Power Efficient Gigabit Communication Over Capacitively Driven RC-Limited On-Chip Interconnects

Eisse Mensink, *Member, IEEE*, Daniël Schinkel, *Member, IEEE*, Eric A. M. Klumperink, *Senior Member, IEEE*, Ed van Tuijl, *Member, IEEE*, and Bram Nauta, *Fellow, IEEE*

Abstract—This paper presents a set of circuit techniques to achieve high data rate point-to-point communication over long on-chip RC-limited wire-pairs. The ideal line termination impedances for a flat transfer function with linear phase (pure delay) are derived, using an s-parameter wire-pair model. It is shown that a driver with series capacitance on the one hand and a resistive load on the other, are fair approximations of these ideal terminations in the frequency range of interest. From a perspective of power efficiency, a capacitive driver is preferred, as the series capacitance reduces the voltage swing along the line which reduces dynamic power consumption. To reduce cross-talk and maintain data integrity, parallel differential interconnects with alternatingly one or two twists are used. In combination with a low offset dynamic sense amplifier at the receiver, and a low-power decision feedback equalization technique with analog feedback, gigabit communication is demonstrated at very low power consumption. A point-to-point link on a 90 nm CMOS test chip achieves 2 Gb/s over 10 mm long interconnects, while consuming 0.28 pJ/bit corresponding to 28 fJ/bit/mm, which is much lower than competing designs.

Index Terms—Capacitive coupling, CMOS, communication techniques, decision feedback equalization, de-emphasis, equalization, low power electronics, low-swing, networks on chip, NoC, on-chip interconnects, on-chip wires, pre-emphasis, RC-limited interconnects.

I. INTRODUCTION

GLOBAL on-chip interconnects, spanning a large portion of CMOS system chips, are a well-known speed, power and reliability bottleneck for digital CMOS systems [1]. Such interconnects are for example used for on-chip buses to connect the different parts of a microprocessor or a System on a Chip (SoC) [2] and in memories, as global address or data-lines. Even if a Network on Chip (NoC) [3] architecture is used to relieve the problems, the availability of fast global interconnects will be desirable. For example, a NoC can benefit from circular network topologies, such as torus or folded torus configurations [4], which require longer interconnects than the standard mesh topology.

Interconnects have an RC-limited bandwidth roughly proportional to the area of the metal cross section and inversely proportional to the squared length [5]. With increasing speeds and reduced metal dimensions, wires are becoming more and more problematic. From a circuit design perspective, a general solution to the limited interconnect bandwidth is the use of repeaters, which make the repeated wire delay linear with length instead of quadratic [5], increasing the achievable data rate [6]. However, for optimal repeater insertion [5] many repeaters are needed, which costs area and power, and makes floor-planning more difficult as portions of active area all over the chip have to be reserved for large repeater circuits. Moreover, the long chain of inverters also creates a high static variation in delay for different process, voltage and temperature (PVT) corners, limiting the achievable data rate again [7].

In recent years, alternatives to repeater insertion were proposed in order to decrease the delay of global on-chip interconnects. Many of these techniques also increase interconnect bandwidth and thus the achievable data rate. Current-mode sensing [7]–[11], for instance, decreases the wire delay by a factor of three and increases the bandwidth by the same factor. The equalization techniques in [7], [10]–[12], referred to as dynamic overdriving or pre-emphasis equalization, also both decrease wire delay and increase the achievable data rate. However, the downside to most of these techniques is a significant increase in power consumption (both static and dynamic power).

During ISSCC 2007, both Ho *et al.* [13] and our research group [14] independently proposed to use a data transmitter with capacitive coupling to the interconnect, resulting in increased bandwidth and low power consumption due to a reduced voltage swing. In [20] Ho *et al.* presented more details on their work and provided a qualitative intuitive explanation of the capacitive driver technique. This paper extends our work in [14]. With the use of an s-parameter model, we will analyze the transfer function of RC-limited interconnect showing that an *ideal source or load impedance* exists for which the *transfer function becomes flat*, as is desired for inter symbol interference and bandwidth. In the frequency range of interest, a capacitance will appear to be a fair approximation of the ideal source impedance, while also reducing the dynamic power consumption (reduced voltage

swing). In order to cope with the lower signal swing, it is important to mitigate cross-talk via twisting [16] and to design a receiver with sufficient sensitivity and speed [15]. We will discuss circuit techniques that achieve this goal at record low power consumption, especially a low offset dynamic sense amplifier and a low-power decision feedback equalizer exploiting analog feedback. Compared to previous publications, this paper also provides more information on the operation of the circuits and the associated power consumption. Moreover, we compare our system with other solutions to demonstrate that we can achieve a higher data rate, while also reducing the energy per bit per mm line length. Recently, we also proposed to use the capacitive driver for Networks on a Chip [19] with shorter line lengths, where bandwidth limitations are less pressing. The main point of [19] is to reduce the power consumption in the lines by reducing voltage swing without requiring an additional supply.

The paper is organized as follows. First, in Section II, an s-parameter model for bandwidth analysis of RC-limited interconnects is derived. Then, Section III addresses methods to increase the achievable data rate. We do this by looking at the termination impedances of the interconnect to derive ideal source and load impedances for a flat transfer. We will also look at an equalization scheme that can further increase the achievable data rate. In Section IV we consider the power consumption of the techniques that are described in Section III and compare them with each other. Section V describes the circuit implementation of the most promising techniques from Sections III and IV and Section VI gives measurement results of a test chip. Finally, in Section VII, the implemented techniques are compared with solutions as found in literature with respect to both speed and power consumption.

II. INTERCONNECT MODEL

The interconnect is RC-limited and the inductance can be neglected as long as the *RC* time constant $(\text{length}^2 \cdot R' \cdot C')$ is much larger than the L/R time constant (L'/R'), which is true in our case for lines longer than about 0.7 mm (R', L' and C' being the wire resistance, inductance and capacitance per unit length) [17]. We will now derive a model for the RC-limited interconnects with which we are able to calculate both the achievable data rate and the power consumption. The communication structure is assumed to consist of a point-to-point bus with all signals traveling in the same direction. Assuming the thick top layers are reserved for clock and power routing, we place the bus in one of the lower metal layers indicated with x (see Fig. 1). The metal plates in metal $x + 1$ and metal $x - 1$ model high density perpendicular interconnects (Manhattan routing style). We further assume high density metal use in all metal layers, thus the interconnect has capacitances to all sides. The width and spacing of the interconnects are chosen to *maximize* the *bandwidth per cross-sectional area* (see Fig. 1), as derived in [7] and [17]. This is achieved by choosing the width (w) and spacing (s) about equal to the height (h) and vertical spacing (d) of the interconnects (see Fig. 1).

The model that we will use in the analysis is given in Fig. 2. The transmitter is modeled by a voltage source with source impedance Z_S and the receiver is modeled by a load impedance Z_L. In later sections, we will see how these two termination
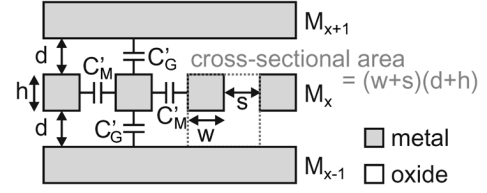

Fig. 1. Interconnect dimensions and definition of the cross-sectional area. The metal plates in M_{x-1} and M_{x+1} model perpendicular interconnects. The cross-sectional area $A_c = (w + s)(h + d)$.
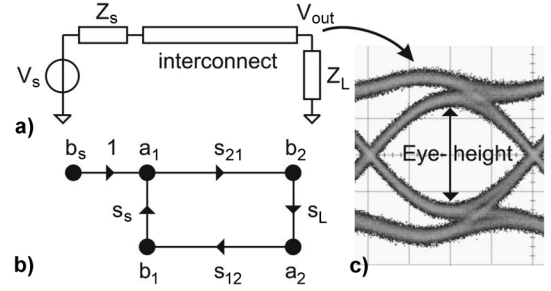

Fig. 2. (a) Interconnect with source impedance Z_S and load impedance Z_L; (b) corresponding s-parameter model; (c) a possible eye-diagram at V_{out} with the definition of eye-height.

impedances influence the achievable data rate (Section III) and the power consumption (Section IV). The interconnect and the termination impedances are modeled with s-parameters, as also shown in Fig. 2. The s-parameters are

$$s_S = \frac{Z_S - Z_C}{Z_S + Z_C} \tag{1}$$

$$s_L = \frac{Z_L - Z_C}{Z_L + Z_C} \tag{2}$$

$$s_{21} = s_{12} = e^{-\gamma \cdot l} \tag{3}$$

with l the length of the interconnect, Z_C its characteristic impedance, and γ the propagation constant:

$$Z_C = \sqrt{\frac{R'}{j \cdot \omega \cdot C'}} \tag{4}$$

$$\gamma = \sqrt{j \cdot \omega \cdot R' \cdot C'}. \tag{5}$$

With this s-parameter model, the transfer function from V_S to V_{out} can readily be calculated [17]:

$$\frac{V_{out}}{V_S} = \frac{Z_C}{Z_S + Z_C} \cdot \frac{s_{21} \cdot (1 + s_L)}{1 - s_S \cdot s_{21} \cdot s_L \cdot s_{12}}. \tag{6}$$

From the transfer function, the impulse response can be calculated (inverse FFT) and from the impulse response the symbol response of the interconnect (convolution of impulse response and symbol). From the symbol response in turn, the worst-case eye-height of V_{out} [see Fig. 2(c)] can be determined with the same method as used in [7]. For higher data rates, the worst-case eye-height will become smaller and eventually become zero. If the minimum eye-height that is needed at the receiver side to reliably detect the transmitted symbols is known, the achievable data rate can be determined, as done in the next section.
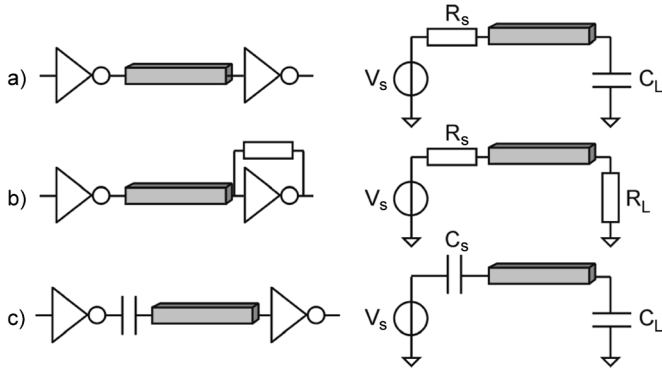
Fig. 3. Three line-termination schemes and their equivalent circuit representations. (a) Conventional scheme; (b) current-sensing scheme; (c) capacitive transmitter scheme. In simulations we use $R_S = 100$ ohm, $C_L = 10$ fF, $R_L = 233$ ohm and $C_S = 311$ fF, while the line is characterized by $R' = 0.20$ kΩ/mm and $C' = 0.28$ pF/mm.
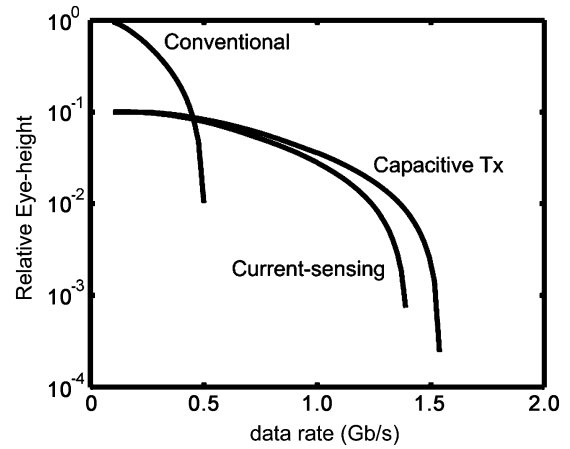

Fig. 4. Calculated worst-case relative eye-height as a function of the data rate for the three different termination schemes in Fig. 3 for 10 mm line length and $R' = 0.20$ kΩ/mm and $C' = 0.28$ pF/mm.

Although the calculation procedure is reasonably straightforward, deriving analytical expressions is intractable. Therefore, we resort to numerical simulations to evaluate the achievable data rate. We will use data from a 90 nm CMOS process with 7 metal layers which is also used for our test chip (see Section VI). We used metal 4 as M_x in Fig. 1, with a width of 0.54 μm and a spacing of 0.32 μm between neighboring interconnects, and derived $R' = 0.20$ kΩ/mm and $C' = 0.28$ pF/mm from measurements [17]. The length of interconnects is chosen to be 10 mm, which represents a typical global interconnect and allows for easy comparison with prior work.

III. ACHIEVABLE DATA RATE

A. Termination Schemes

In this section, we will look at the achievable data rate as a function of the source and load impedance. First, we will look at the case where the interconnect is driven by an inverter and also the receiver consists of an inverter. Supposing the drive-inverter is large enough, we can model the transmitter with a small resistive source impedance and the receiver with a capacitor, as shown in Fig. 3(a). We will call this the conventional termination scheme. The worst-case eye-height of this scheme, relative to the input swing, has been calculated as discussed in the previous section and is shown in Fig. 4. If we assume that about 10 mV eye-opening is needed for detection and 1 V corresponds to 100%, we can roughly use data rates up to an eye opening of 1E-2 (the exact assumption is not critical as the eye-opening curves fall off steeply). The achievable data rate is about 0.5 Gb/s for the conventional termination scheme. In this case the -3 dB bandwidth is about 62 MHz, and we see the eye opening gradually drops above that bandwidth.

If current-sensing is used [7]–[11], the load impedance Z_L is not capacitive anymore, but a small resistor [Fig. 3(b)]. The resulting worst-case eye-height curve in Fig. 4 shows almost three times increase in achievable data rate. Of course, due to the low load impedance, the maximum value of the eye-height at low data rates is smaller than with the conventional scheme (resistive division).

Another way of increasing the achievable data rate is to use the capacitive transmitter of Fig. 3(c) [13], [14]. Now, the achievable data rate has increased about three times, slightly more than for current-sensing, again at the cost of a reduced maximum voltage swing.

What we can learn from Fig. 4 is that by choosing a suitable load impedance or a suitable source impedance, the achievable data rate can be increased significantly. We will use this result in the next paragraph.

B. Ideal Termination Schemes

We have seen that by choosing either a resistive load impedance or a capacitive source impedance, the achievable data rate can be increased. The question can be asked: what are the theoretically ideal termination impedances? As large bandwidth without inter-symbol interference is desired, we aim for a flat transfer function with linear phase (no dispersion, only delay t_d), i.e.,

$$H_{ideal} = A \cdot e^{-j\omega t_d}. \qquad (7)$$

For a fixed load impedance, we can find the ideal source impedance $Z_{S,ideal}$, for which the transfer function of the interconnect is equal to H_{ideal}. If we assume that the receiver is a small capacitive load, thus $Z_L = 1/j\omega C_L$, we find the $Z_{S,ideal}$ of Fig. 5, assuming $C_L = 10$ fF, $t_d = 1$ ns, and $A = 0.1$. For $A = 1$ (not shown), the ideal source impedance is a negative resistance over a large frequency range, which intuitively makes sense as it should compensate the resistive losses in the line. Interestingly, for $A = 0.1$, the ideal source impedance resembles a capacitor for the lower frequencies, which explains the results of the previous paragraph. As the ideal source impedance is not equal to a capacitor anymore for frequencies roughly above 200 MHz, the transfer does not remain flat, which explains degrading eye openings in Fig. 4 above 200 MHz.

Of course, we can also fix the source impedance, for instance use a small resistor R_S, and calculate the ideal load impedance. We now find, with $t_d = 1$ ns, $R_S = 100$ Ω and $A = 1$, that the ideal load impedance resembles a negative capacitance. If we choose $A = 0.1$, the ideal load impedance for the lower frequencies resembles a resistance of about 233 Ω (see Fig. 6). This again holds roughly up to 200 MHz. Again, this is in agreement
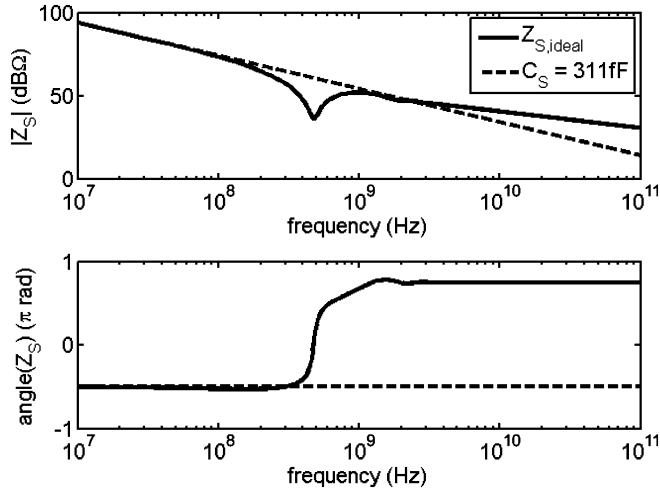
Fig. 5. Calculated absolute value and phase angle of the ideal source impedance that renders a flat transfer function with linear phase for $t_d = 1$ ns, $C_L = 10$ fF and $A = 0.1$. The ideal source impedance resembles a capacitor $C_S = 311$ fF (also shown) at low frequencies.
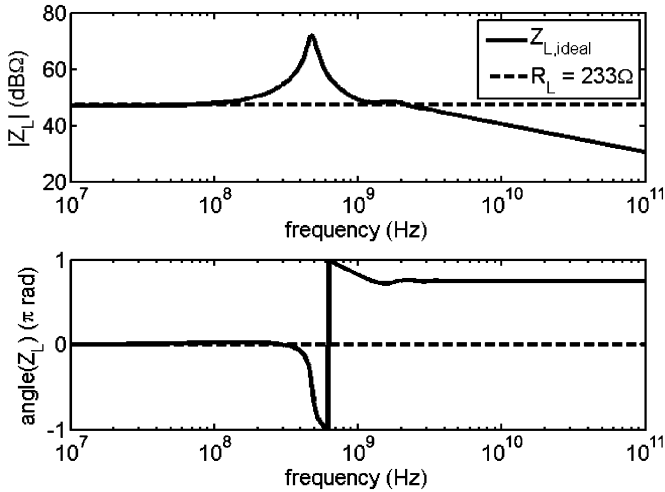

Fig. 6. Calculated absolute value and phase angle of the ideal load impedance that renders a flat transfer function with linear phase for $t_d = 1$ ns, $R_S = 100$ Ω for $A = 0.1$. The ideal load impedance resembles a resistor of 233 Ω (also shown) at low frequencies.

with the previous paragraph, where a small current-sensing resistance as load impedance increases the achievable data rate by about a factor of three compared to the conventional termination scheme.

If we change the nominal delay t_d we still see a capacitor-alike optimum Z_S and resistor-alike optimum Z_L, but the difference between amplitude and phase of the practical case (capacitor/resistor) compared to the theoretical optimum varies, where amplitude deviations and phase deviations can be traded to some extend. We chose $t_d = 1$ ns as it is close the actual delay found for Fig. 3(b) and 3(c).

Whereas the theoretical improvements for a capacitive driver and current sensing receiver are similar (a factor 3), the practical implementation problems are quite different, potentially impairing the achievable improvement. To realize a current-sensing amplifier with low-ohmic input impedance $1/g_m$, large g_m and hence significant static bias current is needed

(e.g., 1.5 mW to realize $1/g_m = 150$ Ω [7]). A capacitive driver does not need such high g_m, but has other implementation challenges, e.g., to realize a well defined DC-bias after capacitive AC-coupling and equalize the DC and AC-path. In Sections IV and V we will address these issues and show that these problems can be solved in a robust and power efficient way.

C. Equalization

A capacitive transmitter increases the bandwidth of the interconnect, but we can even improve the performance further using equalization at the receiver. In [14] we proposed to use decision feedback equalization (DFE), and will show here that this can be implemented at very low power penalty. DFE is well known in other communication areas, for example in inter-chip communication [18]. Fig. 7 shows two alternative implementations of DFE. At the receiver a decision is made by a comparator or sense amplifier whether the received symbol is a 'one' or a 'zero'. The result of this decision is fed back to the end of the interconnect via a filter. This filter can be either discrete-time [Fig. 7(a)] or continuous-time [Fig. 7(b)] and is used to remove the long tail from the symbol response, as also shown in Fig. 7(c) and 7(d). A simple analog RC-feedback filter is used here because it fits almost perfectly to the dominant RC low-pass response of the line [see the solid line in Fig. 7(b)]. A discrete-time version of this filter would require many taps and cost more power consumption.

The worst-case eye-height for a scheme with a simple capacitive transmitter and DFE at the receiver with continuous-time feedback is shown in Fig. 8 along with the situation where only a capacitive transmitter is used (without DFE). The figure shows that the DFE equalization makes higher data rates possible, provided that the receiver can cope with small relative eye-height (e.g., 50% increase in data-rate for a relative eye-height of 0.05). In the next section we will show that this can be done at very small power penalty.

IV. ENERGY CONSUMPTION

The previous section showed that by choosing suitable termination impedances, the achievable data rate can be increased. In this section, we will compare these solutions in terms of power consumption. We aim at minimum energy consumption and hope to spend energy only when useful information is transferred. Therefore, we will look at the energy per bit. We will plot this energy per bit as a function of data activity or transition probability (p_{trans}). Ideally, we would like the energy per bit to be linearly dependent on p_{trans} with zero energy consumption if $p_{trans} = 0$ (no static power consumption). Fig. 9 shows the energy consumption for the three different schemes shown in Fig. 3, assuming a binary (zero-mean) Markov source as derived in [17, pg. 37–38].

Note that the DFE equalization is not included in the figure. The reason for this is that it can be used in combination with all other termination schemes and adds only very little extra power (about 0.02 pJ/bit, see Section VI). Fig. 9 shows that only the resistive termination scheme (current-sensing) has a large static energy consumption, as it requires a static current to maintain
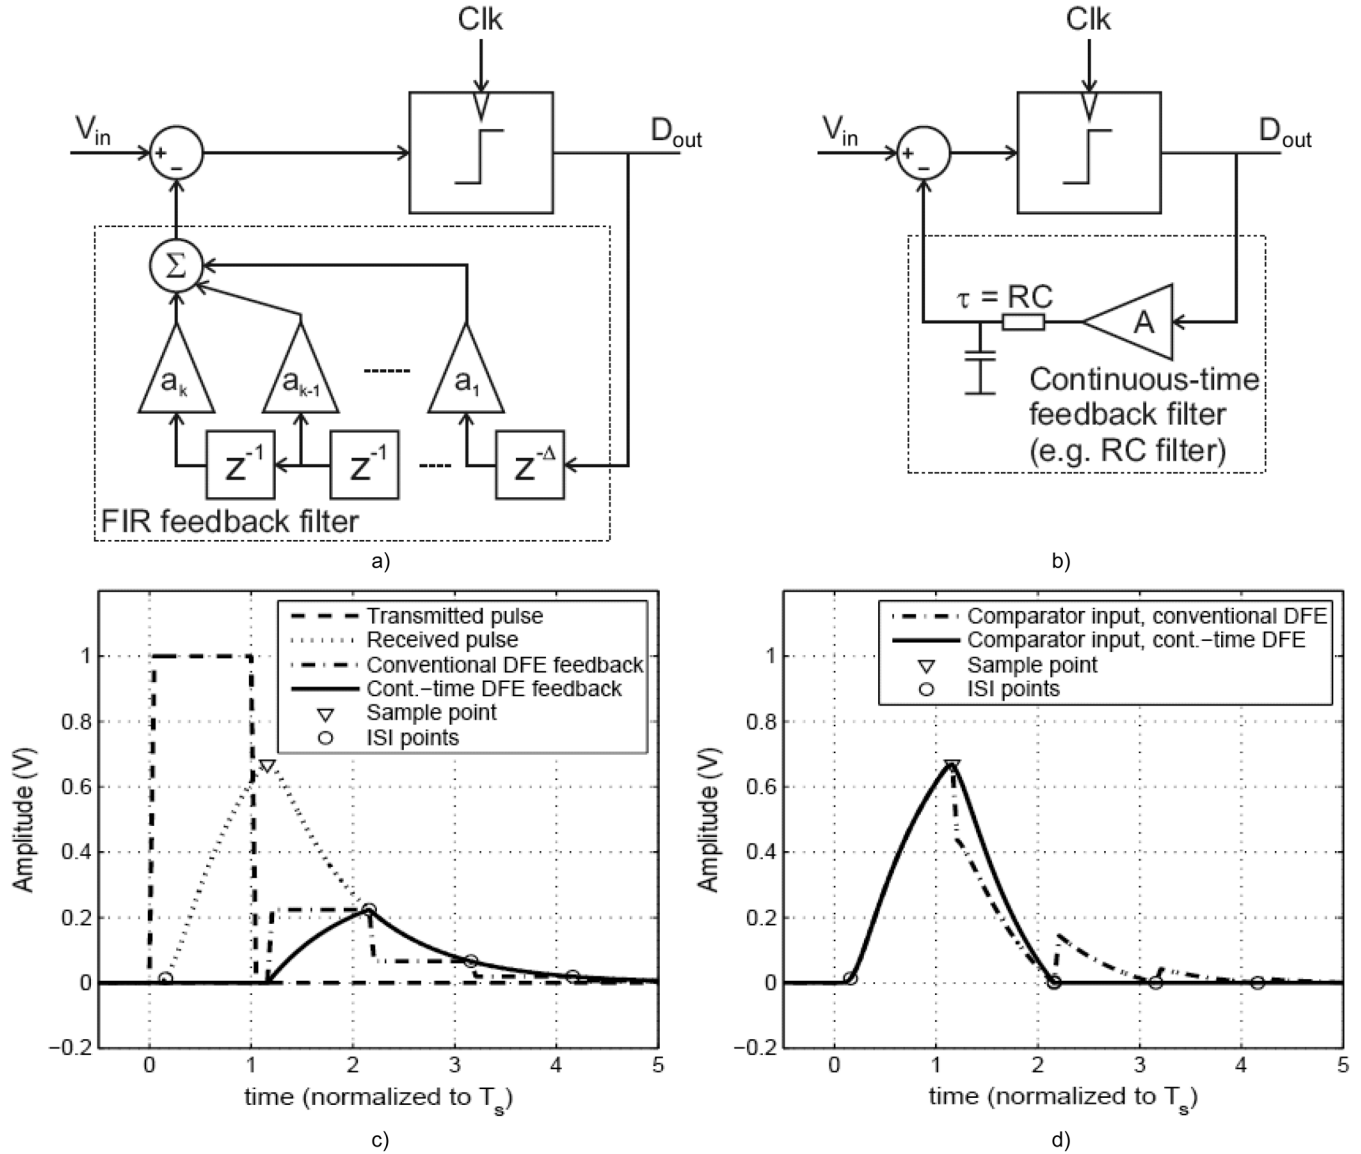
Fig. 7. Circuit diagram and simulated bit response for decision feedback equalization (DFE) with a discrete-time feedback filter (circuit a) $\Leftrightarrow$ response c) or a continuous-time (analog) feedback filter (circuit b $\Leftrightarrow$ response d).

a non-zero voltage across a resistor. Still it is more energy efficient than a conventional scheme with high swing along the whole line, except for very low p_{trans}, where static power dominates. The capacitive transmitter scheme has the lowest energy consumption (lowest slope and no static energy consumption). From Fig. 9 it can be concluded that the capacitive transmitter scheme has much lower energy consumption than the resistive termination scheme, although both increase the achievable data rate by about the same factor. There are two reasons for this lower energy consumption. The first is the static energy consumption of the current-sensing scheme that is not present in the capacitive transmitter scheme. This is because there is a resistive path from V_s to ground in Fig. 3(b), which lacks in Fig. 3(a) and 3(c), leading to a static current of about ± 0.3 mA for a static 1 or 0 ($Vdd/2 = 0.6$ V divided by the total series resistance of about 2 Kohm). This leads to at least 0.3 mA $*$ 0.6 V/1.25 Gbps $= 0.14$ pJ/bit static energy dissipation for current-sensing. If the current-sensing amplifier is not modeled

by a simple resistor but with a transconductor (e.g., and inverter) with resistive feedback, substantial current is also needed to realize sufficient transconductance to create low-ohmic current-sensing. The second reason for the attractiveness of the capacitive transmitter is the associated lower voltage swing on the interconnect, which reduces dynamic power. Although for both cases, the voltage swing at the receiver end of the interconnect is the same, the capacitive transmitter scheme has this low voltage swing along the entire interconnect, while the resistive termination scheme has a linearly increasing voltage swing towards the transmitter. This is shown in Fig. 10, where the voltage step responses of both schemes are given for different positions along the interconnect.

V. CIRCUIT IMPLEMENTATIONS

Our goal is to achieve a high data rate over 10 mm on-chip interconnects, aiming at minimal area and energy consumption.
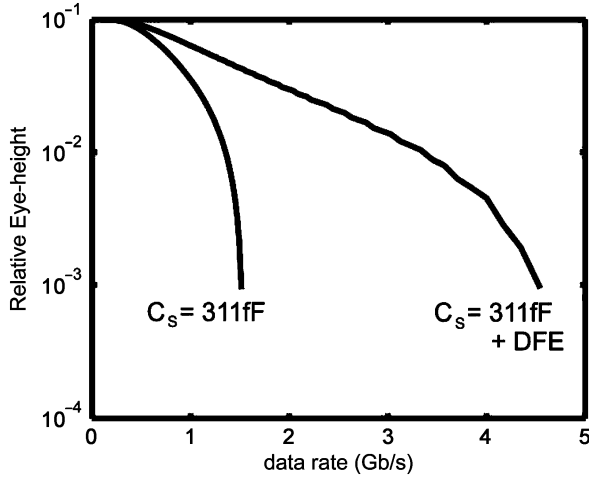
Fig. 8. Simulated relative eye-height as a function of data rate for a capacitive transmitter [Fig. 3(c)] and a capacitive transmitter with DFE at the receiver side [see Fig. 7(b)].

Fig. 10. Voltage step responses at different positions along the interconnect ($z = 1$ is at the receiver end of the interconnect) for: a) capacitive transmitter; b) current-sensing scheme.
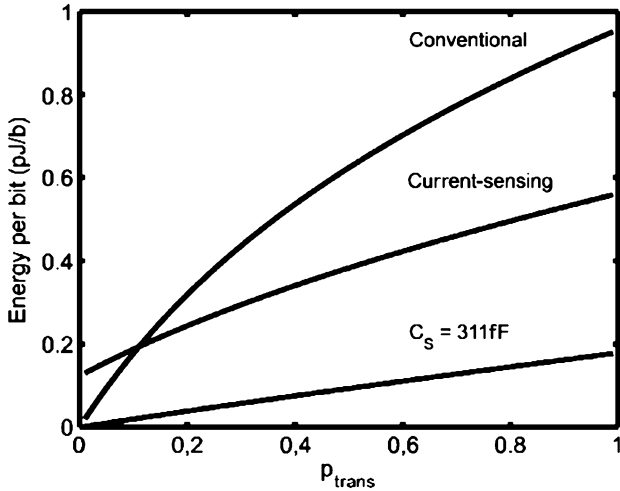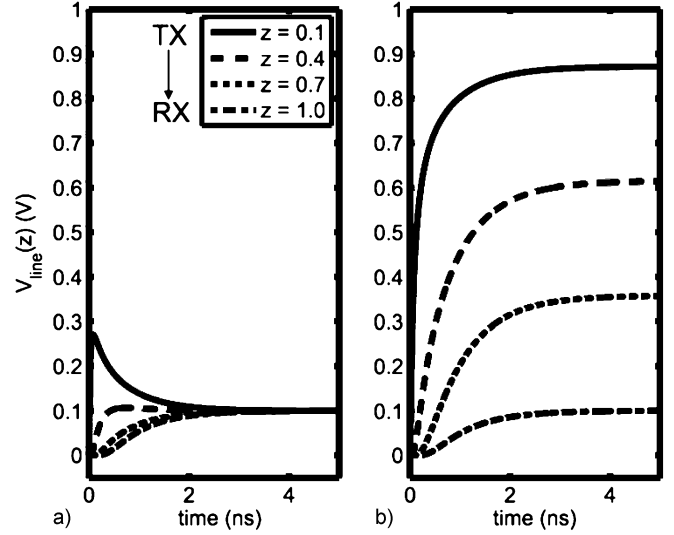

Fig. 9. Simulated energy consumption as a function of transition probability for the three different termination schemes of Fig. 3 working at about 1% eye-opening according to Fig. 4 (0.5 Gbps for conventional and 1.25 Gbps for the two improved schemes).
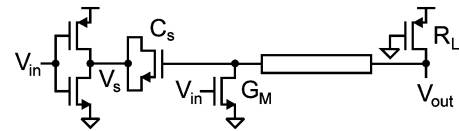

Fig. 11. Circuit implementation of the capacitive transmitter. The $G_M R_L$ combination is used to define the DC potential on the interconnect.

some kind of shield at the edge of the bus might be needed, causing only a relatively low area overhead for wide busses.

A. Capacitive Transmitter

The transmitter circuit is implemented as shown in Fig. 11. The series capacitance is made with an nMOS transistor. Due to the thin gate oxide, the area of the capacitor can be kept rather small compared to the area of the interconnect. A possible problem of the capacitive transmitter is the ill-defined DC potential on the interconnect. In order to define this DC potential for high and low V_{in}, a voltage-controlled current source with current $G_M V_{in}$ is added at the transmitter side and a resistance R_L at the receiver side. If V_{in} switches between 0 and V_{DD}, V_{out} switches between V_{DD} and $V_{DD} - G_M V_{DD} R_L$. The low frequency voltage swing on the interconnect is thus $G_M R_L V_{DD}$, which is chosen at $0.08 \cdot V_{DD}$ ($V_{DD} = 1.2$ V). By choosing G_M small (narrow NMOST) and R_L large (narrow long PMOST), the static energy consumption is kept small. We chose $G_m \approx 5\ \mu S$ and $R_L \approx 16$ kΩ, which renders about 1.2 V $\cdot$ 5 μS $\cdot$ 16 kΩ $\approx$ 100 mV swing at a current which is switched between 0 and 6 μA. For a differential line, one current is on while the other is off, so a total continuous bias current of 6 μA. For gigabit communication this leads to a negligible power overhead, e.g., < 0.01 pJ/b@1 Gbps), while the dynamic power consumption is in the order of 0.15 pJ/b for 50% transition probability.

With this setup, the transfer function at low frequencies is controlled by G_M and R_L, while at high frequencies the capacitive path via C_S and C_{wire} dominates. The question may now

We chose to implement the capacitive transmitter in combination with decision feedback equalization (DFE) at the receiver, as in this way a high data rate is possible with minimal energy consumption.

We use thin wires for optimum bandwidth per area [7]: Metal4 lines of 0.54 μm width and 0.32 μm spacing in a 90 nm technology. We used minimum sized inverters consisting of a 0.72/0.1 pMOS and 0.24/0.1 nMOS transistor. In order to be robust against crosstalk from wires in other metal layers and against supply and substrate noise, we make use of differential interconnects. If wires cross orthogonally, the effective coupling capacitance is small, and the differential receiver can handle the resulting common-mode cross-talk. Crosstalk from neighboring interconnects belonging to the same bus is minimized by alternatingly placing one or two twists in the differential interconnects, as analyzed in [16]. Running full-swing lines in parallel to low-swing lines might cause problems, and
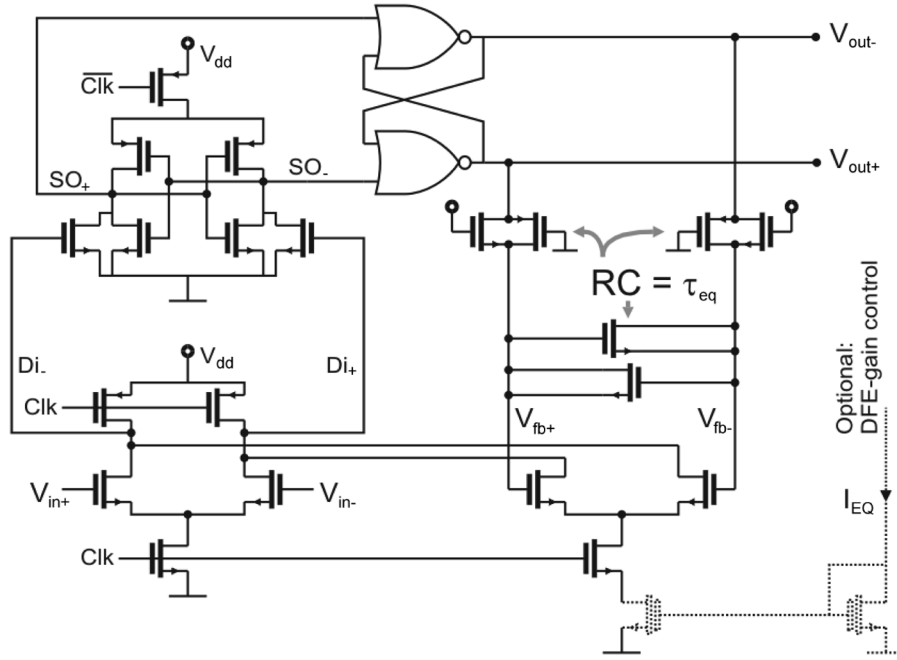
Fig. 12.   Sense amplifier with decision feedback equalization using an dynamically biased analog feedback path with a passive RC-filter.

arise how we can match the low- and high-frequency transfer functions and get a smooth transition. To get some first order insight, we analyzed the frequency transfer function of the *RC* line driven by both the current source and capacitive driver. If we assume that the interconnect can be approximated by a first-order *RC* model with an equivalent resistance of R_T and an equivalent capacitance of C_T, the transfer function from V_S to V_{out} in Fig. 11 is shown in the equation at the bottom of the page.

This transfer function has two poles and a zero. In order to get a first-order *RC* response, but now with extended bandwidth, the C_S has to be chosen as

$$C_S = \frac{G_M \cdot R_L \cdot C_T \cdot (1 + G_M \cdot R_T)}{1 - G_M \cdot (R_T + R_L)}.$$

For small $G_M R_T \ll 1$ and $G_M R_L \ll 1$, this equation can be approximated by

$$\frac{C_S}{G_M} = R_L \cdot C_T.$$

Thus, to match the low- and high-frequency transfer function, the two time constants C_S/G_M and $R_L C_T$ should be equal, with C_S the source capacitance and C_T the total capacitance of the interconnect. Simulations showed that inequality of the time constants has only a modest effect on the eye-opening, so that process variations can be tolerated if nominally equal time constants are chosen at design time.

The total area of the transmitter is 226 μm^2, where about 100 μm^2 is used for two MOSFET-line-driver capacitors, each

$C_S \approx 311$ fF. This is 5 times smaller than the metal capacitors used in [20] (40 × 20 metal tracks to implement a capacitive line driver, taking already about 480 μm^2). It comes at the cost of a more nonlinear capacitor, but eye diagram simulations show that the linearity of the capacitance is not very critical. Thus, a MOS capacitance with much higher capacitance/area can be used instead of metal-metal capacitance.

B. Sense Amplifier With Decision Feedback

Due to the reduced signal swing and high data rate, a sensitive receiver is needed. We implemented these receiver circuits as dynamic circuits to realize low power consumption, as shown in Fig. 12.

The left part of the circuit constitutes a clocked comparator, often also called a 'sense amplifier based flip-flop'. It only consumes power during a short time after a clock edge ("dynamic circuit"). Compared to traditional topologies, charge kick-back is reduced because there is an extra MOS transistor between the Di and So nodes, acting as a shield [15]. The sense amplifier can work at high common mode input voltage with a low offset and high speed (18 ps setup+hold time), as described in more detail in [15]. The SR-latch behind the sense amplifier is used to convert the dynamic (pre-charged) signals at the SO nodes to static CMOS signals that are valid for a whole clock-period. The outputs of the SR-latch are directly used to drive a low-pass *RC* filter for the DFE. The feedback voltage from the low-pass filter is coupled back into the sense amplifier via a second differential input stage, as shown on the right

$$\frac{V_{out}}{V_S} = \frac{j \cdot \omega \cdot C_S \cdot R_L + G_M \cdot R_L}{1 + j \cdot \omega \cdot ((R_T + R_L) \cdot C_S + R_L \cdot C_T) + (j \cdot \omega)^2 \cdot R_T \cdot R_L \cdot C_S \cdot C_T}$$

of Fig. 12. The DFE gain-factor "A" in Fig. 7(b) is defined by the transconductance-ratio of the feedback and main amplifier, based on: 1) the attenuation of the capacitive divider; 2) the ratio of the desired sample and ISI point [see Fig. 7(c)]. A 'dynamic' differential feedback pair is used with a switched tail current, so again no power will be consumed when the clock is inactive. The fact that the feedback output V_{fb} is full-swing, while a differential pair is usually only linear over a small input range, poses no real problems in this circuit if it is dimensioned properly. The linear range of the feedback differential pair is maximized by giving the transistors a high overdrive voltage (meaning small W and large L). The fact that the tail MOST operates in its triode region also helps to increase the input range, as the R_{ds} of the tail transistor acts as a degeneration resistance when only one of the two transistors of the differential pair is active.

The feedback gain-factor A can be controlled by the tail-current of the feedback differential pair. Usually, it is sufficient to set this gain-factor at design-time, through proper dimensioning of the clocked tail transistor. If desired, the tail-current can also be controlled at run-time, for example through a current-mirror configuration, as shown with the dashed transistors in Fig. 12.

The components that determine the time-constant, the resistor and capacitor, can be implemented in various ways, but we aimed for small area consumption. That is why the resistors and capacitor producing V_{fb} (see Fig. 12) have been implemented with MOS transistors, with pass-gates and antiparallel gate-capacitances respectively. The gate-capacitances of a MOST have a very high capacitance per area, but are also quite nonlinear due to the channel-capacitance. The use of an antiparallel configuration reduces the nonlinear effects to tolerable levels.

The total area of the receiver is 117 μm^2 with 32 μm^2 for the DFE part. The simulated power consumption is 0.12 pJ/b with 0.02 pJ/b for the DFE part at 2 Gbps.

C. Clocking Strategy

Both the capacitive transmitter and the sense amplifier with DFE require a clock, and hence a clocking strategy is needed to align the receiver to the eye of the incoming data. In principle a source synchronous clocking strategy can be used, where the clock is sent along with the data over an additional wire-pair. This is shown in [19] where we proposed a transceiver system for capacitively driven 2 mm long lines for a NoC. In this case a full-swing clock is used and the line losses are low enough to allow a single inverter to restore a full-swing clock, especially when a half-rate clock can be used [19]. Here, we use 10 mm long lines with much more high-frequency loss so that an attenuated sine-wave-like signal would result at the clock receiver. Restoration to a full-swing square-wave clock with a cascade of inverters is challenging in this case, as delay uncertainties for instance due to random offsets "eat" into the available eye width. In this case the use of a local copy of the global clock seems more attractive, where an appropriate skew depending on the line length is implemented at design-time, to align the receiver to the middle of the eye. We will show in the measurements section that the eye width in the receiver is large, leaving quite some margin for random spread of the clock skew.
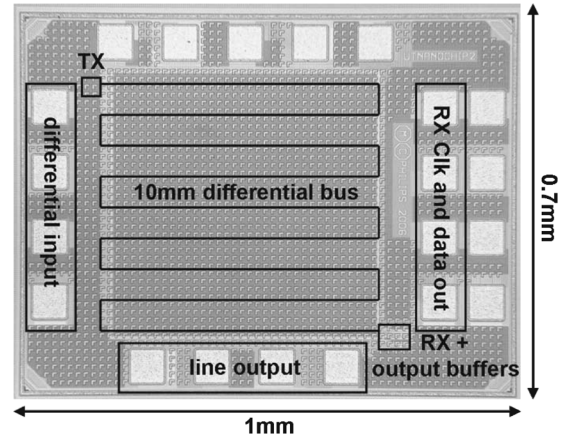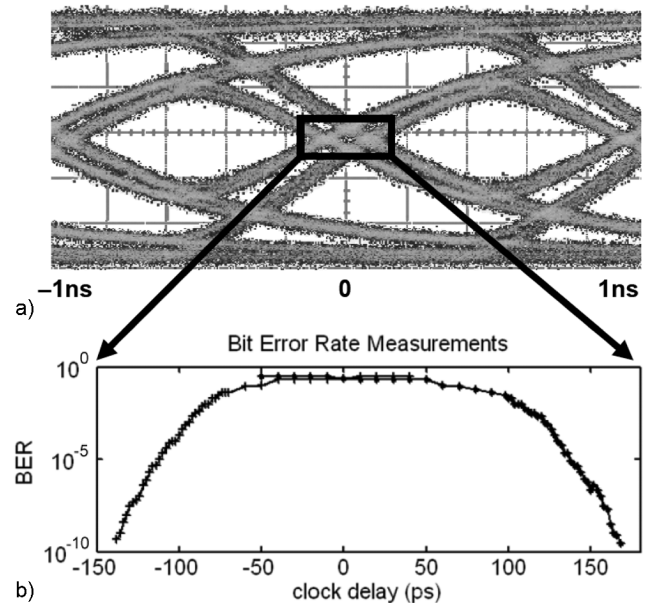

Fig. 13. Chip micrograph.

Fig. 14. Measured eye-diagram for a capacitively driven 10 mm line at (a) 1 Gb/s and (b) BER. DFE is not used.

VI. MEASUREMENTS

A demonstrator IC was fabricated in a CMOS 90 nm process. The chip micrograph is in Fig. 13. An external pattern generator/analyzer is used for data 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. Eye-diagrams are measured via 50 Ω output buffers that are connected to the output of a differential interconnect.

The measured interconnect parameters are $R' = 0.20$ kΩ/mm and $C' = 0.28$ pF/mm for a differential interconnect. A measured eye-diagram for the capacitively driven line at a data rate of 1 Gb/s is shown in Fig. 14. The measured BER at the edges of the eye is also shown. The BER drops rapidly below a clock skew of -150 ps and above 180 ps, giving an eye-opening of 670 ps. Data rates up to 1.35 Gb/s are achieved without decision feedback equalization (DFE) at the receiver side (DFE-gain control current $I_{EQ} = 0$). The one-σ offset of the total
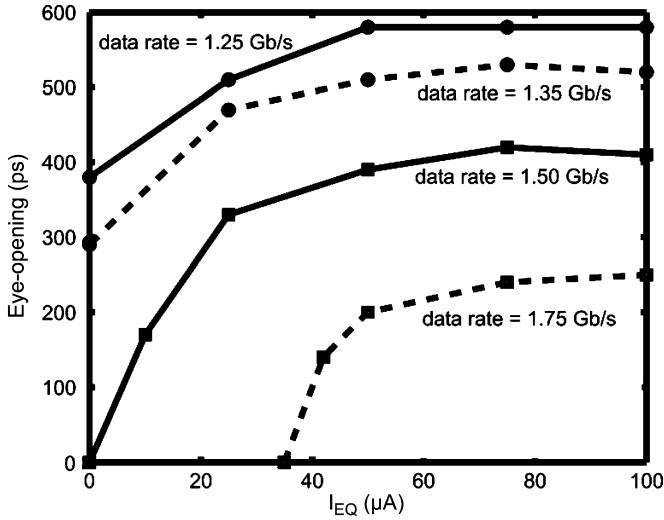
Fig. 15. Measured eye-opening as a function of I_{EQ} for different data rates.

Fig. 17. Energy consumption per bit per mm line length versus the achieved data rate [Gb/s] per cross-sectional area [μm^2] multiplied by the squared line length [mm^2] (see text for motivation).

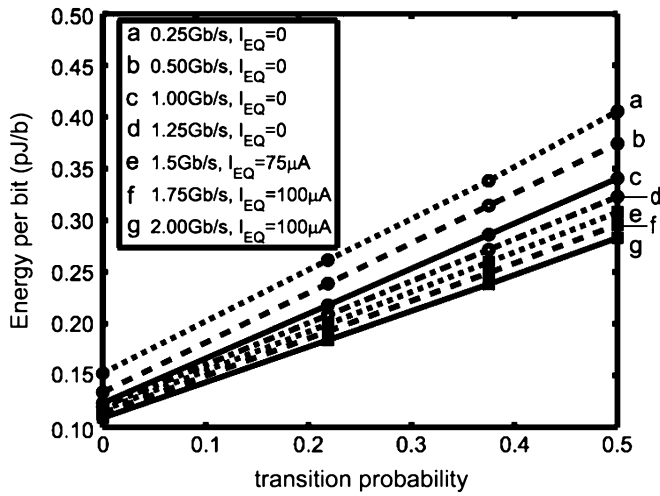Fig. 16. Measured energy consumption per transmitted bit as a function of the transition probability for different data rates, with and without DFE.

inactive periods. The DFE part of the circuit requires less than 7% of the total transceiver power, while it increases the achievable data rate here with a factor 1.5.

VII. COMPARISON

We will now compare the results of our demonstrator IC with other solutions found in literature, considering both the achievable data rate and the energy consumption. The energy consumption depends roughly linearly on the line length as line capacitance scales linear with line length. For a meaningful comparison, we will divide the energy consumption by the line length. As the bandwidth of RC-limited interconnects depends inversely on the square of the line length, we will divide the achievable data rate by the squared line length. As data rate can be increased by using more cross-sectional area (either by using parallel wires or increasing the bandwidth by reduction of the wire resistance), it makes sense to divide the achievable data rate by the cross-sectional area of the interconnect [7]. 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 the vertical spacing to other metal layers (see Fig. 1). For those papers where the parameters s, h and d are not given, they are estimated based on IC technology parameters from a similar process.

Fig. 17 compares the energy per bit per mm line length achieved by different published designs. On the x axis the achievable data rate divided by the cross-sectional area and multiplied by the length squared is indicated 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 normalized data rate and much lower energy consumption than all other previous proposals.

Recently, we proposed to use the capacitive driver for medium length interconnects for Networks on Chip [19]. Simulations for 2 mm line length predict an achievable data rate of 9 Gbps at similar energy per bit per mm as the transceiver published in this paper. This is faster and more power efficient

transceiver is 11 mV, measured over 20 samples. Due to this offset, not all samples achieve 1.35 Gb/s, but all samples do achieve a slightly lower data rate of 1 Gb/s. If desired, area up-scaling could further reduce the offset at the expense of power ($P \propto 1/\sigma_{os}^2$) [15]. Offset compensation schemes can be a good alternative if the application allows for the added complexity, which is probably not the case for most on-chip buses. However, simulations over process corners indicate that the circuit is robust to PVT variations at a rate slightly lower than the maximum achievable data rate. Data rates up to 2 Gb/s are measured with DFE. Note that DFE reduces ISI, making the system less vulnerable to offset. Fig. 15 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.

The measured energy consumption at different data rates is shown in Fig. 16. With random data at 2 Gb/s, only 0.28 pJ/b is dissipated. The energy dissipation of 0.12 pJ/b at zero data activity is mainly due to the energy consumption in the sense amplifier, which has large transistors to get a low offset. Clock-gating can be used to eliminate its energy consumption during
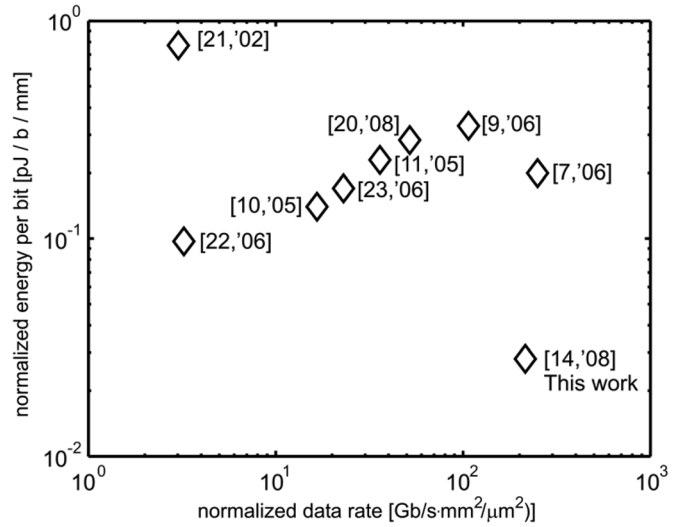
TABLE I
COMPARISON OF PUBLISHED ON-CHIP DATA LINKS USED AS A BASIS FOR FIG. 17

Property	Units	[21] Ch'02	[11] Zh'05	[10] Ka'05	[9] Ba'06	[22] Jo'06	[23] Jo'07	[20] Ho'08	[7] Sc'06	This Me'08
Technology node	[nm]	180	180	130	350	180	180	180	130	**90**
Single / Differential		Diff	Single	Single	Single	Diff.	Diff.	Diff.	Diff.	**Diff.**
Supply Voltage	[V]	1.8	1.8	1.2	2.5	1.8	1.8	1.8	1.2	**1.2**
Line length l	[mm]	20	10	10	17.5	3	14	10*	10	**10**
Line width w	[μm]	2·16	4.5	0.6	2	2·(4+4)	2·8	2·0.3	2·0.4	**2·0.54**
Line spacing s	[μm]	2·1*	1*	0.63	1*	2·4	2·8	2·0.3	2·0.4	**2·0.32**
Metal height h	[μm]	2	0.5*	0.35	0.5*	0.53	0.53	0.5*	0.35	**0.33**
Oxide thickness d	[μm]	1.9	0.5*	0.36	0.5*	0.5*	0.5*	0.5*	0.46	**0.27**
Cross-sectional Area $A_c = (w+s)(h+d)$	[μm²]	133	5.5	0.87	3	24.7	33	1.2	1.3	**1.03**
Achieved data rate f_D	Gbps	1	2	0.2	1	8	3	1	3	**2**
Energy per bit E_b	pJ/b	16.1	2.3	1.7	5.8	0.29	2	2.24	2.0	**0.28**
Normalized Speed $= f_D \cdot l^2 / A_c$	Gbps· mm²/μm²	3.0	36	23	102	2.9	18	53	231	**194**
Normalized Energy $= E_b / length$	fJ/b/mm	805	230	170	331	97	143	280	200	**28**

* parameter not given in the paper; estimated values based on typical technology data

than achieved with a multi-VDD low-swing design in the same technology [19], while only requiring a single supply voltage.

VIII. CONCLUSION

To explore limits of increasing the achievable data rate of RC-limited interconnects, theoretical ideal source and load impedances for a flat transfer function with linear phase (pure delay) were derived, using s-parameters to model the interconnects. Either a capacitance as source impedance, or a resistance as load impedance increase the achievable data rate by a factor of three. Both variants come close to the theoretical ideal termination impedances. However, with respect to power consumption, a transceiver that uses the capacitive transmitter outperforms a transceiver that uses a resistive load as both its static and its dynamic power consumption are lower. To further increase the achievable data rate, DFE can be used at the receiver. By using an analog feedback filter, DFE only costs little extra area and power. We presented experimental results for a 90 nm CMOS transceiver chip incorporating the mentioned techniques to communicate over 10 mm lines with a small cross-sectional area of about 1 μm^2. We achieve error free operation at 2 Gb/s, comparable to the fastest solutions found in literature, but at a much lower power consumption of 0.28 pJ/bit corresponding to 28 fJ/b/mm (see Fig. 17).

REFERENCES

[1] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proc. IEEE*, vol. 89, pp. 490–504, Apr. 2001.

[2] J. Nurmi, H. Tenhunen, J. Isoaho, and A. Jantsch, *Interconnect-Centric Design for Advanced SoC and NoC*. Boston, MA: Kluwer Academic, 2004.

[3] L. Benini and G. De Micheli, "Networks on chips: A new SoC paradigm," *IEEE Computer*, vol. 35, pp. 70–78, Jan. 2002.

[4] W. J. Dally and B. Towles, "Route packets, not wires: On-chip interconnection networks," in *Proc. Design Automation Conf.*, 2001, pp. 684–689.

[5] H. Bakoglu, *Circuits, Interconnections and Packaging for VLSI*. Reading, MA: Addison-Wesley, 1990.

[6] P. Larsson-Edefors, "Investigation on maximal throughput of a CMOS repeater chain," *IEEE Trans. Circuits and Systems I: Fundamental Theory and Applications*, vol. 47, pp. 602–606, Apr. 2000.

[7] D. Schinkel, E. Mensink, E. A. M. Klumperink, E. van Tuijl, and B. Nauta, "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.

[8] E. Seevinck, P. J. van Beers, and H. Ontrop, "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.

[9] 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," *IEEE J. Solid-State Circuits*, vol. 41, pp. 461–473, Feb. 2006.

[10] A. Katoch, H. Veendrick, and E. Seevinck, "High speed current-mode signaling circuits for on-chip interconnects," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, May 2005, pp. 4138–4141.

[11] L. Zhang, J. Wilson, R. Bashirullah, L. Lei, X. Jian, and P. Franzon, "Driver pre-emphasis techniques for on-chip global buses," in *Proc. Int. Symp. Low Power Electronics and Design (ISLPED)*, Aug. 2005, pp. 186–191.

[12] K. Chang-Ki, R. Kwang-Myoung, and L. Kwyro, "High speed and low swing interface circuits using dynamic over-driving and adaptive sensing scheme," in *Proc. Int. Conf. VLSI and CAD*, Oct. 1999, pp. 388–391.

[13] R. Ho, T. Ono, F. Liu, R. Hopkins, A. Chow, J. Schauer, and R. Drost, "High-speed and low-energy capacitively-driven on-chip wires," in *IEEE Int. Solid State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2007, pp. 412–413, 612.

[14] E. Mensink, D. Schinkel, E. Klumperink, E. van Tuijl, and B. Nauta, "A 0.28 pJ/b 2 Gb/s/ch transceiver in 90 nm CMOS for 10 mm on-chip interconnects," in *IEEE Int. Solid State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2007, pp. 414–415, 612.

[15] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "A double-tail latch-type voltage sense amplifier with 18 ps setup+hold time," in *IEEE Int. Solid State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2007, pp. 314–315, 605.

[16] E. Mensink, D. Schinkel, E. Klumperink, E. van Tuijl, and B. Nauta, "Optimal positions of twists in global on-chip differential interconnects," *IEEE Trans. VLSI Systems*, vol. 15, pp. 438–446, Apr. 2007.

[17] E. Mensink, "High-speed global on-chip interconnects and transceivers" Ph.D. dissertation, University of Twente, Enschede, The Netherlands, 2007 [Online]. Available: http://purl.org/utwente/57868, 978-90-365-2504-6

[18] V. Stojanovic, A. Ho, and B. Garlepp *et al.*, "Adaptive equalization and data recovery in a dual-mode (PAM2/4) serial link transceiver," in *Symp. VLSI Circuits Dig.*, Jun. 2004, pp. 348–351.

[19] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "Low-power, high-speed transceivers for network-on-chip communication," *IEEE Trans. VLSI Systems*, vol. 17, no. 1, pp. 12–21, Jan. 20, .

[20] R. Ho, T. Ono, R. D. Hopkins, A. Chow, J. Schauer, F. Y. Liu, and R. Drost, "High speed and low energy capacitively driven on-chip wires," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 52–60, Jan. 2008.

[21] R. T. Chang, C. P. Yue, and S. S. Wong, "Near speed-of-light on-chip electrical interconnect," in *Symp. VLSI Circuits, Dig. Tech. Papers*, Jun. 2002, pp. 18–21.

[22] 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, pp. 772–780, Apr. 2006.

[23] A. P. Jose and K. L. Shepard, "Distributed loss compensation for low-latency on-chip interconnects," in *IEEE Int. Solid State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2006, pp. 516–517.
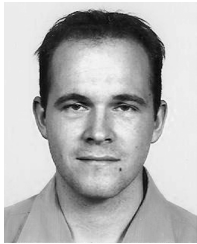
Eisse Mensink (S'03–M'07) was born in Almelo, The Netherlands, in 1979. He received the M.Sc. degree in electrical engineering (with honors) from the University of Twente, Enschede, The Netherlands, in 2003. In 2007 he received the Ph.D. degree from the same university on the subject of high-speed on-chip communication.

He is currently an ASIC design engineer at Bruco B.V., Borne, The Netherlands.

Daniël Schinkel (S'03–M'08) was born in Finsterwolde, The Netherlands, in 1978. He received the M.Sc. degree in electrical engineering (with honors) from the University of Twente, The Netherlands, in 2003. From 2003 to 2007, he worked as a Ph.D. student at the same university at the IC-design group headed by Bram Nauta. During this period he also occasionally worked as a freelance consultant on the subject of sigma-delta converters. He is currently writing his thesis on high-speed on-chip communication.

He is one of the founders of Axiom IC, an IC-design company that started in 2007 and focuses on the design of state-of-the-art analog and mixed signal circuits. His research interests include analog and mixed-signal circuit design, sigma-delta data converters, class-D power amplifiers and high-speed communication circuits. He holds two patents and is author or coauthor of 17 papers.

Eric A. M. Klumperink (M'98–SM'06) was born on April 4, 1960, in Lichtenvoorde, The Netherlands. He received the B.Sc. degree from HTS, Enschede, The Netherlands, in 1982.

After a short period in industry, he joined the Faculty of Electrical Engineering of the University of Twente (UT) in Enschede, in 1984, participating in analog CMOS circuit design and research. This resulted in several publications and a Ph.D. thesis, in 1997 ("Transconductance based CMOS circuits"). After his PhD, he started working on RF CMOS circuits, and he is currently an Associate Professor at the IC-Design Laboratory which participates in the CTIT Research Institute (UT). He holds several patents and has authored and co-authored more than 80 journal and conference papers.

In 2006 and 2007, Dr. Klumperink served as Associate Editor for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II, and since 2008 for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I. He was a co-recipient of the ISSCC 2002 Van Vessem Outstanding Paper Award.

Ed (A. J. M.) van Tuijl (M'97) was born in Rotterdam, The Netherlands, on June 20, 1952.

He joined Philips Semiconductors, Eindhoven, The Netherlands, in 1980. As a Designer, he worked on many kinds of small-signal and power audio applications, including A/D and D/A converters. In 1991, he became Design Manager of the audio power and power-conversion product line. In 1992, he joined the University of Twente, Enschede, The Netherlands, as a part-time Professor. After many years at Philips Semiconductors, he joined Philips Research, Eindhoven, The Netherlands, in 1998 as a Principal Research Scientist. He is one of the founders of Axiom IC, an IC-design company that started in October 2007 and focuses on the design of state-of-the-art analog and mixed signal circuits. His current research interests include data conversion, high-speed communication, and low-noise oscillators. He is an author or coauthor of many papers and holds many patents in the field of analog electronics and data conversion.

Bram Nauta (M'91–SM'03–F'07) was born in Hengelo, The Netherlands, in 1964. In 1987 he received the M.Sc. degree (cum laude) in electrical engineering from the University of Twente, Enschede, The Netherlands. In 1991 he received the Ph.D. degree from the same university on the subject of analog CMOS filters for very high frequencies.

In 1991 he joined the Mixed-Signal Circuits and Systems Department of Philips Research, Eindhoven The Netherlands, where he worked on high speed AD converters and analog key modules. In 1998 he returned to the University of Twente, as full professor heading the IC Design group, which is part of the CTIT Research Institute. His current research interest is high-speed analog CMOS circuits. He is also part-time consultant in industry and in 2001 he co-founded Chip Design Works.

His Ph.D. thesis was published as a book: *Analog CMOS Filters for Very High Frequencies* (Springer, 1993) and he received the Shell Study Tour Award for his Ph.D. work. From 1997 until 1999 he served as Associate Editor of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II—ANALOG AND DIGITAL SIGNAL PROCESSING. After this, he served as Guest Editor, Associate Editor (2001–2006), and since 2007 as Editor-in-Chief for the IEEE JOURNAL OF SOLID-STATE CIRCUITS. He is also a member of the technical program committees of the IEEE International Solid State Circuits Conference (ISSCC), the European Solid State Circuit Conference (ESSCIRC), and the Symposium on VLSI Circuits. He was a co-recipient of the ISSCC 2002 Van Vessem Outstanding Paper Award, and is a distinguished lecturer of the IEEE and an elected member of IEEE SSCS AdCom.