

6.6 Reference-Noise Compensation Scheme for Single-Ended Package-to-Package Links

Xi Chen¹, Nikola Nedovic¹, Stephen G. Tell², Sudhir S. Kudva¹, Brian Zimmer¹, Thomas H. Greer², John W. Poulton², Sanquan Song¹, Walker J. Turner², John M. Wilson², C. Thomas Gray²

¹NVIDIA, Santa Clara, CA

²NVIDIA, Durham, NC

A recent trend in high-performance systems is distribution of computing across many chips and packages to sustain performance scaling while achieving high yield and alleviating power delivery. High-end data center systems and new applications like deep neural network (DNN) accelerators with scalable architecture [1] may extend the system from large chip-scale computing not just to package-scale multi-chip modules (MCM), but also to PCB-scale computing systems. An essential requirement for these distributed systems is a highly scalable low-power and high-bandwidth interconnect system, that can cover a wide range of integration and channel distances. Ideally, the same low-power links designed for ultra-short reach interconnects in MCM should be used in more challenging board-level environments.

With better energy and area efficiency than differential links, single-ended serial links are an attractive signaling technique in short-reach high-bandwidth I/O applications. One major drawback of single-ended links is their sensitivity to environment noise, because the accuracy of data recovery at the receiver relies heavily on good correlation of references at the two ends of the link. Modern GPUs and CPUs can ramp hundreds of amperes of supply current within tens of nanoseconds. Given typical power delivery network resonances, these current transients can cause local supply voltage and ground noise amplitudes of hundreds of millivolts at tens of megahertz and create large reference offset between packages on the same printed circuit board (PCB, Fig. 6.6.1). Traditional methods to solve the reference matching problem include using differential signaling [2] and data pattern coding [3], but both methods have relatively large power and pin overhead compared to simple single-ended design.

We describe a short-reach 25Gb/s clock-forwarded link that includes a reference noise compensation technique; this design enables a low-power single-ended link to operate over noisy PCB channels with negligible power overhead. The noise compensation mechanism extracts the common-mode noise information from the received clock and compensates it at the tunable RX front-end dynamically (Fig. 6.6.2). In this link, the forwarded clock is half-rate (equivalent to a 1010... data pattern), thus the TX-RX reference error is encoded in the clock duty cycle for any finite clock transition time. The clock duty cycle is detected by self-sampling, where the clock is sampled with a 1UI-delayed version of itself (Clk_ck). This results in a sampled "1" for duty cycles above 50% (interpreted as positive TX-to-RX reference offset) and sampled "0" for duty cycles below 50% (negative reference offset). The compensation logic observes the clock lane deserialized outputs for duty-cycle error and adjusts the RX front-end offset tuning to drive the clock duty cycle to 50%, thereby tracking the reference noise and closing the control loop. The same adjustment is combined with the statically calibrated offset settings and forwarded to all data lanes' RX front-ends. During bring-up, the data path delay (δ) in the forwarded clock lane has been trimmed to match the delay of the Rxclk buffers (Fig. 6.6.2). A delay-locked loop (DLL) is formed using the clock lane sampler, deserializer, and logic to dynamically adjust the 1UI clock delay for environment changes. The residual error of the offset compensation is primarily caused by loop latency and offset tuning step size.

Figure 6.6.3 illustrates the error detection conditions for the offset compensation loop and DLL. The phase sense error for the DLL is obtained by comparing even and odd samples of the same deserialized clock as the reference noise loop. In both loops, we implement programable sense thresholds to control the open loop gains and tolerate uncalibrated errors. Voltage regulation and bring-up calibration can put the default setting of the 1UI block around the correct delay value. Small residual error of the 1UI block will not affect the offset compensation loop when it has enough threshold margin. Larger error of the delay may be caused by fast environment changes (e.g. VDD noise), and it could temporarily increase the dead zone in the offset loop transfer function. However, a delay error detector will sense it and move the control codes to tune the delay back to an acceptable range. The control logic of both loops was synthesized in a standard digital flow. The digital domain, including compensation logic, runs at the parallel clock frequency (1.56GHz). The routing latency and digital processing time may limit the

compensation loop response speed, so a programable parameter called Hold cycles was added to slow down the tuning code output rate from the pclk (1.56GHz) to a fraction of this rate. This feature improves the noise tracking accuracy and minimizes the jitter in both offset and delay loops at the cost of potentially lower bandwidth. Delay and offset counters adjust control codes by 1LSB each permitted cycle.

The schematic of the 1UI delay block is shown in Fig. 6.6.4. The main requirements of this block are 1) monotonic tuning, 2) fine resolution to match offset tuning step, and 3) wide enough range (>25% or 10ps) to tolerate environment changes such as VDD noise. To meet these requirements, we optimized for linearity by distributing each bit to all delay stages thereby achieving +/-0.4LSB DNL with 0.8ps average step size. Process variation of the 1UI delay is compensated by tuning the clock lane data-path delay. In the compensation loops, the sampler behaves as a single bit quantizer, which may cause unfavorably high gain in low-noise conditions. To avoid overshooting in the dynamic response of the offset compensation loop, clock dithering is included to linearize the noise transfer function, by toggling the delay control for a few LSBs in each parallel clock (Pclk) cycle.

A reference noise compensation experiment was implemented in a 25Gb/s/pin ground-referenced serial link like [4] and fabricated in 16nm FINFET process (Fig. 6.6.7). The high-speed links are part of a scalable DNN accelerator and perform low-power communication with neighboring chips on the same or a nearby package. The link consists of four data lanes and one forwarded clock lane for each unidirectional port. The same design can support up to eight data lanes. The reference noise compensation performance was validated in a package-to-package link, which has larger noise amplitude.

To test the system performance, we inject high-amplitude low-frequency noise in the PCB ground, mimicking the environment in a large multi-package board. Figure 6.6.5 shows the BER performance of the link over a 50mm PCB channel (80mm including packages) when ground noise was injected. With low ground impedance on the PCB, 60A current causes 122mV_{pp} reference noise at 60Hz, and reduces the eye opening by 66.7%. The compensation loop successfully recovered 94% of the lost time margin. For high frequency characterization of the compensation loop, we create noisy clock and data by adding sinusoidal noise to a noise-free clock and pseudorandom data and inject these noisy signals to the transmitter end PCB pads via RF probes. The link performance at 10MHz noise in the RF probe experiment is also shown in Fig. 6.6.5. With even higher noise amplitude, the compensation loop can still recover most of the lost margin while uncompensated eyes were closed. By reducing the hold cycle value, which increases the loop bandwidth, we were able to open the closed eyes with up to 30MHz of reference noise, at the expense of low-frequency compensation. The clock dithering improved noise compensation performance in all tested cases.

Figure 6.6.6 shows the comparison between the single-ended link with our reference noise compensation scheme and other state-of-the-art short-reach links, which are potentially tolerant to the reference noise. Our work demonstrates the capability of reference noise compensation to extend the reach of single-ended links to noisy package-to-package channels. The overall power overhead of the noise compensation circuit is around 1%, and its area overhead is negligible.

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] B. Zimmer et al., "A 0.11 pJ/Op, 0.32-128 TOPS, Scalable Multi-Chip-Module-based Deep Neural Network Accelerator with Ground-Reference Signaling in 16nm," *IEEE Symp. VLSI Circuits*, pp. 300-301, June 2019.
- [2] M. Erett et al., "A 126mW 56Gb/s NRZ Wireline Transceiver for Synchronous Short-Reach Applications in 16nm FinFET," *ISSCC*, pp. 274-275, Feb. 2018.
- [3] 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.
- [4] J. Poulton et al., "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", *IEEE JSSC*, vol. 54, no. 1, pp. 43-54, Jan. 2019.

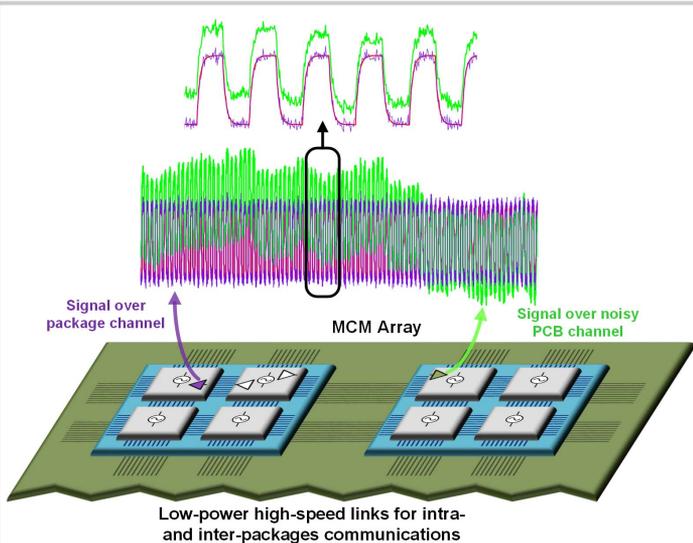


Figure 6.6.1: Reference noise problem in high-speed link.

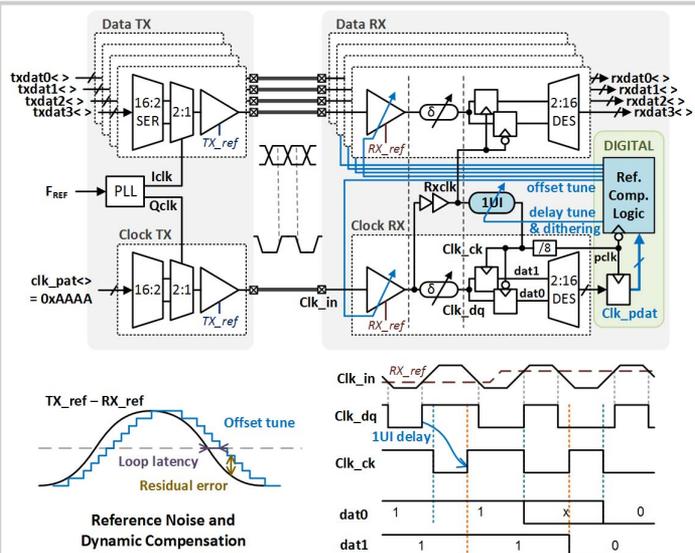


Figure 6.6.2: Reference noise compensation based on self-sampling clock.

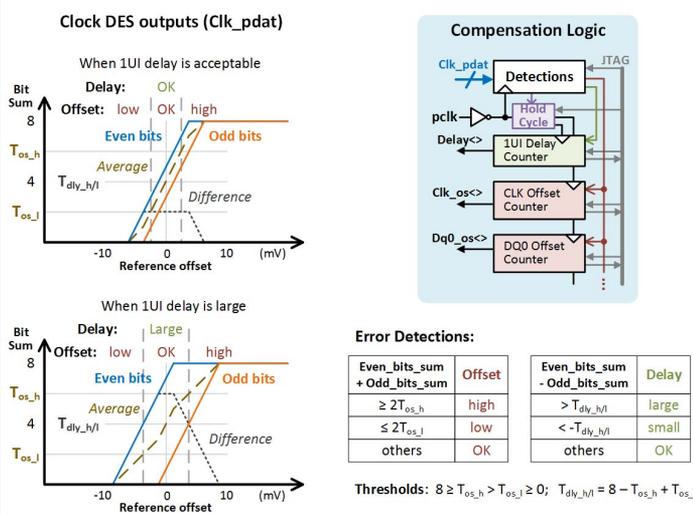


Figure 6.6.3: Error detection conditions (16b DES), and compensation logic design.

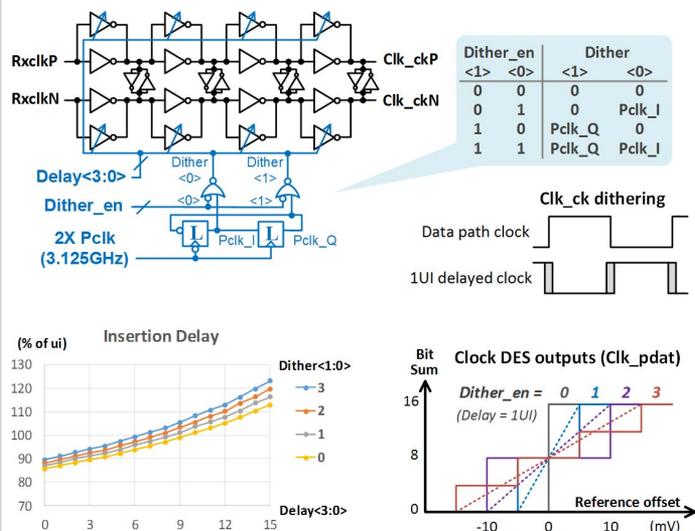


Figure 6.6.4: 1UI delay block: schematic, clock dithering, and performance.

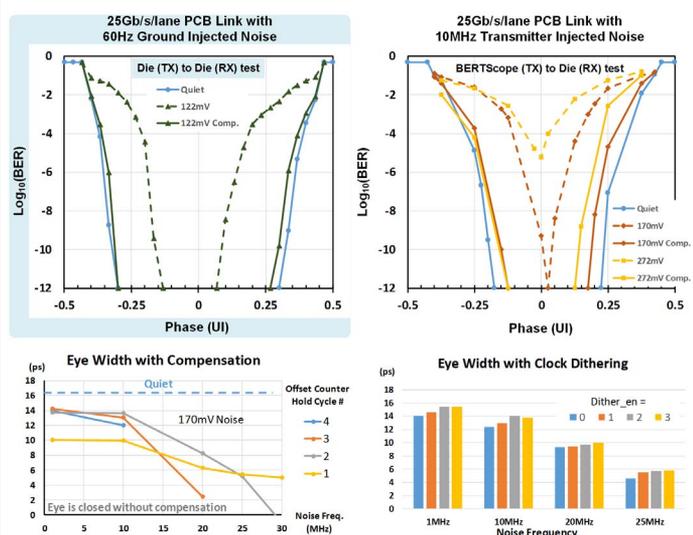


Figure 6.6.5: 25Gb/s PRBS-31 measurement results for PCB channel.

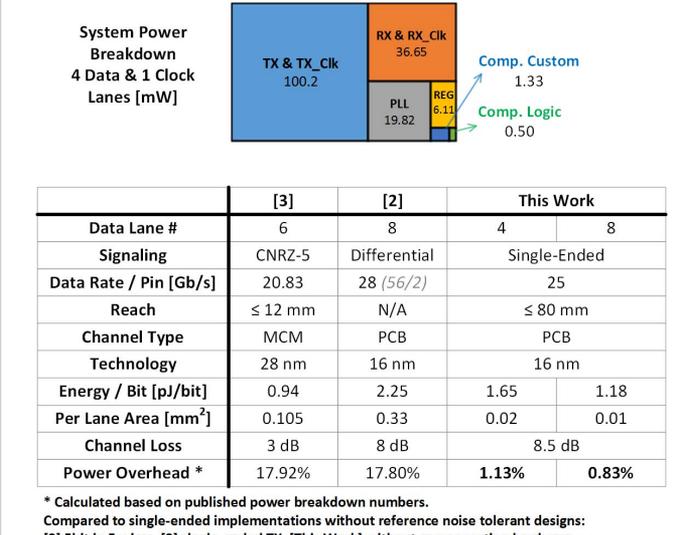


Figure 6.6.6: Power breakdown for this work and comparison to prior work.

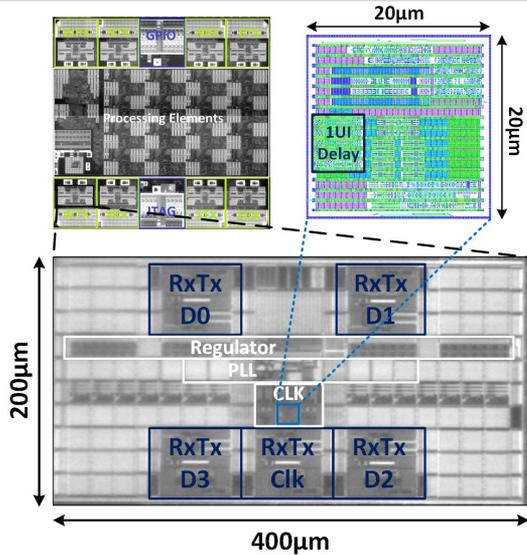


Figure 6.6.7: Die photo.

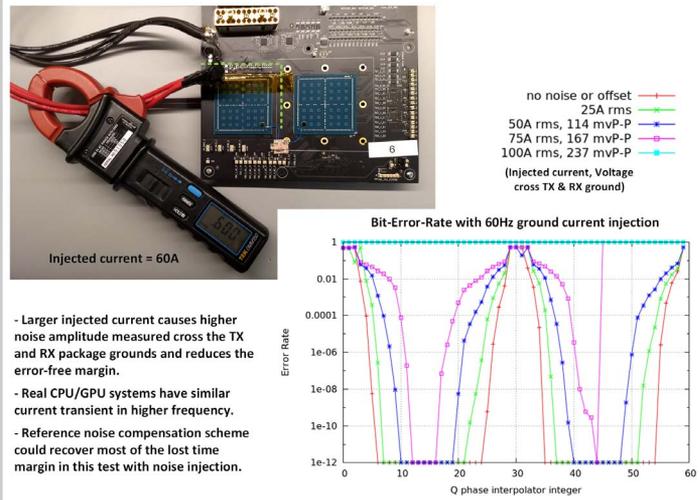


Figure 6.6.S1: Package-to-package link test setup and measured BER with 60Hz ground current injection.

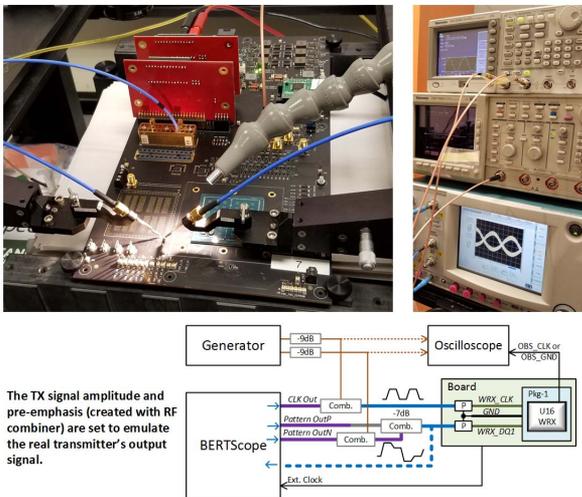


Figure 6.6.S2: RX compensation test uses BERTScope and signal generator as transmitter to deliver the 25Gb/s PRBS-31 pattern with embedded 1~30MHz noise through TX side PCB pad probing.

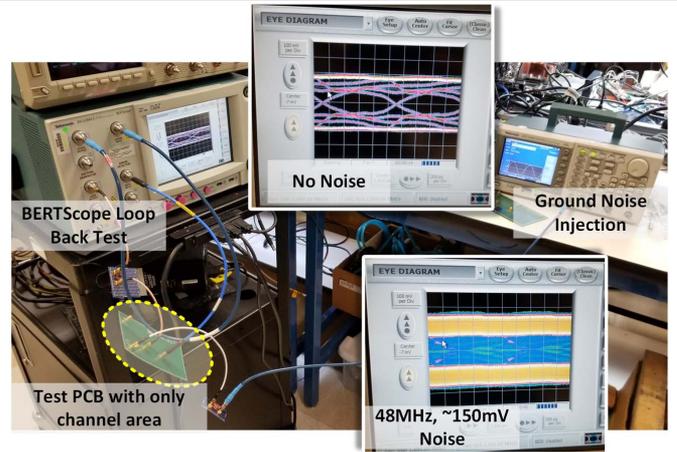


Figure 6.6.S3: Ground noise injection test with channel-only PCB and BERTScope loop-back. Since it's difficult to push high frequency noise into the real test board due to low ground impedance, we made this replica channel part PCB to verify the reference noise problem still exist at frequencies of up to tens of megahertz.