26.3 An 800MS/s 10b/13b Receiver for 10GBASE-T Ethernet in 28nm CMOS

Jan Mulder[1], Davide Vecchi[1], Yi Ke[1], Stefano Bozzola[1], Mark Core[2], Nitz Saputra[3], Qiongna Zhang[1], Jeff Riley[1], Han Yan[1], Mattia Introini[1], Sijia Wang[1], Christopher M. Ward[1], Jan Westra[1], Jiansong Wan[1], Klaas Bult[1]

[1]Broadcom, Bunnik, The Netherlands,
[2]Broadcom, Irvine, CA, [3]Qualcomm, San Diego, CA

The IEEE 802.3an standard describes full-duplex 10Gb/s Ethernet transmission over four pairs of up to 100m UTP cable. The performance required from the analog front end (AFE) of a 10GBASE-T Ethernet transceiver strongly depends on the length of the cable connected to it. Maximum-length cables require the highest performance, and hence, determine the worst-case power dissipation of the transceiver. In practice, however, the vast majority of cable lengths used are below 30m. For these shorter cables, the standard specifies the transmitted power level to be lowered, inherently leading to a reduction in power consumption of the transmitter (TX). In most designs, the power consumption of the receiver (RX), unfortunately, does not benefit from the shorter cable lengths [1,2]. This paper presents a power-efficient 13b RX, implemented in 28nm CMOS. By switching to a 10b mode for short cables, 143mW is saved in the AFE for one complete Ethernet port, comprising four receivers. In addition, to further reduce power, the RX heavily relies on calibrations.

The architecture of the AFE is depicted in Fig. 26.3.1. The TX directly drives the UTP cable. A hybrid circuit, comprising a replica DAC and a resistor R_{hyb}, cancels the transmitted signal at the input of the local RX [1,2]. After the hybrid, a programmable-gain amplifier (PGA) is used to optimize the loading of the subsequent ADC and to provide some filtering of the input signal. Separate PGAs are used for the two modes that the RX can operate in. In 13b mode, PGA_1 connects to the input of a 13b 800MS/s SHA-less pipeline ADC, which consists of two 4b MDAC stages followed by a 7b subranging ADC. At short cable lengths, higher noise floor and distortion levels can be tolerated. Therefore, in 10b mode, considerable power can be saved by using a dedicated low-power PGA_2 that bypasses the first MDAC stage and directly connects to $MDAC_2$ through a MUX. Whereas PGA_1 uses a two-stage amplifier to obtain sufficient linearity (similar to [2]), a single-stage design is sufficient for PGA_2. To save more power in 10b mode, a series resistor $R_{in,2}$ is inserted to increase the impedance of the feedback network of PGA_2. The MUX is implemented by bootstrapped thick-oxide NMOS transistors.

Each MDAC stage shares a residue amplifier (RA) between two sets of input capacitors that are 2× time interleaved. Offset, gain and timing mismatches due to the time interleaving are calibrated in the background. The offset and gain errors are corrected in the digital domain; the timing mismatch can be corrected in the analog domain through 6b control of the timing of the sampling circuit of $MDAC_1$. To save area, the 4b flash ADCs embedded in the MDACs and the subranging ADC do not use time interleaving, but run at 800MS/s. $MDAC_1$ uses a flip-around single-stage RA with a nominal 6× gain. $MDAC_2$ is scaled down in power and has a gain of 8×. $MDAC_2$ uses a two-stage RA, comprising a complementary input stage, folded cascode, and two parallel source-followers at the output [3], one driving the feedback capacitor and one driving the subranging ADC.

Calibration of the gain of the RAs allows their open-loop gain and settling bandwidth requirements to be relaxed. As a result, their power consumption can be reduced significantly. To calibrate $MDAC_1$, a calibration signal V_{cal} is intermittently (controlled by a PRBS sequence) subtracted from the residue voltage of $MDAC_1$. Correlation between the output of $MDAC_1$ and the PRBS is used to drive the calibration loop. To help the accuracy and the convergence speed of the calibration, V_{cal} equals half the size of a nominal subrange. During the foreground calibration, performed at startup, the magnitude of V_{cal} is not a problem, as no other signals are being processed. However, during background calibration, V_{cal} can take up the complete overrange budget of $MDAC_1$. Doubling the flash ADC resolution to solve this issue is an expensive solution in terms of area and power. As an alternative, only two comparators and one preamplifier are added to the 4b flash ADC, as illustrated in Fig. 26.3.2. The background calibration is now performed in only two of the 16 subranges of $MDAC_1$, which is sufficient for this application. As shown in Fig. 26.3.2, when $V_{th,6b} < V_{in} < V_{th,7}$, V_{cal} can safely be subtracted from V_{in} without reducing the overrange budget at all.

In the reported ADC, two techniques for correcting the RA gain errors are implemented. The first method performs the correction in the digital domain. To save power, the second approach performs the correction in the analog domain. Figure 26.3.3 shows the implementation. By tuning the reference voltage $V_{ref,2}$ of $MDAC_2$, the gain error of $MDAC_1$ can be corrected. An op amp and a class-AB output stage are used to generate $V_{ref,1}$ for $MDAC_1$. A (scaled-down) replica output stage produces $V_{ref,2}$. It is driven by a voltage $V_{g,2}$, which is equal to the op amp output voltage $V_{g,1}$ minus a calibration voltage $V_{g,cal}$. A 9b DAC is used to tune $V_{g,cal}$, such that $V_{ref,2}$ is compliant with the actual gain of $MDAC_1$. A third replica output stage, not shown, creates $V_{ref,3}$ for the subranging ADC; it is calibrated through another 9b DAC.

By design, pipeline ADCs are mostly insensitive to comparator offsets. Calibration of the comparator offsets, however, brings additional benefits. Besides improving robustness and power savings in the preamplifiers driving the comparators, calibration reduces the average output swing of the RAs. As a result, the distortion contributed by the nonlinearity of the RAs is reduced, leading to power savings in their design.

Since a calibration DAC (calDAC) is required in each of the comparators, the implementation of this calDAC must be very area-efficient. The implementation of the 7b calDAC used is shown in Fig. 26.3.4. At the input of the comparator, an additional PMOS differential pair connects to the NMOS latch and is driven by a calibration voltage V_{cal}. The differential voltage V_{cal} comprises two voltages, i.e., $V_{cal} = \pm(V_{cal,MSB} - V_{cal,LSB})$. Four switches, driven by the MSB of a 7b counter, determine the polarity of V_{cal}. The next three bits of the counter drive a multiplexer that selects the voltage $V_{cal,MSB}$ from a set of eight voltages $V_{MSB,i}$, where $i = [0\ldots7]$. Likewise, the three LSBs of the counter are used to select $V_{cal,LSB}$ from eight voltages $V_{LSB,i}$. Six XOR gates are required to invert the 6 LSBs when the MSB of the counter changes. The voltages $V_{MSB,i}$ and $V_{LSB,i}$ are generated only once, by a single resistive divider, and are shared by all comparators. The switches and multiplexers use only NMOS transistors, and hence require very little area.

At startup, a foreground calibration of the comparator offsets is performed. To this end, the inputs of the preamplifiers are shorted and each comparator output directly drives the input of its up/down counter. As a result, the comparator offset voltage converges to zero. During normal operation, the outputs of the flash ADCs are occasionally analyzed for usage of the MDAC over-range. The analysis is performed in the background by firmware running on the on-chip processor. If too much of the over-range is used in some sub-ranges, the counters in the corresponding comparators are adjusted.

The RX is integrated as part of a dual-port 10GBASE-T Ethernet transceiver, fabricated in 28nm CMOS. Figure 26.3.7 shows a die photograph of the RX, which occupies an area of $0.23mm^2$. The area of the ADC core is $0.13mm^2$. In the measurement setup, the on-chip TX is used as the input signal to the RX. The ADC clocks are provided by the on-chip PLL. All calibration algorithms are implemented on-chip.

The RX operates from a dual 1V/1.8V supply. Figure 26.3.5 shows the measured SNDR, SNR, THD and SFDR in 13b mode as a function of the input frequency F_{sig}, while running at 800MS/s. At low frequencies, the SNDR drops due to the bandpass characteristic of the PGA. The RX achieves a peak SNDR of 9.54b at F_{sig} = 150MHz. By shorting the output of the PGA, the dynamic range of the ADC was measured to be 10.6b. Figure 26.3.6 shows the RX output spectrum in 10b and 13b mode, respectively, for a two-tone input signal centered around 350MHz. The RX consumes 76.4mW per channel in 13b mode, which is 62% lower than the RX in [2]. In 10b mode, it consumes 40.6mW, thus saving another 143.2mW for one complete Ethernet port.

References:

[1] R. Farjad, *et al.*, “A 48-port FCC-compliant 10GBase-T transmitter with Mixed-Mode Adaptive Echo Canceller,” *IEEE J. Solid-State Circuits*, vol. 47, no. 12, pp. 3261-3272, Dec. 2012.
[2] J.R. Westra, *et al.*, “A sub-1.75W Full-Duplex 10GBASE-T Transceiver in 40nm CMOS,” *ISSCC Dig. Tech Papers*, pp. 146-147, Feb. 2014.
[3] C.-K. Lee, *et al.*, “A Replica-Driving Technique for High Performance SC Circuits and Pipelined ADC Design,” *IEEE Trans. Circuits and Systems-II*, vol. 60, no. 9, pp. 557-561, Sept. 2013.

 978-1-4799-6224-2/15/$31.00

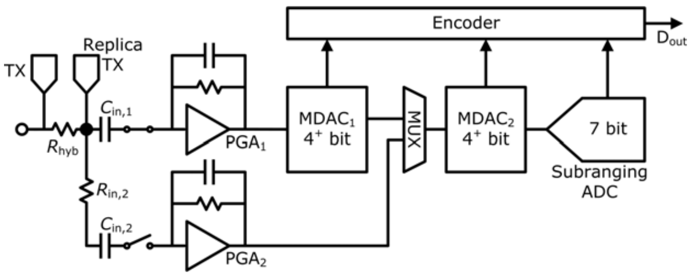

Figure 26.3.1: Dual-mode 10GBASE-T receiver.

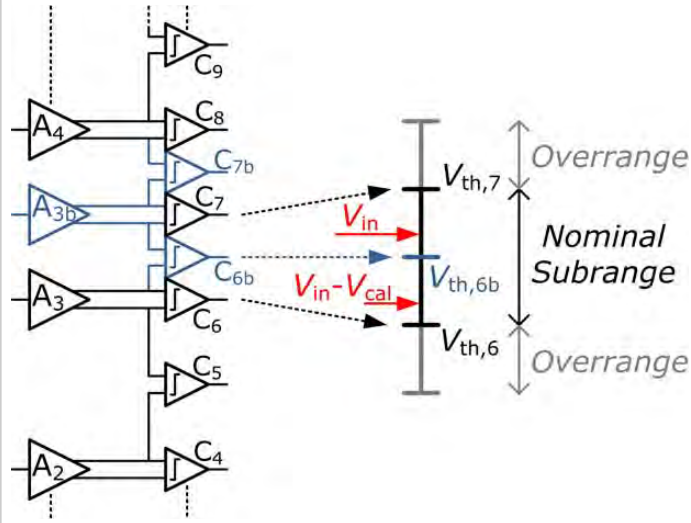

Figure 26.3.2: Two additional comparators create room for the calibration signal.

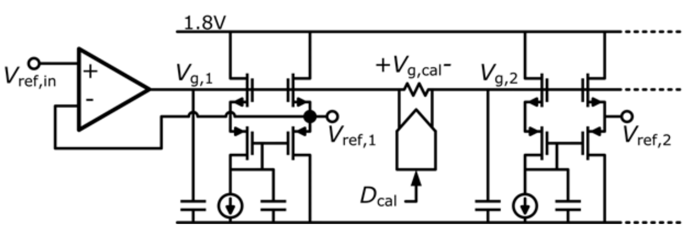

Figure 26.3.3: Reference buffers used for analog calibration of the RA gain errors.

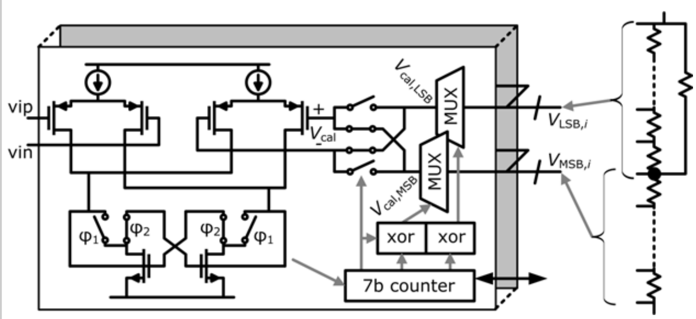

Figure 26.3.4: Comparator with offset calibration.

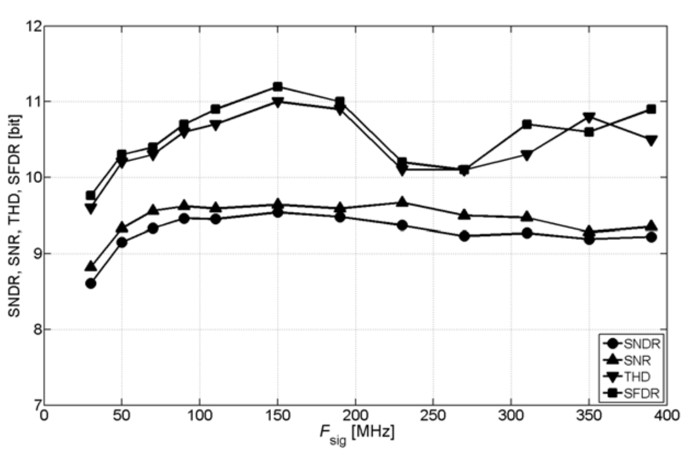

Figure 26.3.5: Sweep of SNDR, SNR, THD and SFDR vs F_{sig}.

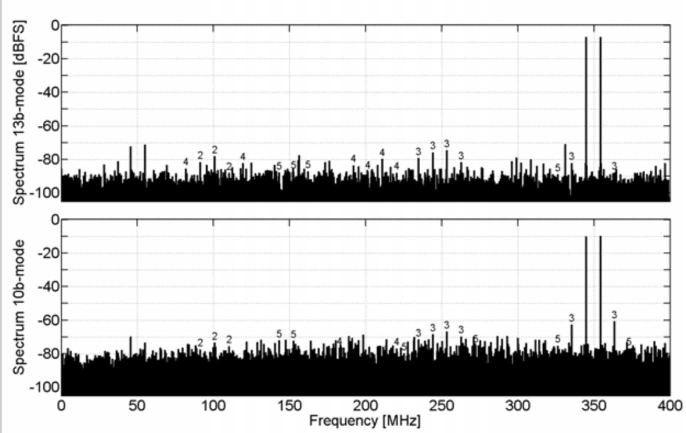

Figure 26.3.6: Output spectrum of the RX in 13b mode and 10b mode.

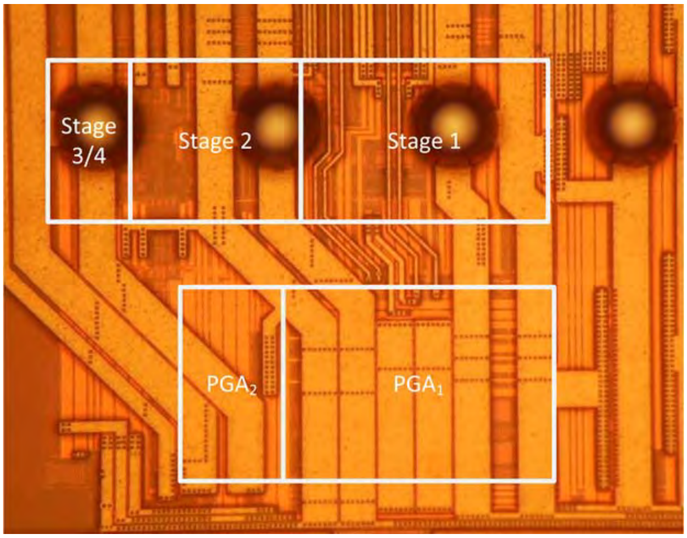

Figure 26.3.7: Die photograph of the RX.