

22.9 A 0.28pJ/b 2Gb/s/ch Transceiver in 90nm CMOS for 10mm On-Chip Interconnects

Eisse Mensink, Daniël Schinkel, Eric Klumperink, Ed van Tuijl, Bram Nauta

University of Twente, Enschede, The Netherlands

The bandwidth of global on-chip interconnects in modern CMOS processes is limited by their high resistance and capacitance [1]. Repeaters that are used to speed up these interconnects consume a considerable amount of power [2] and area. Recently published techniques [1-4] increase the achievable data rate at the cost of high static power consumption, leading to relatively high energy per bit for low data activity. On the other hand, low-swing schemes [5] often sacrifice bandwidth for power reduction, or make use of an extra low-voltage power supply. More ideally, a transceiver would combine low dynamic and static power with a high achievable data rate.

The bandwidth and power consumption of an RC-limited interconnect depends on its source (Z_s) and load impedances (Z_L). In Fig. 22.9.1, a conventional case with an inverter used as both a transmitter ($Z_s = 100\Omega$) and a receiver ($Z_L = 10\text{fF}$) has only 62MHz bandwidth and high power consumption. Current-sensing schemes ($Z_L = 190\Omega$ in Fig. 22.9.1) increase the bandwidth up to $3\times$ [1,4], but with increased power at low data activities. We propose using a capacitive transmitter ($Z_s = 255\text{fF}$ in Fig. 22.9.1), which has the same bandwidth improvement as current sensing, but with lower power and without static power consumption.

This paper presents a transceiver for 10mm long interconnects in a 1.2V 90nm 6M CMOS process, shown in Fig. 22.9.2. A capacitive pre-emphasis transmitter both increases the bandwidth and decreases the voltage swing, without the need for an additional power supply. As low-swing signaling is more susceptible to crosstalk, we use differential interconnects with twists [1], of which only a single-ended half is shown. In contrast to the wide interconnects used in [2,3], we use relatively small width (0.54 μm) and spacing (0.32 μm) [1,4] and assume high metal-density surroundings. The receiver uses decision feedback equalization (DFE) [6] to further increase the achievable data rate. The DFE, with a continuous-time feedback filter, consumes almost no extra power. The bandwidth-increasing pre-emphasis effect of the transmitter is shown at the bottom right of Fig. 22.9.2: every transition is emphasized by the transmitter by injecting a charge via capacitance C_s .

With only a series capacitor (AC-coupling), the DC voltage on the interconnect is not well defined as there is no DC path to one of the supplies. To control the DC voltage, a load resistor R_L and a transconductance G_m controlled by V_{in} are added (see Fig. 22.9.2). By having the time constants C_s/G_m and $R_L C_{\text{wire}}$ equal, the transfer function resembles the transfer function of the capacitive transmitter in Fig. 22.9.1. If a small G_m (5 μS) and a large R_L (16k Ω) are chosen, the static current is kept small (6 μA) and also the power consumption remains similar. G_m and R_L are implemented with MOS transistors as visible in the bottom part of Fig. 22.9.2. For C_s , the gate capacitance of an NMOS transistor is used. As the gate oxide is much thinner than the oxide between interconnects, the area consumed by C_s is relatively small (6 \times 6 μm^2). The signals, with a voltage swing of 100mV, are chosen close to V_{DD} of 1.2V, because the capacitance of the NMOS transistor is highest for a high gate-source voltage. The total area of the differential transmitter is 226 μm^2 .

The receiver concept is also shown in Fig. 22.9.2. A clocked comparator restores the low-swing line output to full swing. DFE further increases the achievable data rate. Instead of the often-used FIR filters [6], a continuous-time filter operates as the decision feedback filter. This filter cancels most of the ISI with a simple and power-efficient first-order implementation, whereas an FIR filter requires many taps.

The schematic of the receiver implementation is shown in Fig. 22.9.3. The left of the diagram shows a clocked comparator, a sense-amplifier-based flip-flop (SAFF), which consists of a differential input stage, cross-coupled inverters and an SR-latch. The outputs of the SR-latch drive the low-pass feedback filter; in this case an RC filter, implemented with pass-gates and anti-parallel gate capacitances. The filter output is coupled back into the SAFF via a second differential input stage, as shown on the right of Fig. 22.9.3. I_{eq} sets the feedback gain A (see Fig. 22.9.2). The total area of the receiver is 117 μm^2 (32 μm^2 for the DFE part).

The chip micrograph is shown in Fig. 22.9.7. The 10mm-long interconnects, placed in metal 4, have a total distributed resistance of 2k Ω and a capacitance of 2.8pF. The other metal layers are filled with GND- and V_{DD} -connected metal stripes. An external pattern generator/analyzer generates data and measures BER. The receiver clock is generated externally to adapt its phase to the eye position and to be able to measure eye widths. In an application, a simple skew circuit or a source-synchronous approach could be used to generate the proper clock phase. Eye-diagrams are measured via 50 Ω output buffers that are connected to the output of a differential interconnect.

Figure 22.9.4 shows a measured eye diagram at a data rate of 1Gb/s. The measured BER at the edges of the eye is also shown. The BER drops rapidly below a clock skew of -150ps and above 180ps, giving an eye-opening of 670ps. Data rates up to 1.35Gb/s are achieved without DFE ($I_{\text{eq}}=0$). The one- σ offset of the total transceiver is 11mV, measured over 20 samples. Due to this offset, not all samples achieve 1.35Gb/s, but all samples do achieve a slightly lower data rate of 1Gb/s. Simulations over process corners also indicate that the circuit is robust to PVT variations at a rate slightly lower than the maximum achievable data rate. Data rates up to 2Gb/s are measured with DFE. Fig. 22.9.5 shows that DFE improves the eye opening for a wide range of I_{eq} . In an application, I_{eq} can therefore be fixed at design time.

In Fig. 22.9.6, the measured energy per bit is plotted as a function of transition probability at different data rates. With random data at 2Gb/s, only 0.28pJ/b is dissipated, which is $7\times$ lower than earlier work [1,4]. The power dissipation of 0.12pJ/b at zero data activity is mainly due to the power dissipation in the SAFF, which has large transistors to get a low offset ($\sigma_{\text{off}} = 8\text{mV}$). Clock-gating can be used to eliminate power consumption during inactive periods. The DFE part of the circuit requires less than 7% of the total transceiver power, while it can increase the achievable data rate $1.5\times$.

With the presented transceiver, the same high data rates over small RC bandwidth limited on-chip interconnects are possible as with previous solutions, but with a $7\times$ lower power consumption. By using both a capacitive pre-emphasis transmitter and continuous-time DFE, a data rate of 2Gb/s is achieved over a 10mm long interconnect. The transceiver consumes only 0.28pJ/b.

Acknowledgements:

We thank Philips Research for chip fabrication, the Dutch Technology Foundation (STW, project TCS.5791) for funding and Gerard Wienk for assistance.

References:

- [1] D. Schinkel, E. Mensink, E. A. M. Klumperink, et al., "A 3-Gb/s/ch Transceiver for 10-mm Uninterrupted RC-limited Global On-Chip Interconnects," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 297-306, Jan., 2006.
- [2] A. P. Jose, G. Patounakis, and K. L. Shepard, "Pulsed Current-Mode Signaling for Nearly Speed-of-Light Intrachip Communication," *IEEE J. Solid-State Circuits*, vol. 41, no. 4, pp. 772-780, Apr., 2006.
- [3] A. P. Jose, and K. L. Shepard, "Distributed Loss Compensation for Low-Latency On-Chip Interconnects," *ISSCC Dig. Tech. Papers*, pp. 516-517, Feb., 2006.
- [4] L. Zhang, J. Wilson, R. Bashirullah, et al., "Driver Pre-Emphasis Techniques for On-Chip Global Buses," *ISLPED*, pp. 186-191, Aug., 2005.
- [5] H. Zhang, V. George, and J. M. Rabaey, "Low-Swing On-Chip Signaling Techniques: Effectiveness and Robustness," *IEEE Trans. VLSI Systems*, vol. 8, pp. 264-272, Jun., 2000.
- [6] V. Stojanovic, A. Ho, B. Garlepp, et al., "Adaptive Equalization and Data Recovery in a Dual-Mode (PAM2/4) Serial Link Transceiver," *Symp. VLSI Circuits*, pp. 348-351, Jun., 2004.

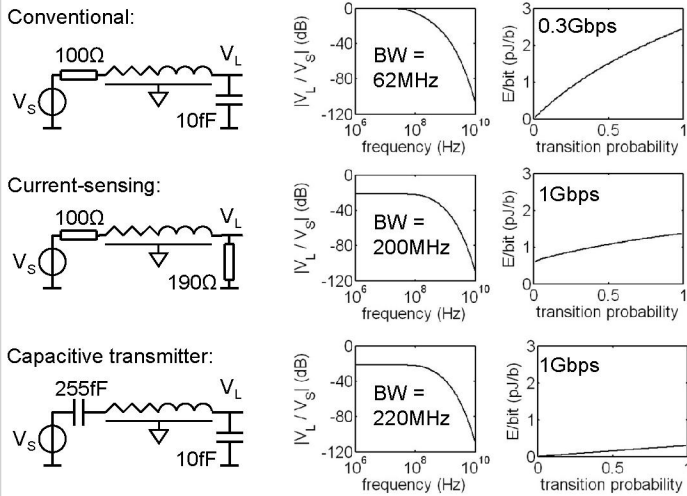


Figure 22.9.1: Bandwidth and energy per bit versus transition probability (= data activity) for three different termination schemes. The results are for 10mm differential interconnects with a distributed resistance of 2kΩ and a distributed capacitance of 2.8pF.

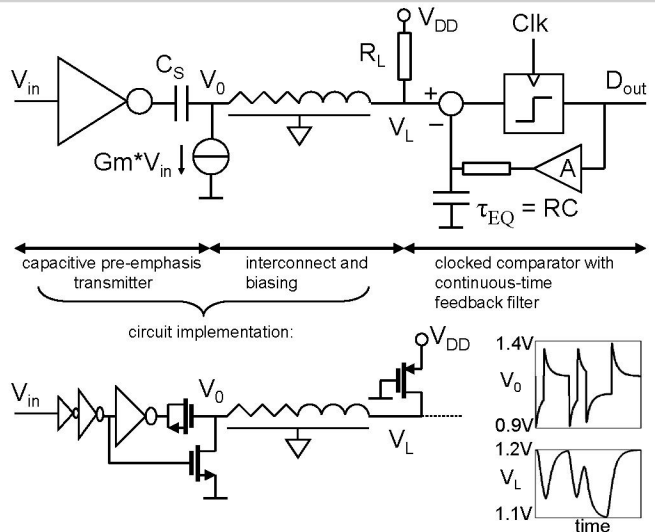


Figure 22.9.2: Concept of transceiver and circuit implementation of the capacitive pre-emphasis transmitter.

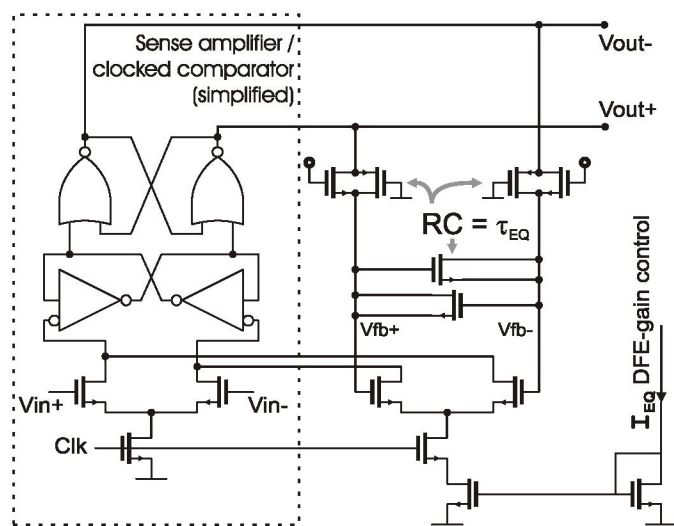


Figure 22.9.3: Implementation of the clocked comparator with continuous-time feedback filter.




Figure 22.9.4: Eye-diagram at the input of the receiver at 1Gb/s and measured Bit Error Rate at the edges of the eye.

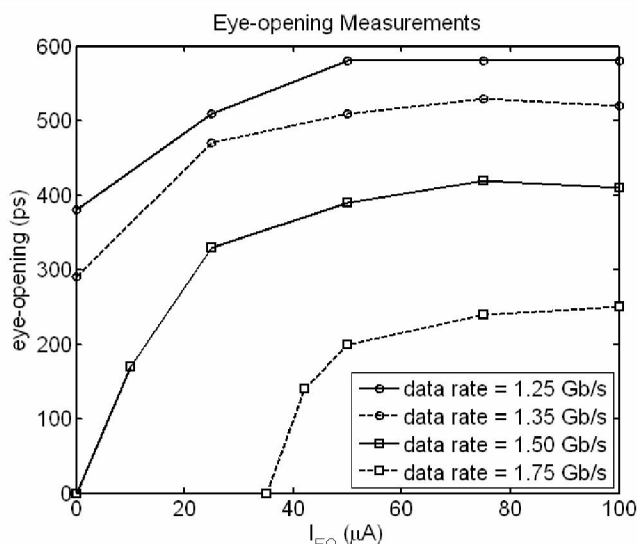


Figure 22.9.5: Measured eye-opening for different data rates as a function of I_{EQ} .

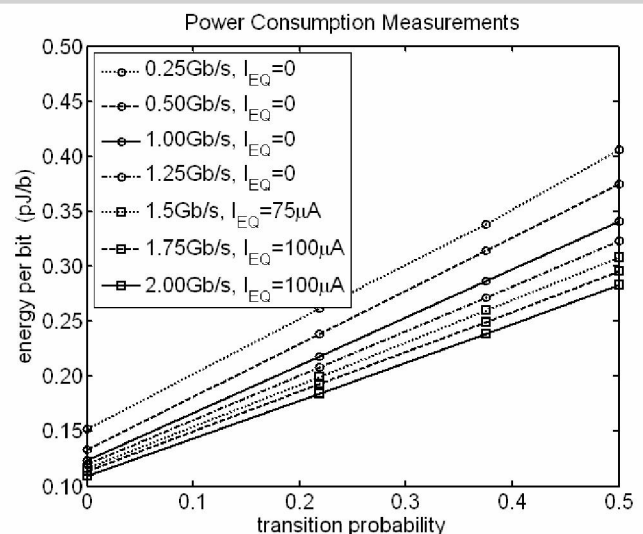


Figure 22.9.6: Measured power consumption for different data rates as a function of transition probability (=data activity).

Continued on Page 612




Figure 22.9.7: Chip micrograph.