T. O. Dickson, H. A. Ainspan, and M. Meghelli, “A 1.8 pJ/b 56 Gb/s PAM-4 transmitter with fractionally spaced FFE in 14nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 118–119, Paper 6.5.
- [6] J. Im *et al.*, “A 40-to-56 Gb/s PAM-4 receiver with 10-tap direct decision-feedback equalization in 16nm FinFET,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 114–115, Paper 6.3.
- [7] T. Shibusaki *et al.*, “A 56 Gb/s NRZ-electrical 247 mW/lane serial-link transceiver in 28 nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Jan./Feb. 2016, pp. 64–65, Paper 3.5.
- [8] M. O'Connor. (Jun. 14, 2014). Highlights of the High Bandwidth Memory (HBM) Standard. The Memory Forum. [Online]. Available: <http://www.cs.utah.edu/thememoryforum/mike.pdf>
- [9] E. Depaoli *et al.*, “A 4.9 pJ/b 16-to-64 Gb/s PAM-4 VSR transceiver in 28 nm FDSOI CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 112–113, Paper 6.6.
- [10] M. Erett *et al.*, “A 126 mW 56 Gb/s NRZ wireline transceiver for synchronous short-reach applications in 16 nm FinFET,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 274–275, Paper 16.7.
- [11] A. Shokrollahi *et al.*, “A pin-efficient 20.83 Gb/s/wire 0.94 pJ/bit forward clock CNRZ-5-coded SerDes up to 12 mm for MCM packages in 28 nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Jan./Feb. 2016, pp. 182–183, Paper 10.1.
- [12] J. W. Poulton *et al.*, “A 0.54 pJ/b 20 Gb/s ground-referenced single-ended short-reach serial link in 28 nm CMOS for advanced packaging applications,” *J. Solid-State Circuits*, vol. 48, no. 12, pp. 3206–3218, Dec. 2013.
- [13] J. Wilson *et al.*, “A 1.17 pJ/b 25 Gb/s/pin ground-referenced single-ended serial link for off- and on-package communication in 16 nm CMOS using a process- and temperature-adaptive voltage regulator,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 276–277, Paper 16.8.
- [14] W. J. Turner *et al.*, “Ground-referenced signaling for intra-chip and short-reach chip-to-chip interconnects,” in *Proc. Custom Integr. Circuits Conf.*, Apr. 2018, pp. 1–8, Paper 18-5.
- [15] S. S. Kudva, S. Song, J. Poulton, J. Wilson, W. Zhao, and C. T. Gray, “A switching linear regulator based on a fast-self-clocked comparator with very low probability of meta-stability and a parallel analog ripple control module,” in *Proc. Custom Integr. Circuits Conf.*, Apr. 2018, pp. 1–4, Paper 6-5.
- [16] M. G. Johnson, “MOSFET sense amplifier circuit,” U.S. Patent 4523110, Sep. 30, 1983.
- [17] L. Fick, D. Sylvester, J. Poulton, J. Wilson, and T. Gray, “A 25 Gb/s 470 μ W active inductor equalizer for ground referenced signaling receivers,” in *Proc. Int. Symp. Circuits Syst.*, May 2017, pp. 1–4.
- [18] F. Yuan, *CMOS Active Inductors and Transformers: Principle, Implementation, and Applications*. New York, NY, USA: Springer, 2008.
- [19] F. O'Mahony *et al.*, “47 \times 10 Gb/s 1.4 mW/(Gb/s) parallel interface in 45 nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2010, pp. 156–157.



John W. Poulton (M'85–SM'90–F'12–LF'16) received the B.S. degree in physics from the Virginia Polytechnic Institute and State University, Blacksburg, VA, USA, in 1967, the M.S. degree in physics from the State University of New York, Stony Brook, NY, USA, in 1969, and the Ph.D. degree in physics from the University of North Carolina, Chapel Hill (UNCCH), NC, USA, in 1980.

From 1981 to 1999, he was a Researcher with the Department of Computer Science, UNCCH, where from 1995 he held the rank of Research Professor.

He performed research on VLSI-based architectures for graphics and imaging. He was a principal contributor to the design and construction of the PixelPlanes and PixelFlow computer graphics systems, and designed custom beam-forming chips for the first commercial 3-D medical ultrasound scanner. From 1999 to 2003, he was Chief Engineer with Velio Communications, where he developed multi-gigabit chip-to-chip signaling systems. From 2003 to 2009, he was a Technical Director with Rambus Inc., Chapel Hill, where he led an effort to build power-efficient multi-gigabit I/O systems, demonstrating a system in 2006 with the lowest energy per bit published up to that time. He is currently Senior Scientist with NVIDIA, Inc., Durham, NC, USA, where he is working on low-energy on- and off-chip signaling.



John M. Wilson (S'98–M'03–SM'16) received the B.S., M.S., and Ph.D. degrees in electrical engineering from North Carolina State University (NCSU), Raleigh, NC, USA, in 1993, 1995, and 2003, respectively.

From 2003 to 2006, he was a Research Professor at NCSU leading projects in advanced packaging, low-power capacitive and inductive coupled transceivers for 3-D ICs, and circuits for low-power on-chip global signaling. From 2006 to 2012, while with Rambus Inc., he worked on low-power, high-speed

I/O circuit design, and methods to mitigate signal and power integrity problems in single-ended memory interfaces. In 2012, he joined NVIDIA, Inc., Durham, NC, USA, where he is a Senior Research Scientist in the Circuits Research Group. He has authored over 66 publications and more than 30 granted patents. His current research interests include high-speed I/O circuit design, on-chip signaling, signal integrity, advanced packaging, and chip/package co-design.



Walker J. Turner (M'15) received the B.S., M.S., and Ph.D. degrees in electrical and computer engineering from the University of Florida, Gainesville, FL, USA, in 2009, 2012, and 2015, respectively.

In 2014, he was a Contractor with the U.S. Army Research Laboratory, Adelphi, MD, USA, developing wirelessly powered systems and integrated low-noise amplifiers for piezoelectric *E*-field sensors. Since 2015, he has been with NVIDIA, Inc., Durham, NC, USA, where he works as a Senior Research Scientist with the Circuits Research Group.

His current research interests include low-power integrated circuit design for mixed-signal and high-speed signaling applications.



Brian Zimmer (S'09–M'15) received the B.S. degree in electrical engineering from the University of California, Davis, CA, USA, in 2010, and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 2012 and 2015, respectively.

He is currently with the Circuits Research Group, NVIDIA, Inc., Santa Clara, CA, USA. His current research interests include soft error resilience and energy-efficient digital design, with an emphasis on low-voltage SRAM design and variation tolerance.



Xi Chen (S'08–M'12) received the B.S. and M.S. degrees in electrical engineering from Zhejiang University, Hangzhou, China, in 2003 and 2006, respectively, and the Ph.D. degree in electrical engineering from North Carolina State University (NCSU), Raleigh, NC, USA, in 2011.

From 2003 to 2006, he was with the Institute of VLSI Design, Zhejiang University. From 2006 to 2007, he was an Analog IC Design Engineer with Accel Semiconductor, Shanghai, China. From 2007 to 2011, he was a Research Assistant with the

Department of Electrical and Computer Engineering, NCSU. Since 2012, he has been a Research Scientist with NVIDIA, Inc., Durham, NC, USA. Later in 2012, he moved to NVIDIA, Santa Clara, CA, USA. His current research interests include high-speed signaling and clocking, 3-D IC design methods, and analog mixed-signal ICs.



Sudhir S. Kudva received the B.E. degree in electrical engineering from the National Institute of Technology, Suratkal, India, in 2004, the M.E. degree in electrical engineering from the Indian Institute of Science, Bangalore, India, in 2006, and the Ph.D. degree in electrical engineering from the University of Minnesota, Twin Cities, MN, USA, in 2013.

From 2006 to 2008, he was a Design Engineer with the AMD India Engineering Centre, Bangalore, India, designing ROMs in the 65-nm and 45-nm

silicon-on-insulator technology. In 2013, he joined NVIDIA Research, Santa Clara, CA, USA, where he is currently a Senior Research Scientist and a member of the Circuits Research Group. His current research interests include modeling and design of power delivery networks, fully integrated on-die voltage regulators, high step-down ratio, point of load converters for high-power graphics cards, and variation tolerant voltage regulator design.



Sanquan Song (S'05–M'11) received the Ph.D. degree from the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2011.

From 2010 to 2013, he was with Intel Corporation. From 2013 to 2015, he was with the Samsung Display R&D Laboratory. In 2015, he joined the Circuits Research Group, NVIDIA, Inc., Santa Clara, CA, USA, focusing on SerDes. He has authored multiple papers and patents and actively reviewed for

multiple journals, including JSSC, TCAS-I, TCAS-II, and TVLSI. His current research interests include SerDes modeling, design, and implementation.



Stephen G. Tell (M'97) received the B.S.E. degree in electrical engineering from Duke University, Durham, NC, USA, in 1989, and the M.S. degree in computer science from the University of North Carolina, Chapel Hill (UNC/Chapel Hill), NC, USA, in 1991.

From 1991 to 1999, he was a Senior Research Associate with UNC/Chapel Hill, working on parallel graphics systems and high-speed signaling. In 1999, he joined Velio, Inc., developing circuits and control systems for high-speed SerDes products.

This work continued at Rambus Inc., until 2009, where he designed the logic for a SerDes with the lowest energy per bit demonstrated up to that time. In 2009, he joined NVIDIA, Inc., Durham, NC, USA, as a Founding Member of the Circuits Research Group. He holds more than 10 U.S. patents. His current research interests include custom circuit design and the surrounding logic for intra- and inter-chip communication.



Nikola Nedovic (S'99–M'03) received the Dipl. Ing. degree in electrical engineering from the University of Belgrade, Serbia, in 1998, and the Ph.D. degree from the University of California, Davis, CA, USA, in 2003.

In 2001, he joined Fujitsu Laboratories of America, Inc., Sunnyvale, CA, USA, where he worked on high-speed communications and high-performance and low-power circuits for electrical and optical communications. In 2016, he joined NVIDIA Research, Santa Clara, CA, USA, where he works on system and circuit design for low-power high-speed links. His current research interests include a range of aspects of high-speed electrical and optical communications, from devices and signal integrity to adaptive filtering and system design and modeling.



Wenxu Zhao received the B.S. degree in micro-electronics from Fudan University, Shanghai, China, in 2010, and the M.S. and Ph.D. degrees in electrical engineering from North Carolina State University, Raleigh, NC, USA, in 2012 and 2017, respectively.

He was an Analog Design Intern at Cirrus Logic, Inc., in the summer of 2013. In 2016, he was a Research Intern with NVIDIA Research, Santa Clara, CA, USA, working on high-speed serial link design. In 2017, he joined Broadcom, Irvine, CA, USA, where he works on physical layer products for coherent optical communications.



Sunil R. Sudhakaran received the B.S. degree in computer engineering from the University of Madison–Wisconsin, Madison, WI, USA, and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 2004 and 2006, respectively.

He is currently the Director of Signal and Power Integrity with NVIDIA, Inc., Santa Clara, CA, USA. His current research interests include high-speed I/O link modeling and investigating ways to extend the bandwidth of memory interfaces.



C. Thomas Gray (M'92–SM'16) received the B.S. degree in computer science and mathematics from the Mississippi College, Clinton, MS, USA, and the M.S. and Ph.D. degrees in computer engineering from North Carolina State University, Raleigh, NC, USA.

From 1993 to 1998, he was an Advisory Engineer with IBM, Research Triangle Park, NC, USA, working on transceiver design for communication systems. From 1998 to 2004, he was a Senior Staff Design Engineer with the Analog/Mixed Signal Design Group, Cadence Design Systems, working on SerDes system architecture. From 2004 to 2010, he was a Consultant Design Engineer with Artisan/ARM and Technical Lead of SerDes architecture and design. In 2010, he joined Nethra Imaging as a System Architect. His work experience includes digital signal processing design and CMOS implementation of DSP blocks as well as high-speed serial link communication systems, architectures, and implementation. In 2011, he joined NVIDIA, Inc., Durham, NC, USA, where he is currently the Senior Director of Circuit Research, leading activities related to high-speed signaling, low-energy and resilient memories, circuits for machine learning, and variation-tolerant clocking and power delivery.



William J. Dally (M'80–SM'01–F'02) received the B.S. degree in electrical engineering from Virginia Tech, Blacksburg, VA, USA, in 1980, the M.S. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 1981, and the Ph.D. degree in computer science from the California Institute of Technology (Caltech), Pasadena, CA, USA, in 1986.

He is currently Chief Scientist and Senior Vice President of Research at NVIDIA, Inc., Santa Clara, CA, USA, and a Professor (Research) and former Chair of computer science at Stanford University. He is currently working on developing hardware and software to accelerate demanding applications, including machine learning, bioinformatics, and logical inference. He has a history of designing innovative and efficient experimental computing systems. While at Bell Labs, Murray Hill, NY, USA, he contributed to 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. At the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, 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 developed the Imagine processor, which introduced the concepts of stream processing and partitioned register organizations, the Merrimac super-computer, which led to GPU computing, and the ELM low-power processor. He currently leads projects on computer architecture, network architecture, circuit design, and programing systems. He has published over 250 papers in these areas and is an author of the textbooks *Digital Design: A Systems Approach*, *Digital Systems Engineering*, and *Principles and Practices of Interconnection Networks*, and holds over 160 issued patents.

Dr. Dally is a Member of the National Academy of Engineering, a Fellow of the IEEE, and a Fellow of the ACM, and a Fellow of the American Academy of Arts and Sciences. He has received the ACM Eckert-Mauchly Award, the IEEE Seymour Cray Award, the ACM Maurice Wilkes Award, the IEEE Computer Society Charles Babbage Award, and the IPSJ FUNAI Achievement Award.