

24.5 25Gb/s 3.6pJ/b and 15Gb/s 1.37pJ/b VCSEL-Based Optical Links in 90nm CMOS

Jonathan Proesel, Clint Schow, Alexander Rylyakov

IBM T. J. Watson, Yorktown Heights, NY

Future high-performance computing systems require sub-2pJ/bit power efficiencies at >10Gb/s [1-2]. The best reported optical link efficiencies at these data rates are ≥ 2.5 pJ/bit [1-4]. This paper describes two VCSEL-based multimode (MM) fiber optical links achieving sub-2pJ/bit power efficiency from 15Gb/s to 22Gb/s. The links realized in 90nm CMOS share the same TX but use two different RXs that explore different aspects of the power/BW/area tradeoff.

The system-level block diagram of the optical link and the TX circuits are shown in Fig. 24.5.1. The laser diode driver (LDD) TX consists of a 2-stage preamplifier and a main driver stage (DRV). The preamplifier stages are Cherry-Hooper (CH) amplifiers with 1.6V supply V_{DD_PA} . The DRV is an inductively-peaked differential amplifier, one side wire-bonded to a 5 μ m diameter 850nm GaAs laser diode (VCSEL) [5] and the other side terminated to 60 Ω on chip. The DRV is powered by a 1V supply, V_{DD_DRV} . The VCSEL anode is biased by V_{VCSEL} and I_{LDB} tunes the VCSEL DC bias current. MM fiber carries the optical signal from VCSEL to a 25 μ m diameter GaAs photodiode (PD). Lensed MM fiber probes provide optical coupling between MM fiber and VCSEL and PD. The PD capacitance is 80fF and the responsivity is 0.55A/W [5]. Bond wire inductances for VCSEL and PD are estimated to be 300pH. The PD cathode is biased by a 3V supply, V_{PD} .

The two RX designs in this paper are called the T-coil RX and the CMOS inverter RX. The T-coil RX targets high BW and low power at the expense of area (mostly occupied by the T-coils). The CMOS inverter RX targets low power and small area for high integration density at the expense of BW. Simulation results below include extracted layout parasitics and, for TIA and RX, a model of the wirebonded PD.

The T-coil RX block diagram and circuits are shown in Fig. 24.5.2. The TIA is a pair of CMOS inverters with resistive feedback, one active and one a replica, providing pseudodifferential power-supply noise rejection. The CMOS TIA is chosen because of its excellent gain and noise at low power consumption due to the reuse of current in PMOS and NMOS, maximizing g_m . For DC-offset compensation, V_{FB} steers tail current I_{DC} through differential pair M_{D1} and M_{D2} into the TIA input nodes. The simulated TIA has 193 Ω gain and 30GHz BW. The limiting amplifier (LA) is a 5-stage amplifier with 34dB gain and 23GHz BW in simulation. The LA stages are differential amplifiers with T-coils for BW extension [6]. T-coils provide greater BW extension than inductors at similar area cost and, like inductors, work with low supply voltages. The LPF amplifies the difference between the DC levels at V_{LA_OUT} and returns it to the TIA as V_{FB} for offset compensation. The LPF is a single-pole RC filter using a Miller-boosted 2.2pF capacitor. The Miller-boosting amplifier is a 3-stage amplifier composed of resistively-loaded differential amplifiers with inverse scaling and an 870fF capacitor for frequency compensation. The RX output stage (OUT) is an inductively-peaked differential amplifier; a load resistance of 75 Ω is chosen as a compromise between output signal amplitude and termination to the off-chip 50 Ω environment. The nominal supply voltage for the TIA, LA, OUT, and LPF is 1.2V. The T-coil RX has 8.2k Ω gain, 20GHz BW, 2.4 μ A_{rms} input-referred noise, and 8MHz low frequency cutoff in simulation.

The CMOS inverter RX block diagram and circuits are shown in Fig. 24.5.3. CMOS inverter I1 with resistive feedback forms the TIA, while inverter I2 acts as a transconductor to sink or source current at the input for self-biasing. Simulated TIA gain is 194 Ω and BW is 29GHz. The LA is a 4-stage amplifier with 34dB gain and 14GHz BW in simulation. Each stage of the LA is realized as two inverters, I3 and I4, with resistive feedback around I4 to boost the BW, similar to the CH topology used in the TX. The LPF provides feedback for self-biasing and is realized as a single-pole RC filter with a 1.1pF capacitor Miller-boosted by inverter I5. The output stage is an NMOS source follower (SF) driving an off-chip 50 Ω load. The nominal supply voltage for the CMOS inverters and SF is 1.2V. The CMOS inverter RX has 6.3k Ω gain, 12GHz BW, 1.9 μ A_{rms} input-referred noise, and 2MHz low frequency cutoff in simulation.

Because the CMOS inverter RX is single-ended, it is vulnerable to power supply noise, requiring decoupling or supply regulation for use in noisy environments (e.g., multi-channel RX). Supply regulation can be a net benefit after the power efficiency loss, as power, gain, and BW can be reduced by lowering the supply voltage, allowing the CMOS inverter RX performance to be tuned to usage conditions. No regulator is included on these chips.

The TX and RX sites are wire-bonded to a high-speed custom PCB for testing (Fig. 24.5.7). The LDD TX occupies 80 \times 170 μ m², the T-coil RX occupies 250 \times 390 μ m², and the CMOS inverter RX occupies 40 \times 95 μ m². At nominal supply voltages, the TX uses 36.2mW (PA 16mW, DRV 8mW, and VCSEL 12.2mW); the T-coil RX uses 44.4mW (TIA 6mW, LA and LPF 27.6mW, and OUT 10.8mW); and the CMOS inverter RX uses 25.2mW (TIA, LA, and LPF 20.4mW, SF 4.8mW). The VCSEL biases are hand-tuned. The VCSEL DC bias current is 4.2mA, sufficiently low for long-term reliability. The VCSEL emitted 0.73mW average optical power with 5.1dB extinction ratio (ER), giving 0.77mW (-1.1dBm) optical modulation amplitude (OMA).

Bathtub and sensitivity measurements (Fig. 24.5.4) are taken with a 4m MM fiber and an adjustable optical attenuator. The test pattern is PRBS7. Bathtub curves are measured with 230 μ A DC and 240 μ A_{pp} modulated photocurrent. The two RXs have bathtub eye openings at 10⁻¹⁰ BER within a few percent of each other, pointing to the TX as a limiting factor in the links. The lower sensitivity in the T-coil RX is due to the larger T-coil RX BW allowing more high-frequency noise to degrade sensitivity. The input-referred noise is 2.3 μ A_{rms} for the T-coil RX and 2.0 μ A_{rms} for the CMOS inverter RX, closely agreeing with simulation. Sensitivity losses of 4dB and 2dB are observed moving from PRBS7 to PRBS31 at 15Gb/s for the T-coil RX and CMOS inverter RX, respectively. A low transmission penalty of 1dBm is observed for both RXs after 100m of MM fiber at 20Gb/s.

The measured 25Gb/s TX optical eye, T-coil RX electrical eye, and bathtub curve are presented in Fig. 24.5.5. To reach 25Gb/s, V_{DD_PA} is increased to 2V, resulting in 46mW TX power; T-coil RX power is unchanged. The VCSEL outputs 0.91mW (-0.4dBm) OMA and 5.7dB ER. The RX receives 340 μ A DC and 390 μ A_{pp} modulated photocurrent. The TX optical eye exhibits asymmetric eye closure which is reflected in the RX eye. For both eyes, measured BER is <10⁻¹² in the center of the eye. The RX eye opening is 22% at 10⁻¹⁰ BER. The full-link power efficiency is 3.6pJ/bit at 25Gb/s.

To determine the best achievable power efficiency versus data rate, both links are power optimized at data rates ranging from 10Gb/s to 25Gb/s. Power is optimized by reducing supply voltages and bias currents while maintaining BER <10⁻¹² in the center of the eye and RX single-ended output voltage ≥ 100 mV_{pp}. The results are plotted in Fig. 24.5.6. The full link with CMOS inverter RX achieves a power efficiency of 1.37pJ/bit at 15Gb/s.

Acknowledgements:

The authors thank C. Baks for high-speed custom PCB design, Y. Vlasov for management support, and Emcore Corp. for the VCSELs and PDs.

References:

- [1] I. Young *et al.*, "Optical I/O technology for tera-scale computing," *J. Solid-State Circuits*, vol. 45, no. 1, pp. 235–248, Jan. 2010.
- [2] J. Kash *et al.*, "Optical interconnects in future servers," *Proc. Optical Fiber Communications Conf. (OFC)*, paper OWQ1, Mar. 2011.
- [3] C. Schow *et al.*, "A <5mW/Gb/s/link, 16 \times 10Gb/s bi-directional single-chip CMOS optical transceiver for board-level optical interconnects," *ISSCC Dig. Tech. Papers*, pp. 294–295, Feb. 2008.
- [4] C. Kromer *et al.*, "A 100mW 4 \times 10Gb/s transceiver in 80nm CMOS for high-density optical interconnects," *ISSCC Dig. Tech. Papers*, pp. 334–335, Feb. 2005.
- [5] N. Li *et al.*, "High-performance 850 nm VCSEL and photodetector arrays for 25 Gb/s parallel optical interconnects," in *Proc. Optical Fiber Communications Conf. (OFC)*, paper OTuP2, Mar. 2010.
- [6] T. Toifl *et al.*, "A 23GHz differential amplifier with monolithically integrated T-coils in 0.09 μ m CMOS technology," *IEEE MTT-S Int. Microw. Symp. Dig.*, pp. 239–242, Jun. 2003.

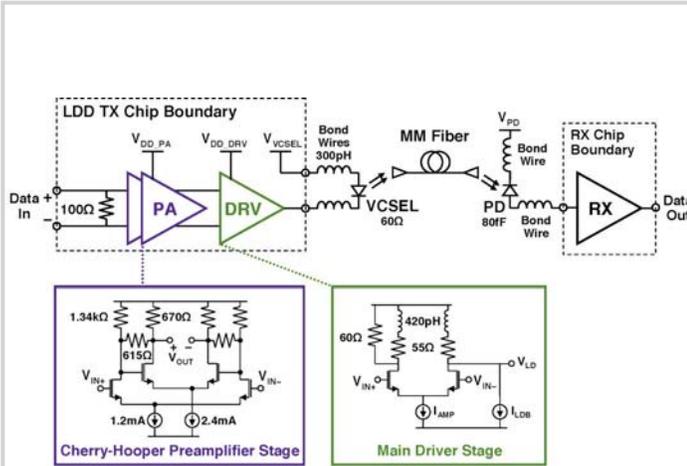


Figure 24.5.1: Top level diagram of the optical links and LDD TX circuits.

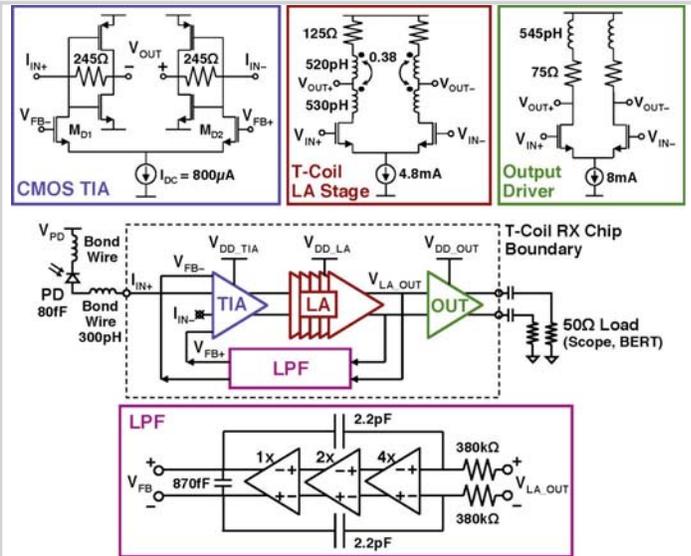


Figure 24.5.2: T-coil RX block diagram and circuits.

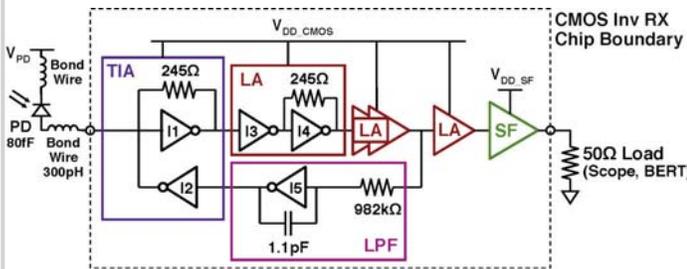


Figure 24.5.3: CMOS inverter RX block diagram.

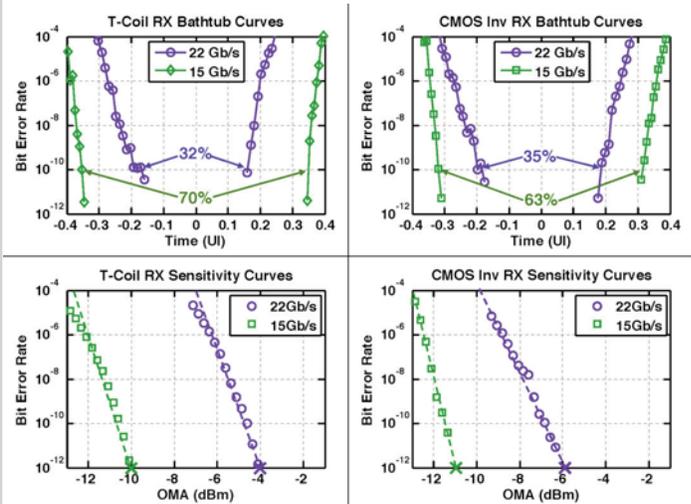


Figure 24.5.4: Measured full-link bathtub and sensitivity curves for the T-coil RX and CMOS inverter RX at nominal supply voltages.

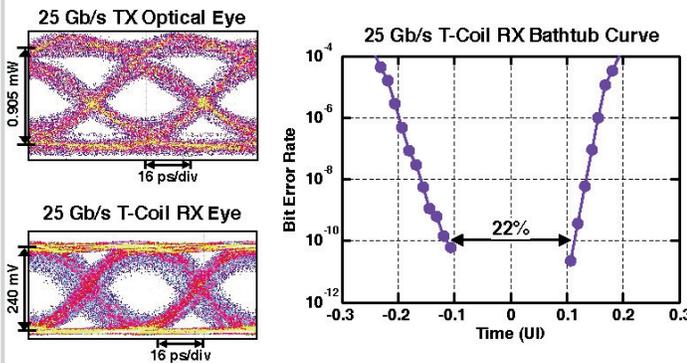


Figure 24.5.5: 25Gb/s eye diagrams and bathtub curve for full link with T-coil RX through 4m MM fiber.

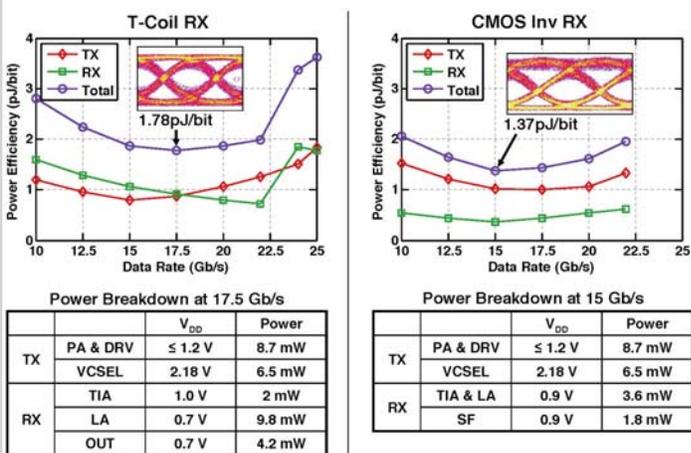


Figure 24.5.6: Optimized power efficiency vs. data rate for full links through 4m MM fiber. RX eyes and power breakdowns are shown at the best power efficiency points.

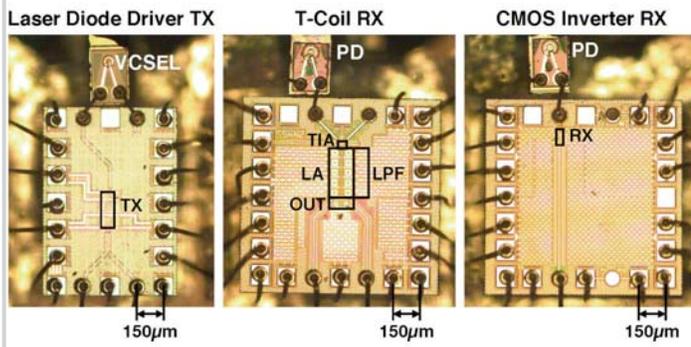


Figure 24.5.7: Photomicrographs of the wirebonded laser diode driver TX, T-coil RX, and CMOS inverter RX. Wirebond pad pitch is 150µm (for scale).