

## 9.2 An 800mW 10-Gigabit Ethernet Transceiver in 0.13 $\mu$ m CMOS

Stefanos Sidiropoulos, Nikhil Acharya, Pak Chau, Joseph Dao, Arnold Feldman, Haw-Jyh Liaw, Marc Loinaz, R. Sekhar Narayanaswami, Clemenz Portmann, Shahriar Rabii, Azmeer Salleh, Siddharth Sheth, Lars Thon, Katelijin Vleugels, Patrick Yue, Don Stark

Aeluros, Mountain View, CA

The higher performance of fine-line fabrication processes has allowed CMOS ICs to be used in communication applications previously dominated by compound semiconductor devices, thus enabling the integration of increased functionality on a single IC. This paper presents a 10 Gigabit Ethernet physical layer IC which includes all the logic necessary to implement the 10Gb/s extender (XGXS) and packet coding (PCS) sub-layers as defined in the IEEE standard [1]. The chip clocking micro-architecture and the integrated 10.3Gb/s and 4x3.12Gb/s custom transceiver blocks support arbitrarily long Ethernet packets as well as SONET applications [2]. Extensive use of CMOS logic circuits and resonant distribution of 10GHz clocks help achieve a power dissipation of 800mW.

Figure 9.2.1 shows a block diagram of the transceiver IC. In the electrical to optical path, four independent CDRs retune and demultiplex four lanes of 3.12Gb/s data. After 8b/10b decoding, the four streams are aligned and written into an elastic FIFO. The data is read out of the FIFO, processed by a 64/66b encoder/scrambler, and serialized into a 10.3Gb/s stream. The reverse operation is performed by the optical to electrical path: the 10.3Gb/s data is retimed by a half-rate interpolating DLL/CDR, de-multiplexed, decoded and written into a second elastic FIFO. The data is read out of the FIFO and serialized into four 3.12Gb/s streams. A clean-up PLL can be used to filter the system reference clock (clkRef), or the recovered clocks (clkRx, clkTx). In the normal mode of operation where all the PLL reference clocks are driven by clkRef, the elastic FIFOs can compensate for up to 500ppm of frequency offset between their read and write clocks, provided that the data stream does not contain packets longer than 10KB. However, for synchronous network applications or for applications requiring unlimited packet size, the reference clocks of the transmit PLLs need to be derived from the corresponding recovered clocks. To achieve this without compromising jitter performance or increasing the system cost, the transceiver can operate in single source plesiochronous (SSP) mode. In SSP mode, the device uses a single VCXO both to aid the CDR acquisition and to filter the jitter of the recovered clock.

Figure 9.2.2-(a) shows a block diagram of the electrical to optical clock path operating in SSP mode. The DLL/CDR [3] uses four clock phases generated by the 1.56GHz PLL and continuously adjusts the phase of the sampling clocks such that it coincides with that of the incoming 3.12Gb/s data. The CDR can acquire lock as long as the frequency of the data is within  $\pm 400$ ppm of the frequency of the local clock generated by the VCXO. The clock recovered by the CDR also serves as the reference for the clean-up PLL, seemingly creating a loop of PLLs in which neither phase nor frequency are precisely defined. However, the limited oscillation range of the VCXO ensures that the CDR will acquire lock. The clock recovered by the CDR will in turn force the clean-up PLL to be synchronous with the incoming data. Figure 9.2.2-(b) shows a linear model of this control loop. Stability is ensured by the fact that the clean-up PLL bandwidth is at least two orders of magnitude lower than that of the 1.6GHz PLL and CDR.

Minimizing power consumption beyond the point achievable by the scaled supply of a CMOS process requires extended use of switching CMOS circuits. The speed of the 0.13 $\mu$ m CMOS gates available (50ps delay for a fanout-4 inverter) enables the use of CMOS switching circuits up to the 5GHz clock domain. CML circuits are used in the 10GHz clock paths of both the 10Gb/s transmitter and receiver. In order to reduce the power dissipated in distributing 10GHz differential clocks, the design makes extensive use of resonant clock buffers.

Figure 9.2.3 shows the circuit diagram of the 10GHz VCO and the first stage resonant buffer used in both 10GHz PLLs; a similar buffer fans-up the output to drive the final 10GHz clock. An NMOS VCO topology is chosen to minimize the tank parasitic capacitance. A binary weighted array of low-loss metal-to-metal capacitors selects one of the eight overlapping frequency bands, and NMOS in well varactors provide continuous fine tuning within the band. The dc-coupled VCO buffer uses small input devices ( $M_6$ - $M_7$ ) to minimize the capacitance of the main VCO tank and cascodes  $M_6$ - $M_7$  to provide good reverse isolation. To increase the output swing of this stage, a positive feedback pair  $M_8$ - $M_9$ , significantly boosts the gain [4]. A calibrating FSM optimizes the swing and common mode of the resonant buffer chain under varying process, environmental and target frequency conditions. Simulation results indicate that the use of purely resonant rather than peaked CML buffers saves approximately 60mW of power, and improves the yield of the 10GHz retiming chains by maintaining a large output swing even under worst case fabrication process and operating conditions.

The transceiver IC is fabricated in a 1.2/2.5V 0.13 $\mu$ m CMOS logic process with one poly and five metal layers, and does not use special options such as deep N-well or thin insulator MiM capacitors. The 2.5x5mm<sup>2</sup> die (Fig. 9.2.4) is packaged in a 144-pin 13x13mm<sup>2</sup> wire-bond plastic BGA. Figure 9.2.5 shows the jitter transfer function measured with the transceiver operating in line loop-back mode with and without SSP clocking enabled. To achieve the same low transfer bandwidth without SSP would require the use of an additional VCXO. Figure 9.2.6a shows the output data eye-diagram with the chip operating with a 1.14V supply at 85°C case temperature while transmitting 2<sup>31</sup>-1 PRBS data. The maximum peak-to-peak jitter is 6.2ps, while its variation over temperature is less than 1.4ps as shown in Fig. 9.2.6b. Figure 9.2.a shows the input stressed eye used to evaluate the 10.3Gb/s receiver jitter tolerance [1]. The jitter tolerance results in Fig. 9.2.7b confirm that the receiver exceeds the IEEE specification while operating with a 20mV peak-to-peak differential input signal. The operating ranges of the electrical and optical interfaces were measured to be 0.5 to 4.5Gb/s and 8.5 to 10.8Gb/s, respectively. Operating from 1.2V and 2.5V supplies the transceiver dissipates 0.8W: 340mW and 310mW on the 10.3Gb/s and 3.12Gb/s interfaces respectively, and 140mW on the digital data-path.

### Acknowledgments:

The authors would like to thank N. Gamini for his packaging expertise, D. Ballinger for the initial implementation of the management interface, B. Cunnie for being a rock-star, and C. Moriyama, M. Speers and J. Chan for layout support.

### References:

- [1] "CSMA/CD Access Method, PHY Specification, MAC, PHY, and Management Parameters for 10Gb/s," IEEE 802.3-ae, 2002.
- [2] SFI-4 Phase 2 (OC-192 4-bit SERDES to framer interface), Optical Networking Forum, OIF-SFI4-02.0, Sept. 2002.
- [3] S. Sidiropoulos and M. Horowitz, "A Semi-Digital Dual Delay-Locked Loop," *IEEE J. Solid-State Circuits*, vol. 32, no. 11, pp. 1683-1692, Nov. 1997.
- [4] K.C. Tsai and P. Gray, "A 1-W CMOS Class-E Power Amplifier for Wireless Communications," *IEEE J. Solid-State Circuits*, vol. 34, no. 7, pp. 962-970, July 1999.

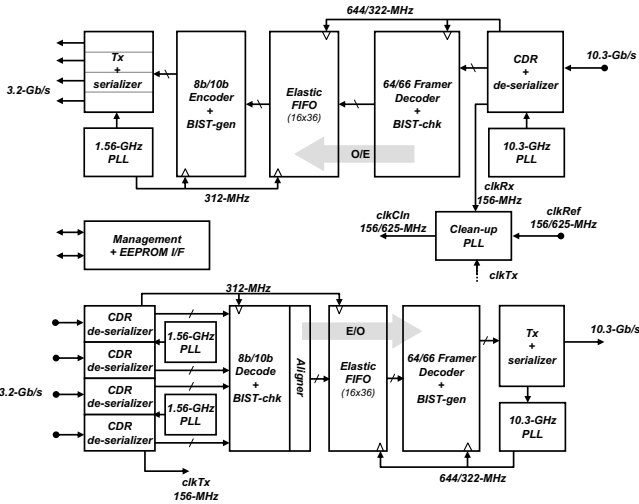




Figure 9.2.1: Transceiver block diagram



(a)



(b)

Figure 9.2.2: SSP clock path (a), and linear model (b)

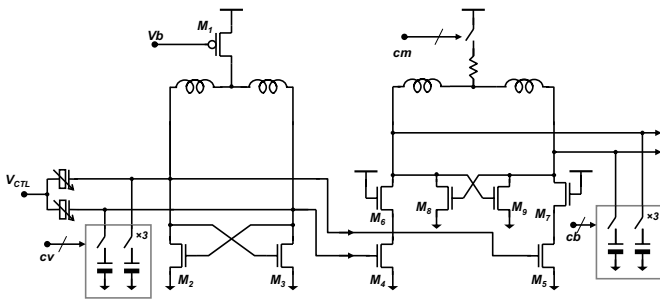


Figure 9.2.3: 10-GHz VCO and 1st stage clock buffer




Figure 9.2.4: Die micrograph






Figure 9.2.5: Measured jitter transfer function: conventional vs SSP




(a)




(b)

Figure 9.2.6: 10.3125-Gbps output data eye (a), and peak-to-peak jitter vs temperature (b)

Continued on Page 000



(a)



(b)

Figure 9.2.7: Jitter tolerance: (a) calibrated input eye, (b) tolerance results