

A Power Efficient 2Gb/s Transceiver in 90nm CMOS for 10mm On-Chip Interconnect

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

Abstract—Global on-chip data communication is becoming a concern as the gap between transistor speed and interconnect bandwidth increases with CMOS process scaling. In this paper a low-swing transceiver for 10mm long 0.54 μm wide on-chip interconnect is presented, which achieves a similar data rate as previous designs (a few Gb/s), but at much lower power than recently published work. Both low static power and low dynamic power (low energy per bit) is aimed for. A capacitive pre-emphasis transmitter lowers the voltage swing and increases the bandwidth using a simple inverter based transceiver and capacitive coupling to the interconnect. The receiver uses Decision Feedback Equalization with a power-efficient continuous-time feedback filter. A low power latch-type voltage sense amplifier is used. The transceiver, fabricated in a 1.2V 90nm CMOS process, achieves 2Gb/s. It consumes only 0.28pJ/b, which is 7 times lower than earlier work.

Index Terms—Global on-chip wires, interconnect, on-chip communication, data bus, intersymbol interference (ISI), pre-emphasis, transceivers

I. INTRODUCTION

THE bandwidth of global on-chip interconnects in modern CMOS processes is limited by their high resistance and capacitance [3]. Therefore, the data rate that can be achieved over these long wires is only small. Repeaters can be used to speed up these interconnects, but they consume a considerable amount of power [4] and area. Recently published techniques [3-6] also increase the achievable data rate, but these techniques have high static power consumption, leading to relatively high energy per bit for low data activity. On the other hand, low-swing schemes [7] 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.

In this paper, a transceiver for 10mm long interconnects in

Manuscript received October 1, 2007. This research was supported by the Technology Foundation STW, applied science division of NWO and the technology programme of the Ministry of Economic Affairs.

E. Mensink was with the University of Twente, Enschede, The Netherlands. He is now with Bruco B.V., Borne, The Netherlands (phone: +31-742406650, fax: +31-742406611, email: eisse.mensink@bruco.nl).

D. Schinkel was with the University of Twente, Enschede, The Netherlands. He is now with Axiom IC, Enschede, The Netherlands (email: daniel.schinkel@axiom-ic.com).

E. Klumperink, E. van Tuijl and B. Nauta are with the University of Twente, Enschede, The Netherlands.

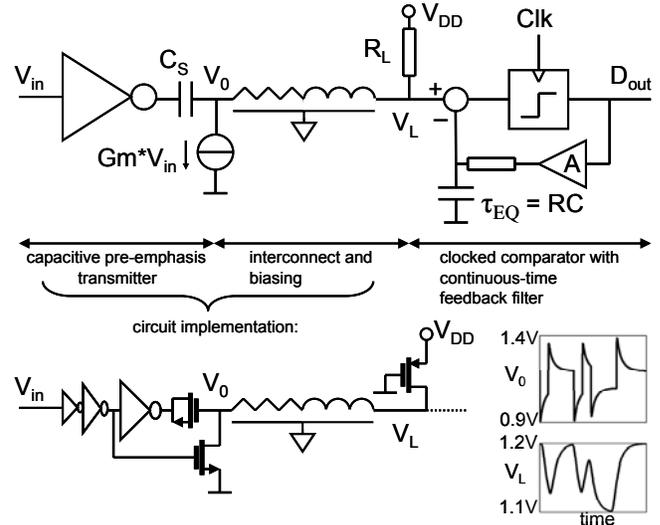


Fig. 1: Concept of the transceiver and circuit implementation of the capacitive pre-emphasis transmitter.

a 1.2V 90nm 6M CMOS process is presented, shown in Fig. 1. A capacitive pre-emphasis transmitter [1] both increases the bandwidth and decreases the voltage swing, without the need for an additional power supply. The receiver uses decision feedback equalization (DFE) [8] to further increase the achievable data rate. The DFE, with a continuous-time feedback filter [1], consumes almost no extra power.

As low-swing signaling is more susceptible to crosstalk, we use differential interconnects with twists [3], of which only a single-ended half is shown. In contrast to the wide interconnects used in [4, 5], we use relatively small widths (0.54 μm) and spacings (0.32 μm) [3, 6] and assume high metal density surroundings.

The paper is organized as follows. We will first describe the techniques that are used to improve the achievable data rate of the interconnect with minimal power consumption. After that we will describe circuit implementations. Measurement results of a test chip are discussed and the results are compared with other transceivers for global interconnects as found in literature.

II. TERMINATION IMPEDANCES AND EQUALIZATION

The bandwidth and power consumption of an RC-limited interconnect depends on its source (Z_S) and load impedances (Z_L). In Fig. 1, a conventional case with inverters as both the transmitter ($Z_S=100\Omega$) and receiver ($Z_L=10\text{fF}$) has only

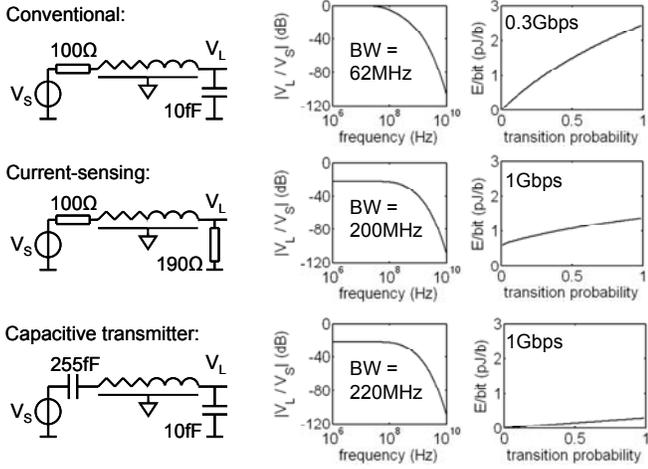


Fig. 2: Bandwidth and energy per bit versus transition probability (=data activity) for three different termination schemes. The results are for 10mm differential interconnects with a distributed resistance of 2k Ω and a distributed capacitance of 2.8pF.

62MHz bandwidth and high power consumption. Current-sensing schemes ($Z_L=190\Omega$ in Fig. 2) increase the bandwidth up to 3 times [3, 6], but with increased power at low data activities. We propose to use a capacitive transmitter ($Z_S=255fF$ in Fig. 2), which has the same bandwidth improvement as current-sensing, but with lower power and without static power consumption. The bandwidth-increasing pre-emphasis effect of the transmitter is shown at the bottom right of Fig. 1: every transition is emphasized by the transmitter by injecting a charge via capacitance C_S .

The receiver concept is also shown in Fig. 1. A clocked comparator [2] is used to restore the low-swing line output to full swing. DFE further increases the achievable data rate. Instead of the often used FIR filters [8], a continuous-time filter is introduced as decision feedback filter. This filter cancels most of the ISI with a simple and power-efficient first-order implementation, whereas a FIR filter would require many taps.

III. CIRCUIT IMPLEMENTATION

With only a series capacitor (AC-coupling), the DC voltage on the interconnect is ill-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. 1). By having the time constants C_S/G_m and $R_L C_{wire}$ equal, the transfer function resembles the transfer function of the capacitive transmitter in Fig. 2. 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. 1. For C_S , the gate capacitance of an NMOS transistor is used. As the gate oxide thickness is much smaller than the oxide between interconnects, the area that is consumed by C_S is relatively small (6x6 μm^2). The signals, with a voltage swing of 100mV, are chosen close to V_{DD} (1.2V), because the capacitance of the NMOS transistor is highest for a high gate-source voltage. The total area of the

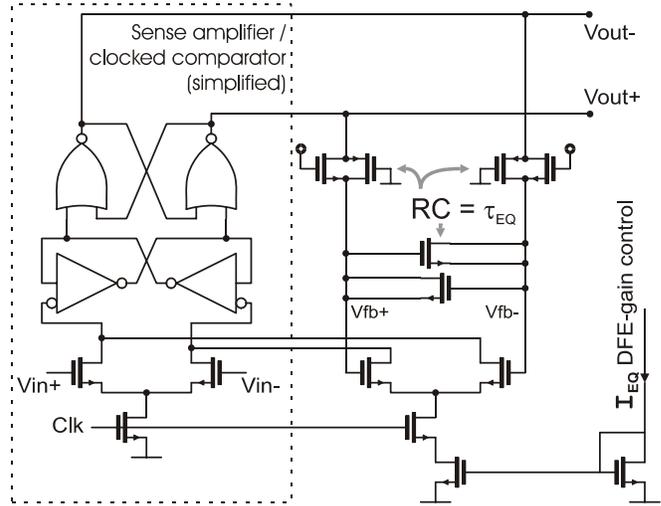


Fig. 3: Implementation of the clocked comparator with continuous-time feedback filter.

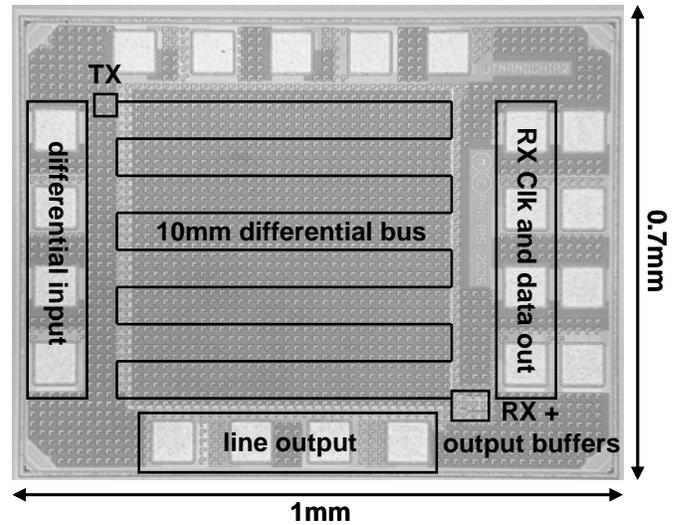


Fig. 4: Chip micrograph.

differential transmitter is 226 μm^2 .

The schematic of the receiver implementation is shown in Fig. 3. The left of the circuit 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 [2]. The outputs of the SR-latch are used to 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. 3. I_{EQ} is used to set the feedback gain A (see Fig. 1). The total area of the receiver is 117 μm^2 (32 μm^2 for the DFE part).

IV. MEASUREMENTS

The chip micrograph is shown in Fig. 4. 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 is used for data

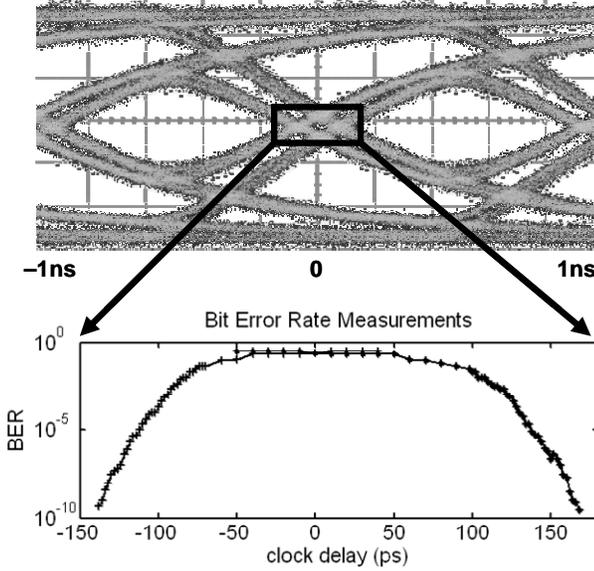


Fig. 5: Eye-diagram at the input of the receiver at 1Gb/s and measured Bit Error Rate at the edges of the eye.

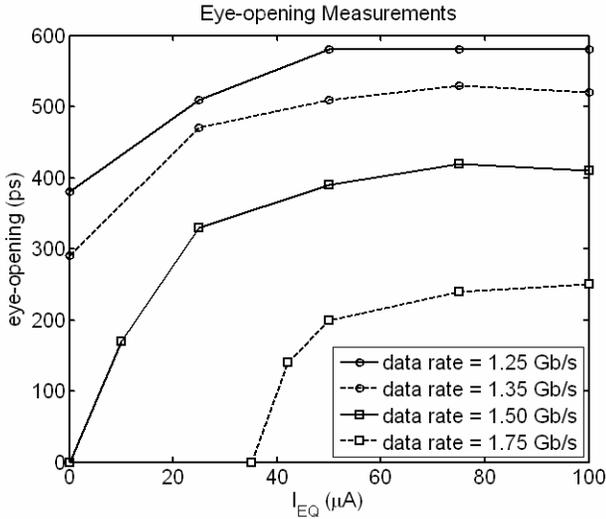


Fig. 6: Measured eye-opening for different data rates as a function of I_{EQ} .

generation and BER measurement. The receiver clock is generated externally in order to adapt its phase to the eye position and 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 50Ω output buffers that are connected to the output of a differential interconnect.

Fig. 5 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 ($I_{EQ}=0$). The one-sigma offset of the total transceiver is 11mV, measured over 20 samples. Due to this offset, not all samples achieve 1.35Gb/s, but a slightly lower data rate of 1Gb/s is achieved by all samples. Simulations over process corners also indicate that the circuit is robust for PVT variations at a rate slightly lower than the maximum achievable data rate. Data rates up to

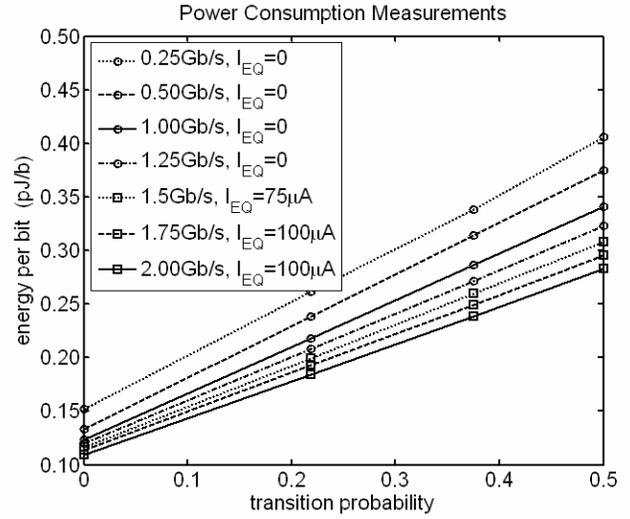


Fig. 7: Measured power consumption for different data rates as a function of transition probability (=data activity).

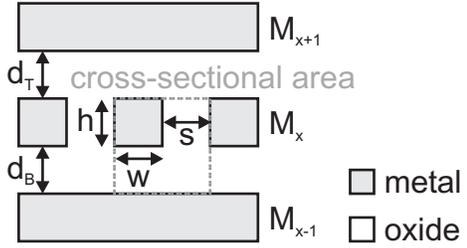


Fig. 8: Definition of cross-sectional area.

2Gb/s are measured with DFE. Fig. 6 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. 7 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 a factor 7 lower than earlier work [3, 6]. 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 ($\sigma_{os}=8mV$). 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 with a factor 1.5.

V. COMPARISON

We will now compare the results of our demonstrator IC with other solutions, as found in literature. We will compare the different interconnect schemes both with respect to achievable data rate and energy consumption. The energy consumption depends linearly on the length (larger length means larger capacitance and hence more energy consumption) and therefore, we will divide the energy consumption by the length. As the bandwidth of RC-limited interconnects depends on the length squared (larger length means smaller bandwidth), we will divide the achievable data rate by the length squared. As we would also like to consume as little chip area as possible, we will also divide the

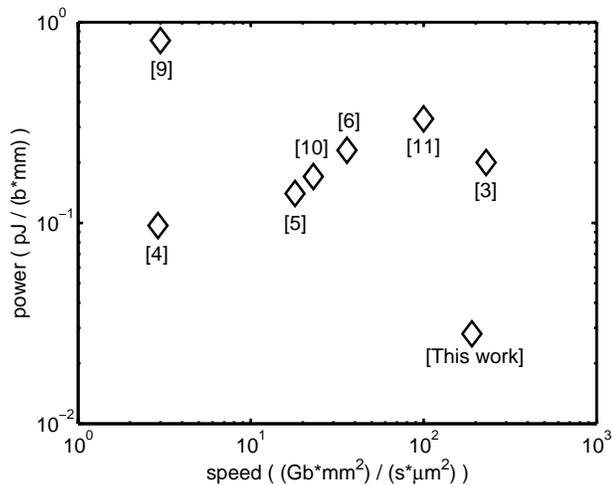


Fig. 9: Comparison of different solutions with respect to speed and power.

achievable data rate by the cross-sectional area (see Fig. 8) of the interconnect. The cross-sectional area is defined as $(w+s)(h+d)$ with w the width of the interconnect, s the spacing, h the height of the interconnect and $d (=d_T=d_B)$ the vertical spacing to other metal layers. The parameters s , h and d are not always given in literature and are in some cases estimated from the used technology process.

Fig. 9 has on the x-axis the achievable data rate divided by the cross-sectional area and the length squared and on the y-axis the energy consumption per transmitted bit divided by the length. The figure shows that the transceiver as presented in this paper has both a high achievable data rate and much lower energy consumption than all other solutions.

ACKNOWLEDGEMENT:

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

REFERENCES:

- [1] E. Mensink, D. Schinkel, E.A.M. Klumperink, A.J.M. van Tuijl, B. Nauta, "A 0.28pJ/b 2Gb/s/ch Transceiver in 90nm CMOS for 10mm On-Chip Interconnects," IEEE Int. Solid-State Circuits Conference (ISSCC) Dig. Tech. Papers, pp. 414-415, Feb. 2007.
- [2] D. Schinkel, E. Mensink, E.A.M. Klumperink, A.J.M. van Tuijl, B. Nauta, "Double-Tail Latch-Type Voltage Sense Amplifier With 18ps Setup+Hold Time," IEEE Int. Solid-State Circuits Conference (ISSCC) Dig. Tech. Papers, pp. 314-315, Feb. 2007.
- [3] D. Schinkel, 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, pp. 297-306, Jan. 2006.
- [4] A. P. Jose, G. Patounakis, K. L. Shepard, "Pulsed Current-Mode Signaling for Nearly Speed-of-Light Intrachip Communication," IEEE J. Solid-State Circuits, vol. 41, pp. 772-780, April 2006.
- [5] A. P. Jose, K. L. Shepard, "Distributed Loss Compensation for Low-Latency On-Chip Interconnects," ISSCC Dig. Tech. Papers, pp. 516-517, Feb. 2006.
- [6] L. Zhang, et al., "Driver Pre-Emphasis Techniques for On-Chip Global Buses," Proc. of the ISLPED, pp. 186-191, Aug. 2005.
- [7] H. Zhang, V. George, J. M. Rabaey, "Low-Swing On-Chip Signaling Techniques: Effectiveness and Robustness," IEEE Trans. on VLSI Systems, vol. 8, pp. 264-272, June 2000.
- [8] V. Stojanovic, et al., "Adaptive Equalization and Data Recovery in a Dual-Mode (PAM2/4) Serial Link Transceiver," Symp. on VLSI Circuits Dig. Tech. Papers, pp. 348-351, June 2004.
- [9] R. T. Chang, C. P. Yue, and S. S. Wong, "Near speed-of-light on-chip electrical interconnect," VLSI Circuits, Digest of Tech. Papers, Symp. on, pp. 18-21, June 2002.
- [10] A. Katoch, H. Veendrick, and E. Seevinck, "High speed current-mode signaling circuits for on-chip interconnects," Circuits and Systems (ISCAS), proc. of the IEEE Intern. Symp. on, pp. 4138-4141, May 2005.
- [11] R. Bashirullah, L. Wentai, R. Cavin, III, and D. Edwards, "A 16 Gb/s adaptive bandwidth on-chip bus based on hybrid current/voltage mode signaling," Solid-State Circuits, IEEE Journal of, vol. 41, pp. 461-473, Feb. 2006.