20.7 A 3Gb/s/ch Transceiver for RC-limited On-Chip Interconnects

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

University of Twente, Enschede, The Netherlands

The on-chip communication is getting more attention, as (global) interconnects are rapidly becoming a speed, power and reliability bottleneck for digital systems [1]. Technological advances such as copper interconnects and low-k dielectrics are not sufficient to let the interconnect bandwidth keep up with the advances in transistor speeds.

From a circuit-design perspective, a general solution is the use of repeaters, but at the expense of area and power. Another proposed solution [2] uses low-swing signaling over differential 10mm aluminum interconnects, but with the requirement of clocked switches along the wire, increasing the already troublesome clock-load. In [3], it is proposed to use 16μm-wide differential wires (20mm long) and exploit the LC regime (transmission-line behavior) of these wires, but at the expense of a significant increase in power consumption and interconnect area. Both papers achieve 1Gb/s/ch in a 0.18μm CMOS technology.

In this paper, a bus transceiver demonstrator IC in a 1.2V 0.13μm 6M copper CMOS process is presented. Simulations and measurements show that pulse-width pre-emphasis in combination with resistive termination can increase the data-rate to 3Gb/s/ch, using 10mm-long, 0.4μm-wide differential interconnects. Without the proposed techniques, these interconnects can only achieve 0.55Gb/s/ch.

The interconnects, as shown in Fig. 20.7.1, are modeled with a 3D EM-field solver and a distributed RLC model is extracted (0.15kΩ/mm, 0.35nH/mm and 0.27pF/mm). The bus is placed in metal 5 as it is assumed that the thick top-metal is reserved for clock and power routing. In the EM-field solver, metal 4 and metal 6 plates approximate the effect of other high-density interconnects. The dimensions of the interconnects are optimized for highest bandwidth per cross-sectional area. Analysis and simulations show that the bandwidth per cross-area peaks when all dimensions (w, s, h, tt and tb) are equal. This results in both a width and a spacing of 0.4μm. Figure 20.7.1 also shows the simulated interconnect transfer function. For these long and narrow interconnects the effect of inductance is negligible. A significant part of the transfer function can be approximated by a first-order RC model. Note that the bandwidth increases 3 times with low-ohmic resistive termination instead of (conventional) capacitive termination [4].

To reduce the overall crosstalk differential interconnects are used [5]. Furthermore, to cancel neighbor-to-neighbor crosstalk between channels in a bus, one twist is placed at 50% of the length in the even channels and two twists are placed at 25% and 75% in the uneven channels, as shown in Fig. 20.7.2.

The dominance of the first-order roll-off makes an on-chip interconnect very suitable for simple pre-emphasis transmission schemes (e.g. 2-taps FIR). Conventional pre-emphasis schemes (overdrive signaling) used for inter-chip communication [6] are less suitable for on-chip implementation. The large drive impedance of simple current-summing transmitters degrades interconnect bandwidth, while low-ohmic voltage transmitters require additional low-impedance voltage levels and their performance is degraded by slew-rate. As a robust alternative, the use of pulse-width (PW) pre-emphasis is proposed. As shown in Fig. 20.7.3, PW pre-emphasis can greatly reduce the amount of ISI, by using the second part of the symbol-time to compensate for the remaining line charge.

The schematic of the PW pre-emphasis transmitter is shown in Fig. 20.7.4. The PW-modulated signal is generated with a clock with adjustable duty-cycle that selects either Data or not(Data). The not(Data) is delayed by half a clock-cycle to increase the timing margin. In the prototype IC, the duty-cycle is controlled by an external current source, to provide programmability. $I_{bias}=0$ results in conventional binary signaling. At a 3GHz clock an I_{bias} of 80μA, 200μA or 400μA results in transmitted symbols with pulse-widths of 75%, 58% or 52%, respectively. In a product application, the I_{bias} can be fixed at design time, as the transmission scheme is robust towards (circuit and wire) parameter deviations. The line-driver inverters are scaled to have an R_{out} of about 60Ω and the size of the differential TX is ~300μm^2.

The schematic of the receiver is shown in Fig. 20.7.2. The input inverters use transmission gates as selectable feedback resistors. In this way, either conventional termination or (active) resistive (R_{in} ≈150Ω) termination can be selected. A clocked comparator followed by a dynamic latch samples the received data. The size of the prototype differential RX is ~1000μm^2 (non-optimized). The complete design is optimized for low mismatch and to function over all process corners. Dynamic latches, low-V_t transistors and small fan-outs (≤3) are used to meet the target data-rate of 3Gb/s even at the slow process corner. The simulated latency of the transceiver is about 650ps at 3Gb/s, composed of 180ps for the TX, 420ps for the channel and 50ps for the RX.

Figure 20.7.7 shows the test chip micrograph, with a 7 channel differential bus, surrounded by GND/Vdd-connected metal stripes. A single-ended bus is placed below the differential bus, providing some intra-bus crosstalk. An external single-channel 3.2Gb/s pattern generator/analyzer is used for the data generation and BER measurement. Large on-chip delay lines (chains of flip-flops) provide all bus-channels with pseudo-independent data. The phase of the RxClk can be adjusted externally to adapt to the eye position and measure its width. The measured line parameters are 0.19kΩ/mm and 0.25pF/mm which agree with simulations given the tolerance bounds of the process.

Figure 20.7.5 shows the measured eye-diagrams at the input of the clocked comparator, both with and without the use of PW pre-emphasis and resistive termination, with data-rates at the edge of immeasurable BER (<1e-12). Note that the achievable data-rate increases 4 times by PW pre-emphasis, 3 times by resistive termination and 6 times by the combination of both.

At 3.2Gb/s, the eye-opening at the RX side is so small that offset and memory effects in the clocked comparator lead to measurable BER (5e-9). At 3Gb/s, error-free operation is possible for all 10 measured samples (with nominal biasing). At 2.5Gb/s, the design is very robust and the BER remains immeasurable with large external parameter deviations: $1.0V<V_{dd}<1.5V$ (nominal 1.2V); 34%<TxClk duty-cycle<62% (nominal 50%); $130μA<I_{bias}<400μA$ (nominal 200μA); –130ps<RxClk skew<+130ps.

Figure 20.7.6 illustrates crosstalk from a non-twisted neighboring interconnect on both single-ended halves of an interconnect with one twist. The reduction in crosstalk on the differential voltage (due to the twist) is apparent.

At 3Gb/s, the total power consumption (TX+RX) for a single channel is 6mW. Conventional repeater systems consume up to 4 times more power [2] and have comparable latency. The TX and RX circuits are well suited for power-management, as the speed-enhancing, but power-consuming techniques can be easily turned on and off dynamically.

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

0-7803-8904-2/05/$20.00 ©2005 IEEE.

References:
[1] R. Ho et al., “The Future of Wires,” *Proc. IEEE*, pp. 490-504, Apr., 2001.
[2] R. Ho et al., “Efficient On-Chip Global Interconnects,” *Symp. VLSI Circuits*, pp. 271-274, June, 2003.
[3] R. Chang, et. al., “Near Speed-of-Light Signaling Over On-Chip Electrical Interconnects,” *IEEE J. Solid-State Circuits*, vol. 38, pp. 834-838, May, 2003.
[4] E. Seevinck et al., “Current-Mode Techniques for High-Speed VLSI Circuits with Application to Current Sense Amplifier for CMOS SRAM’s,” *IEEE J. Solid-State Circuits*, vol. 26, pp. 525-536, Apr., 1991.
[5] Y. Massoud, et. al., “Modeling and Analysis of Differential Signaling for Minimizing Inductive Cross-Talk,” *Proc. DAC*, pp. 804-809, June, 2001.
[6] R. Farjad-Rad, et. al., “A 0.4µm CMOS 10-Gb/s 4-PAM Pre-Emphasis Serial Link Transmitter,” *IEEE J. Solid-State Circuits*, vol. 34, pp. 580 -585, May, 1999.

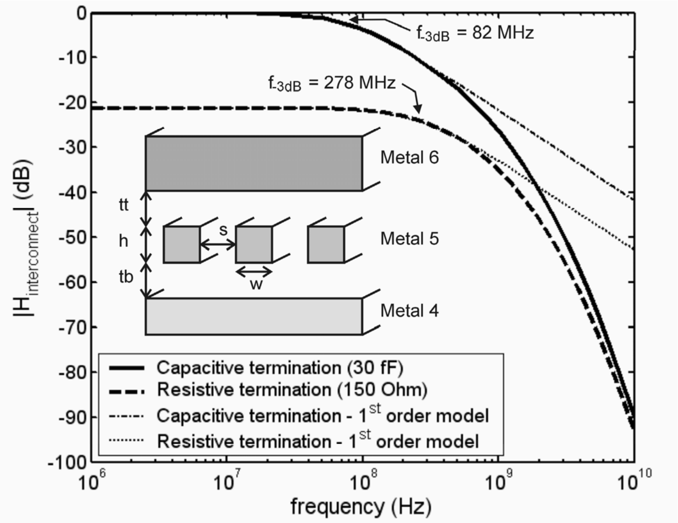

Figure 20.7.1: Interconnect transfer function with resistive and capacitive termination.

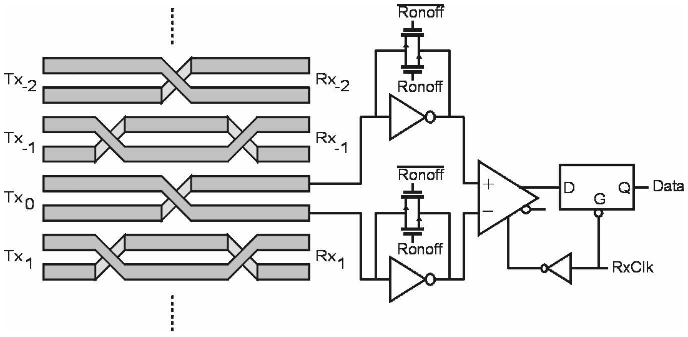

Figure 20.7.2: Differential bus and receiver schematic.

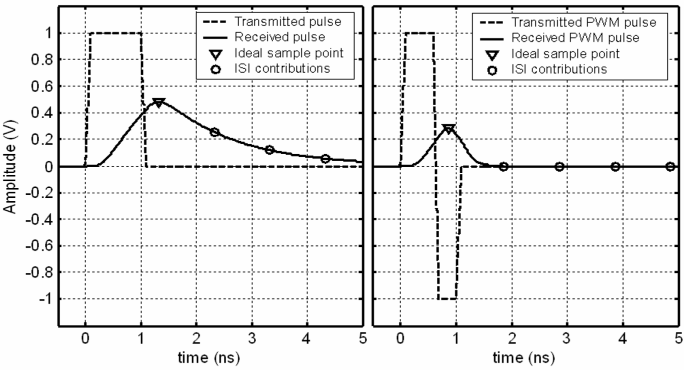

Figure 20.7.3: Symbol responses of (capacitively terminated) 1-cm interconnect with 1ns symbol period.

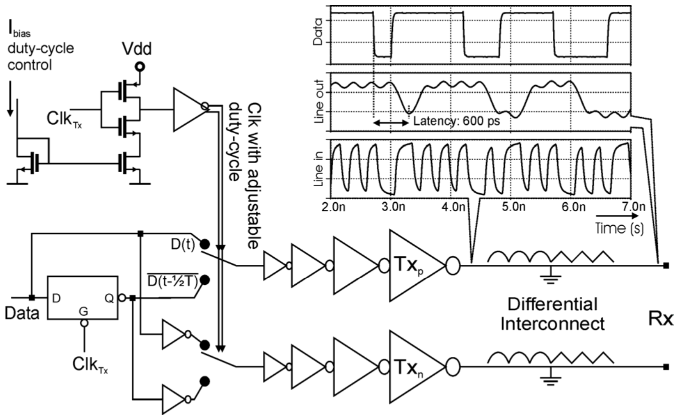

Figure 20.7.4: Transmitter schematic and signal waveforms.

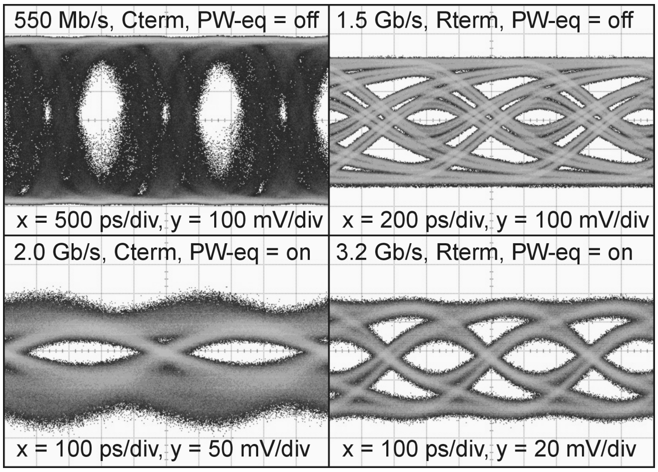

Figure 20.7.5: Eye-diagrams for various configurations. The output buffers compress the vertical scale; on-chip signals are 6 to 9dB larger.

Continued on Page 606

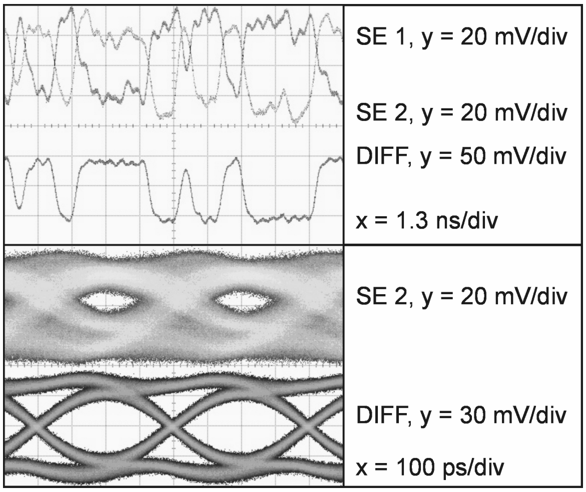

Figure 20.7.6: Effect of crosstalk on single-ended (SE) and differential (twisted) interconnect @2.5Gb/s. The output buffers compress the vertical scale; on-chip signals are 6 to 9dB larger.

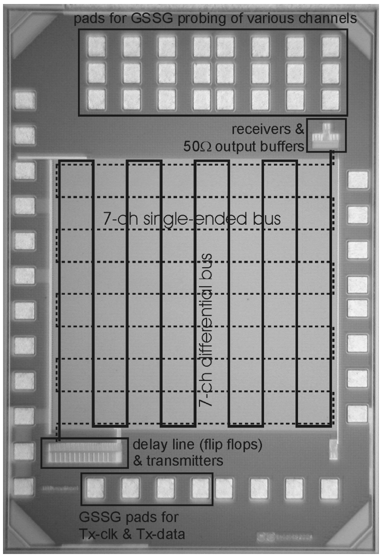

Figure 20.7.7: Chip micrograph.