

Low-Power, High-Speed Transceivers for Network-on-Chip Communication

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

Abstract—Networks on chips (NoCs) are becoming popular as they provide a solution for the interconnection problems on large integrated circuits (ICs). But even in a NoC, link-power can become unacceptably high and data rates are limited when conventional data transceivers are used. In this paper, we present a low-power, high-speed source-synchronous link transceiver which enables a factor 3.3 reduction in link power together with an 80% increase in data-rate. A low-swing capacitive pre-emphasis transmitter in combination with a double-tail sense-amplifier enable speeds in excess of 9 Gb/s over a 2 mm twisted differential interconnect, while consuming only 130 fJ/transition without the need for an additional supply. Multiple transceivers can be connected back-to-back to create a source-synchronous transceiver-chain with a wave-pipelined clock, operating with 6σ offset reliability at 5 Gb/s.

Index Terms—Capacitive pre-emphasis transmitter, globally asynchronous, locally synchronous (GALS), interconnect, low-power design, low-swing, network on chip (NoC), on-chip communication, source synchronous, wave-pipelining.

I. INTRODUCTION

ON-CHIP communication has become an active research area in the past few years. This not only because on-chip interconnects are becoming a speed, power, and reliability bottleneck [1], but also because systems on chips (SoCs) start to become so complex that they require new interconnection approaches [2], [3].

Networks on chips (NoCs) have emerged as the seemingly best candidate to connect the many functional elements on present and future SoCs [2]–[7]. Most of the long (global) interconnects, which have the severest bandwidth limitations and crosstalk problems, are eliminated in a NoC, especially when mesh-like network configurations are used. A NoC also enables easier clock-distribution with alleviated skew requirements and less power consumption as the various processing elements can operate mesochronous [4]–[6] or asynchronous

[7] to each other, using for example the globally asynchronous, locally synchronous (GALS) design style.

Still, even in a NoC configuration, the network interconnects and especially the routers can consume a considerable part of the total power budget. In [6], for example, the on-chip network consumes up to 39% of the total chip power (76 W when operating at 5.1 GHz) [8]. 17% of the network power is consumed in the links (13 W at 5.1 GHz).

A NoC can therefore benefit from link-transceivers that are more advanced than the standard inverters. High-speed, low-power transceivers can for example facilitate network topologies with longer and more wires than the standard mesh topology, such as a (folded) torus or star topology, to simultaneously reduce the interconnect power and the average hop count, and hence also the latency and the associated router power [3], [4].

A number of on-chip transceiver improvements have been proposed in the past, but they usually reduce either the power consumed in the interconnect [4], [9] or improve the data-rate achievable over the interconnect [10], [11]. In a recent paper [12], we presented transceiver techniques for global on-chip interconnects which both increase the achievable data-rate and decrease the transmission power.

In this paper, we will adapt these techniques for NoC applications and compare the resulting transceiver with other common types of transceivers. Other topics that were not covered in previous publications are the optimization of the circuit for yield versus power and the addition of synchronization circuitry. Yield is an important issue given PVT variations, random mismatch, crosstalk, and the fact that many transceivers will be present on a NoC.

A schematic overview of the proposed NoC transceiver is shown in Fig. 1. The transmitter uses a series capacitance to lower the swing on the interconnect, increase its bandwidth and lower the power dissipation. The interconnects consist of twisted differential pairs to be robust towards disturbances such as supply noise and crosstalk [13]. An improved sense amplifier [14] clocks the data at the receiving end and regenerates it to full swing. A clock or strobe channel is present alongside the data-channels to enable source-synchronous operation.

This paper is organized as follows. Section II discusses data links for networks on chip and the drawbacks of conventional transceivers. Section III describes the improved low-swing transmitters and Section IV discusses the accompanying receivers. Section V includes synchronization in the discussion and describes the entire transceiver. The paper ends with the conclusions in Section VI.

Manuscript received August 02, 2007; revised January 08, 2008. First published November 18, 2008; current version published December 17, 2008. This work was supported by the Technology Foundation STW, an applied science division of NWO, and the technology program of the Ministry of Economic Affairs, under project TCS.5791.

D. Schinkel is with Axiom-IC B.V., 7521PT Enschede, The Netherlands (e-mail: daniel.schinkel@axiom-ic.com).

E. Mensink is with Bruco, 7623CS Borne, The Netherlands (e-mail: eisse.mensink@bruco.nl).

E. A. M. Klumperink and B. Nauta are with the IC Design Group, University of Twente, 7500 AE Enschede, The Netherlands (e-mail: e.a.m.klumperink@utwente.nl; b.nauta@utwente.nl).

A. J. M. van Tuijl is with the University of Twente, 7500 AE Enschede, The Netherlands and also with Axiom IC B.V., 7521 PT Enschede, The Netherlands (e-mail: ed.van.tuijl@axiom-ic.com).

Digital Object Identifier 10.1109/TVLSI.2008.2001949




Fig. 1. Overview schematic of the proposed transceiver for NoCs.

II. DATA COMMUNICATION ON A NOC

A. Interconnects for NoCs

The high capacitance and high resistance of on-chip interconnects provide the grounds for the problems associated with interconnects. The high capacitance causes high power consumption and the mutual capacitance causes the dominant part of the crosstalk. The RC product limits the bandwidth. In a dense interconnect environment, the inductance of the interconnects does not play a significant role for lengths larger than a few tenths of a millimeter [10]. To characterize the interconnects, we used 3-D EM-Field solver simulations and measurements. The resulting parameters are used in lumped-element models (100 lumps) for circuit-level simulations.

In this paper, we will focus on interconnects that span one or two processing tiles. A wire length of 2 mm is assumed throughout the paper, but the same techniques apply to a variety of lengths. The transceiver presented in [12] focused on much longer (10 mm) wires and contains some additional equalization circuitry to boost the data-rate. Wires of 2 mm have a much higher intrinsic bandwidth (the RC product scales with the length squared [1]), so we will focus here on slightly simpler transceivers and leave out the receiver equalization.

We also assume that the interconnects are used unidirectional, as bidirectional use of the interconnects complicates the design of fast and power-efficient transceivers. Bidirectional communication can be implemented with a second set of interconnects, as is often done in NoCs.

To maximize the throughput between two routers, it makes sense to use wide data paths [3] with many densely packed interconnects. In [10] it was shown that the cross-sectional dimensions of interconnects should be chosen roughly equal to optimize the bandwidth per cross-sectional area (BW/Area). A bus




Fig. 2. Conventional transceiver schematic.

with these optimized interconnects will have the highest achievable throughput for a certain bus area (also see [15] and [16]). Wires in the thick (reverse-scaled) top-metal layers will have lower resistance and higher bandwidths so it makes sense to use the top metal layers for the link when the data-rate per wire is a limiting factor [2], [3]. However, the BW/Area is roughly the same as for thinner metal layers [10], so one could choose to also use the lower metal layers for the link. In this last case, certain areas of the chip could be dedicated to the link interconnects to enable high throughput in a well defined link environment.

To fully use the available BW/Area, it would also seem best to use single-ended interconnects. But, as will be shown in later sections, differential interconnects enable more robust transceivers that hardly suffer from crosstalk, can operate at higher speeds and at a lower swing, which is why the proposed transceiver uses differential wires.

In the 1.2-V, 6-M, 90-nm CMOS process that is used in this project, metal-4 wires with a width of $0.54 \mu\text{m}$ and a spacing of $0.32 \mu\text{m}$ have the highest BW/area under the assumption that the wires are surrounded by other wires in all directions. Under these conditions, the interconnect parameters are

$$R_{\text{wire}} = 200 \Omega/\text{mm}, \quad C_{\text{wire}} = 280 \text{ fF}/\text{mm} \quad (1)$$

or $C_{\text{wire}} = 240 \text{ fF}/\text{mm}$ for single-ended interconnects [12].

With these dimensions, one differential channel will have a pitch of $1.72 \mu\text{m}$. A link with for example a length of $l = 2 \text{ mm}$ and a width $w = 64 \text{ bits}$ in both directions occupies an area of $2wl \cdot 1.72 \mu\text{m} = 0.44 \text{ mm}^2$ when placed in one metal layer, which can still easily fit above a $2 \times 2 \text{ mm}$ tile. When five metal layers would be available to connect routers in a mesh topology with, e.g., $N \times N = 5 \times 5$ tiles of $2 \times 2 \text{ mm}$ each, then the total link area becomes $2N(N-1) \cdot 0.44/5 = 3.5 \text{ mm}^2$, only 4% of the tile area of 100 mm^2 . The total wire-length is then: $2 \cdot 2N(N-1) \cdot 2wl = 2 \cdot 2 \cdot 5 \cdot 4 \cdot 2 \cdot 64 \cdot 2 \text{ mm} = 20.48 \text{ m}$.

B. Conventional Data Transmission

In conventional digital IC design practice, interconnects that are used for chip-wide data communication are simply treated as part of the normal digital design flow, perhaps with a few additional steps such as the (automated) placement of repeaters, to minimize the delay per interconnect length [17].

An example of a “conventional transceiver” for data communication on a NoC is shown in Fig. 2. It does not have repeaters because delay optimal repeater insertion comes at the price of about 90% increase in power consumption (the C_{gates} add 60% to the total capacitance [1] and the C_{drains} add another 30%). Furthermore, for these relatively short wires, repeaters reduce the delay only marginally [1] as the dominant time-constant of the interconnect itself is still only $1/2R_{\text{wire}}C_{\text{wire}} = 96 \text{ ps}$. To be able to approach this intrinsic wire speed, the transmitter from Fig. 2 does need to use a buffer-cascade with a large

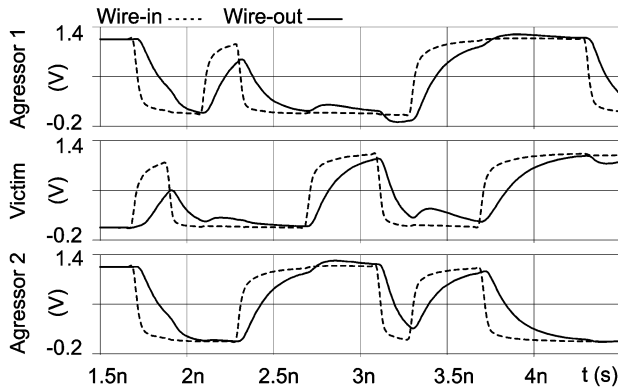


Fig. 3. Signals at 5 Gb/s for three neighboring channels from a conventional transceiver.

and power-hungry driver. In Section III, it will be shown that it is also possible to use a smaller and more power efficient low-swing capacitive transmitter.

In classical synchronous systems, the maximum delay of a combinatorial logic stage is limited to the clock period—or vice versa: the clock-rate is limited by the stage with the maximum delay—and this constraint is usually also imposed on the data transceivers. But such a constraint is not necessary for a communication channel, as is often demonstrated in wireline communication where several bits can be in flight along the channel at any given time. The channel bandwidth is the real limiting factor for the data-rate. For on-chip transceivers, it is also easy to achieve data-rates higher than $1/T_{\text{delay}}$ provided that proper clocking schemes are used, such as pipelined or source synchronous schemes, as will be demonstrated in Section V.

Without additional layout measures, a conventional transceiver is not very suitable as a high-speed transceiver, because its delay can vary widely due to crosstalk [1]. Fig. 3 shows the effect of capacitive crosstalk between neighboring data wires in a bus. The average delay of the transmitter and the 2 mm of interconnect amounts to 205 ps, but the delay speeds up to 160 ps when neighboring aggressors make a transition in the same direction and the delay increases to 262 ps when the neighboring aggressors switch in the opposite direction. Crosstalk not only creates this varying delay (reduced eye-width), but it also decreases the voltage noise margin (reduced eye-height) as is visible in Fig. 3. Above a certain data rate, crosstalk from specific aggressor data patterns can even prohibit proper detection of data bits, as visible for the bit in the victim signal at $t = 1.9$ ns. Quantitatively, crosstalk between neighboring wires in one metal layer can decrease the achievable data-rate by a factor of 1.7 [18]. Crosstalk problems become even worse when the surrounding metal layers are also used as data paths.

A standard method to reduce crosstalk is to increase the spacing between the wires or insert shield-wires and shield-planes [15], [19], where the latter option also helps to define a return path and reduce inductive crosstalk. To enable the highest data-rates for each channel, one would need to place a shield wire between every signal wire, but at the cost of increased wiring resources and possibly a lower BW/area [15].

A conventional transceiver is also not very power efficient as the transmitter needs to fully charge and discharge large wire

and driver capacitances. The setup shown in Fig. 2 consumes 775 fJ per upward transition (consumed mainly to charge the wire) and 65 fJ per downward transition (consumed to charge the driver capacitances), which averages to 420 fJ/transition. As an example for what this would cost on an entire chip, assume the same situation as earlier with 2 mm long 64 bits wide links in both directions, used in a mesh of 5×5 tiles. Furthermore assume a clock-frequency of 5 GHz for the links, with an average switching-activity of about $p_{\text{act}} = 25\%$ (heavy traffic). Then the total link power becomes $P = 2N(N-1) \cdot 2w \cdot p_{\text{act}} \cdot E/\text{trans} \cdot f_{\text{clk}} = 2 \cdot 5 \cdot 4 \cdot 2 \cdot 64 \cdot 0.25 \cdot 420 \cdot 10^{-15} \cdot 5 \cdot 10^9 = 2.7$ W, which is not acceptable for low-power applications such as mobile baseband processors [7]. The reported link power for the 80-tile NoC from [6] is even higher: 13 W at 5.1 GHz.

C. Link Improvements

It is well recognized that low-swing signaling can reduce the interconnect power consumption [9], [20], but at the cost of a reduced noise margin. The degradation of data integrity due to supply- and substrate-noise increases as the swing goes down. Crosstalk also becomes an even more severe problem, especially when a full-swing aggressor interconnect is routed in the vicinity of a low-swing victim.


Fortunately, the regular nature of the top-level wiring in a NoC and the re-usability of the interconnection links justify a slightly higher design-effort to better optimize the wires [3]. In this way, routing of full-swing wires next to low-swing wires can be avoided, as well as the routing of far-end wire parts next to near-end ones. Application of these simple rules leaves only the crosstalk between the different wires from the same bus, with the neighbor-to-neighbor crosstalk as dominant part.

Application of twisted differential wires can effectively mitigate neighbor-to-neighbor crosstalk, needing only one twist in every even wire pair and two twists in every odd pair [13], as indicated in Fig. 1. The optimal positions of these twists depend on the type of wire termination. With equal impedances for transmitter and receiver, intra-bus crosstalk is perfectly canceled and the optimal twist positions are symmetric around the midpoint [13].

The increase in power and area due to the doubling of the number of active wires is actually not that large, among others due to the earlier discussed overhead in shield wires for single-ended channels. Even with shields, single-ended wires are less immune towards disturbances than twisted differential channels, which can even make a differential interconnect more power efficient than a single-ended alternative because a differential transceiver can operate at a lower swing [9].

The immunity to (supply- or ground) disturbances is not only valid for the differential interconnects themselves but also for the receiver, as a differential sense amplifier with a low offset and a high power-supply rejection can be used [9], [14], which can operate reliably at much lower noise margins than a single-ended latch or logic cell. This advantage is shared by other alternatives that use single-ended data wires and a shared reference, such as the pseudo-differential interconnect from [9]. The ability to cancel crosstalk is however not present in pseudo-differential interconnection schemes.

In the presented transceiver, differential interconnects with twists are used. Due to the twists and the capacitive termination

Fig. 4. Low-swing transceiver with multiple V_{DD} 's.

at both transmitter and receiver-side, practically all crosstalk is canceled as will be demonstrated in Section V.

III. LOW-SWING TRANSMITTERS

The energy-cost for a rising edge with swing V equals the well-known $E = CV^2$. Half of this energy is dissipated during charging. The other half is stored in the interconnect and dissipated at a later time when the interconnect is discharged (the resistance of the interconnect prevents efficient charge-recycling techniques). To reduce the link power it hence makes sense to reduce the swing. If only a single supply voltage is available and active circuits are used to reduce the swing, there is no quadratic but linearly relation with the swing ($E = CV_{DD}V_{swing}$). When a dedicated supply voltage is available to generate the low-swing signal, then the power is again quadratically dependent on the swing. Many low-swing techniques with a dedicated supply voltage (either generated on- or off-chip) for the transmitter have therefore been introduced in the past [4], [9], [21], [22].

The need for a dedicated supply voltage is a drawback, but the use of multiple supply grids becomes more accepted now that SoC-designs start to use multiple supplies (multiple voltage islands). SoCs use for example a high voltage (V_{DDH}) for the high performance (logic) parts and a slightly lower voltage (V_{DDL}) for the slower parts of the chip. Low-swing interconnect drivers can switch between these two supplies to generate the low-swing signal, with equal power efficiency as the dedicated supply variant, but without the need for yet another supply grid. An example schematic of such a low-swing transceiver is shown in Fig. 4.

This variant still has several drawbacks. A first drawback is the fact that the noise-margin is directly related to the amount of supply-noise and a short drop in one of the two supplies can easily introduce a bit-error. Tight coupling between the two supplies, to lower the differential ($V_{DDH} - V_{DDL}$) noise, could reduce this problem, but at the expense of area overhead for example for coupling capacitors. A second drawback, which is found in most low-swing transmitters, are the large transistors that are needed to drive the interconnects with sufficient speed. Driving these large transistors costs a lot of power and hence decreases the efficiency.

To circumvent these drawbacks and simultaneously increase the achievable data-rate, we propose to use capacitive pre-emphasis transmitters [12], [23]. The capacitive transmitter uses a series capacitance (C_{TX}) to drive the interconnect, as shown earlier in Fig. 1. This capacitance, together with the wire capacitance, acts as a capacitive divider which reduces the swing by a factor of $C_{TX}/(C_{wire} + C_{TX})$. The capacitive transmitter also





Fig. 5. Proposed low-swing capacitive pre-emphasis transmitter.

Fig. 6. Signals at 5 Gb/s for (a) the multiple V_{DD} transmitter and (b) the capacitive transmitter.

increases the bandwidth of the interconnect [12], [23], as C_{TX} emphasizes each transition with an overshoot. Compared to the low-swing transmitters that switch between supplies, the capacitive transmitter is much less sensitive to supply noise, as the capacitor divider also attenuates this noise. It does furthermore not require a special supply voltage and the lower theoretical efficiency ($E = CV_{swing}V_{DD}$) is more than compensated by the reduction in energy overhead at the driver side.

To illustrate these claims, the capacitive transmitter and the multiple- V_{dd} circuit were simulated and compared. The implementation of the capacitive transmitter that was used for the comparison is shown in Fig. 5. It uses a MOST as C_{TX} , as the high capacitance-density of the gate-oxide makes it very suitable as transmitter capacitance [12]. For the 2-mm interconnects, a MOST with $W = L = 2.7 \mu\text{m}$ gives a swing reduction to 10% of the supply voltage. A PMOST channel-capacitance is used with the gate connected to the driver to avoid loading the driver with the junction capacitances. An NMOST (current-source) at the Tx-side and a PMOST (resistive) load at the Rx-side define the low-frequency behavior and dc-operating point [12] and these are narrow and long transistors to minimize the static current.

Some signal waveforms of both circuits are shown in Fig. 6, which clearly illustrates the pre-emphasis effect of the capacitive transmitter. Numerical results are shown in Table I, which also includes the simulation results of the conventional full-swing transceiver from Fig. 2.

TABLE I
COMPARISON OF THE DIFFERENT TRANSMITTERS

	Conventional full-swing	Multi- V_{DD} low-swing	Capacitive low-swing
Technology	1.2V, 6 metal, 90nm CMOS		
Interconnects	2mm, $R_{wire}=400\Omega$		
	Shielded single ended $C_{wire}=480fF$	Twisted differential $C_{wire}=560fF$	
Supply	1.2V	$V_{DDH} = 1.2V$ $V_{DDL} = 1.08V$	1.2V
Voltage swing	1.2V	120mV	120mV
Driver size	$W_n=8\mu m$ $W_p=20\mu m$	$W_p=20\mu m$	$W_n=1.6\mu m$ $W_p=4\mu m$
Energy/trans			
Total	420fJ	135fJ	105fJ
Wire (theory)	346fJ	8fJ	80fJ
Tx overhead	74fJ	127fJ	25fJ
Static power	$\sim 9nW$ (leakage)	$\sim 10nW$ (leakage)	6 μW
Data-rate	fully shielded		
50% eye opening	5Gb/s	5Gb/s	9Gb/s
zero eye opening	9Gb/s	9Gb/s	12Gb/s
Delay (50%)			
Transmitter			
Nominal /Slow ($T=25^\circ C/T=100^\circ C$)	105ps/150ps	90ps/130ps	60ps/80ps
Interconnect	100ps	115ps	80ps

Both low-swing circuits have the same voltage swing and the driver sizes were chosen such that the circuits can reach 5 Gb/s with an eye-diagram that is at least 50% open. This means that a relatively large driver is needed for the multiple- V_{dd} circuit, which creates a significant overhead of 127 fJ/transition; 16 times more than the energy that is theoretically consumed. The capacitive transmitter has only 25 fJ overhead on top of its theoretical energy as the series capacitance reduces the capacitive load seen by the driver and hence enables a smaller driver-size.

In total, the capacitive transmitter is the most power-efficient (total of 105 fJ/transition). The smaller driver chain also has less delay and the pre-emphasis effect provides a higher achievable data-rate of 9 Gb/s with 50% vertical eye opening versus 5 Gb/s for the other two circuits. The conventional full-swing transmitter can only achieve this 5 Gb/s when every signal wire is fully shielded from any neighbors, to mitigate crosstalk. Compared to the conventional transmitter, the capacitive transmitter operates with four times lower power consumption, despite the fact that it uses two active wires per channel instead of one.

Table I also shows that the delay of the capacitive transmitter increases with 20 ps (33%) at the slow process corner and 100 °C temperature. The delay of the conventional alternatives increases by a larger margin of 42%/44%.

The swing (vertical eye-opening) of the capacitive transmitter is affected by process variations, mainly because the N- and PMOST that define the magnitude of the low-frequency transfer spread with respect to each other (the capacitance ratio C_{TX}/C_{wire} is more stable). This effect can reduce the swing in the worst-case corner to 95 mV. Compared to the other low-swing transmitter, which has to cope with supply variations that can easily amount to ± 100 mV, this is still quite stable behavior.

The transistors that define the low-frequency behavior in the capacitive transmitter also cause some static power consumption. However, the dynamic power easily dominates the static part for data-rates above 90 MHz (assuming random data). When the link is not used, it is easy to stop the static power consumption by setting both the D_{in} and the D_{in} -bar high, to break the current-path from the transmitter NMOSTs through the wire to the PMOST loads at the receiver.

When the link is in use, the receiver PMOSTs operate in triode and act as large resistances, connected to the (local) V_{DD} . Note that this configuration makes the capacitive transceiver well suited to cross (bridge) voltage domains, which can be an advantage in SoCs that operate with multiple voltage islands. This capability is both due to the differential nature and due to the fact that the dc operating point (common-mode voltage) is determined locally at the receiving end, which is good for robust operation of the sense amplifier. This in contrast to the multiple-supply transceiver which has its common-mode defined at the transmitting end.

The PMOST resistances are connected to the highest available reference: the (local) V_{DD} , which is not only simple, but is also beneficial for the channel-capacitance density of the C_{TX} -PMOST which is highest when it reaches strong inversion. Connecting the (PMOST) resistances to the supply does however require that the receiving sense amplifier is able to cope with an input common-mode voltage that is close to V_{DD} . A sense amplifier that tolerates these high common-mode voltages is discussed next.

IV. RECEIVER AND OPTIMAL SWING

In a low-swing transceiver, a latch-type sense amplifier—or in more general terms a clocked comparator—is a very suitable data receiver. A sense amplifier is not only a very fast circuit to regenerate the voltage to full swing, but it also samples the incoming data and realigns it to the clock.

In a recent paper [14], an improved version of a voltage latch-type sense amplifier was presented, which is fast and can operate over a wide common-mode and supply voltage range. Also, The offset of this “double-tail” sense amplifier is stable and does not increase significantly for high input common-mode levels, which is attractive for this application.

The schematic of the double-tail sense amplifier is shown in Fig. 7 together with its signal behavior. The operation is similar to a conventional latch-type voltage sense amplifier, apart from the fact that the input stage and the cross-coupled stage of this sense amplifier have a separate tail and are separated by a third, intermediate stage (M10 and M11). The circuit does need both a clock and a clock-not signal, but in case both complements are not available, a simple inverter can derive one from the other as their relative timing is not critical. To create static output signals, an SR-latch can be added at the output of the circuit or two sense amplifiers can be interleaved as shown in the next section. The clock-to-output delay of a single sense amplifier core is about 70 ps for 50-mV differential input voltage, but the sum of its setup and hold time is only 18 ps.

Offset is the bottleneck for the sense amplifier in this application (the measured rms noise is a factor five lower than the σ_{os}).

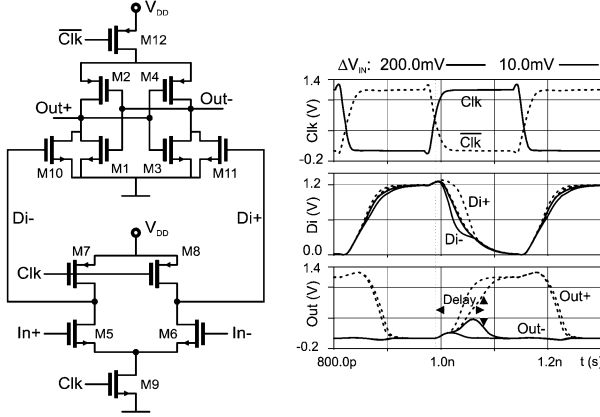


Fig. 7. Double-tail sense amplifier and its signals.

Therefore, the transistor dimensions of the double-tail sense amplifier are optimized relative to each other to get the lowest offset standard deviation (σ_{os}) per unit of power cost. Width scaling (or impedance or area scaling) can subsequently be applied to all the transistors together to match the offset standard deviation to the desired specification ($P \propto 1/\sigma_{os}^2$) [24] while maintaining the original speed characteristics.

The offset specification depends on the signal swing and the required yield and reliability. With a swing that equals for example six times the offset standard deviation ($6\sigma_{os}$), the chance that a sense-amplifier will introduce bit-errors due to its offset is only $\sim 2 \cdot 10^{-9}$. With Q being the cumulative Gaussian distribution function and Y being the yield-factor in terms of sigma (six in this case), this value is calculated as: $p_{\text{defect}} = 1 - \text{yield} = 2 \cdot Q(-Y)$. For the earlier introduced 25-tile NoC example with $2 \times 5 \times 4 \times 2 \times 64 = 5120$ sense amplifiers on a chip, the chance for offset related bit-errors is still only 10 ppm. A double-tail sense amplifier that has an offset standard deviation of 10 mV (according to 1000 Monte Carlo simulations) consumes about 90 fJ/bit. This sense amplifier can be scaled-down to get an offset of 20 mV, when a 6σ yield per sense amplifier is desired at 120-mV swing. The energy times offset-variance ($E\sigma^2$) remains constant, so the corresponding energy consumption will be $90/(20/10)^2 = 22.5$ fJ/bit.

The values for the swing and yield-factor above are not chosen randomly but actually define a power optimum, due to the tradeoff between transmitter and receiver power. The energy that is consumed in the transmitter, including the interconnect, has a more or less fixed overhead part and a part that is proportional to the swing

$$E_{\text{TX}} = p_{\text{activity}}(E_{\text{overhead}} + C_{\text{wire}}V_{\text{DD}}V_{\text{swing}}) \text{ J/bit} \quad (2)$$

where p_{activity} is the data activity (transition probability). The energy consumption of the sense amplifier is inversely proportional to the square of the offset and the required yield parameter Y relates offset to swing

$$E_{\text{RX}} = \frac{E\sigma^2}{\sigma_{\text{allowed}}^2} = \frac{E\sigma^2}{(V_{\text{swing}}/Y)^2} \text{ J/bit}. \quad (3)$$

With the substitution of $E\sigma^2 = 90 \text{ fJ} \cdot (10 \text{ mV})^2$, $Y = 6$, $p = 0.5$ (random bits), and the data from Table I, a graph can be




Fig. 8. Energy consumption versus swing.

plotted of these two equations and their sum, as shown in Fig. 8. The figure clearly emphasizes the advantage of low-swing signaling. At large signal swings, the lowered sense amplifier power can not compensate for the large increase in line power and full-swing signaling would cost over 5 times more power than signaling with the optimal swing. For the given parameters, this optimum is indeed about 120 mV (125 mV to be exact).

The optimum is also analytically solvable by taking the sum of E_{TX} and E_{RX} , differentiate, and solve for zero

$$\begin{aligned} \frac{d(E_{\text{TX}} + E_{\text{RX}})}{dV_{\text{swing}}} &= 0 \Leftrightarrow p_{\text{activity}}C_{\text{wire}}V_{\text{DD}} - 2\frac{E\sigma^2Y^2}{V_{\text{swing}}^3} \\ &= 0 \Rightarrow \\ V_{\text{swing,opt}} &= \sqrt[3]{2\frac{E\sigma^2Y^2}{p_{\text{activity}}C_{\text{wire}}V_{\text{DD}}}}. \end{aligned} \quad (4)$$

The equation shows that the optimum is only weakly dependent (with a third-order root) on properties such as C_{wire} and p_{activity} , so the optimum will not change much for different wire lengths or different data activities. We can make the reasonable assumption that the energy consumed in the sense amplifier is, at a given offset, quadratically proportional to the supply ($E\sigma^2 \propto V_{\text{DD}}^2$). Under that assumption, the optimal swing is proportional to the third-order root of the V_{DD} and a change in supply voltage will also have only a small influence on the optimum.

A change in technology has no influence on the optimum swing when we assume feature size (s) scaling with classical Dennard scaling rules [25]. First, the $C_{\text{wire}}/\text{length}$ does not change significantly over different technologies [1], but in a NoC, the size of the tiles and thereby the lengths of the wires are likely to scale, so $C_{\text{wire}} \propto s$. Second, V_{DD} (ideally) scales with s . Third, the energy scales with $E \propto C_{\text{MOST}}V_{\text{DD}}^2$ and with $C_{\text{MOST}} \propto \text{Area}^2/t_{\text{ox}} \propto s$ this becomes $E \propto s^3$. Fourth, the offset σ^2 scales with $1/s$ as $\sigma^2 \propto t_{\text{ox}}/\text{Area}^2$ [24]. Put altogether in (4), these four factors cancel each other out.

These observations are in line with the results from [20], where an optimum swing is calculated for the case when the receiver would be a ‘‘linear’’ amplifier instead of a latching sense

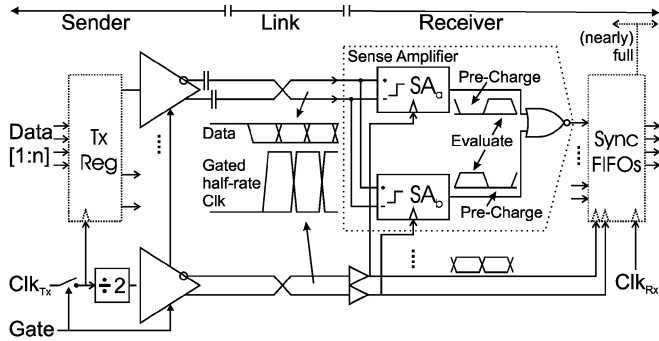


Fig. 9. Complete transceiver.

amplifier. Despite the use of a quite different calculation approach and a different technology, a similar optimal swing is found there.

At the optimal swing of 120 mV, the equations predict 53-fJ/bit energy consumption for the transmitter and interconnect and 22 fJ/bit for the sense amplifier. The actual sense amplifier circuit that is used in the complete transceiver is scaled for this optimum and consumes 24 fJ/bit. This is 10% more than predicted because the minimum width in the technology limits the downsizing of some transistors and because the actual sense amplifier consists of two interleaved instances which creates a slight power overhead of 1 fJ.

V. COMPLETE TRANSCEIVER

A. Transceiver With Synchronization

Section IV discussed the circuits for the data link, but did intentionally not yet mention how the clock is supplied to the receiver, as the data transceiver can operate with many different clocking-schemes, depending on the clocking strategy of the application (the SoC).

In a synchronous NoC, the receiver can simply be clocked with a local copy of the global clock, provided that the link latency does not exceed a clock period. In a completely asynchronous NoC without any clock signals, handshake signals can be used to provide the sense amplifier with a “clock.” But for most NoCs, the transceiver clocking strategy that is likely to be most suitable is a source-synchronous scheme in which the transmitter sends a copy of its local clock (or “stroke” or “sync” signal) alongside the data [4]–[6]. It is a simple and fast technique that is applicable to both synchronous, mesochronous, and GALS systems, as long as each router has a local clock available.

This option will be investigated further in this section and a schematic overview of a source-synchronous transceiver is shown in Fig. 9. At the left side the data words (flits) from the transmitting router enter the transceiver where they are optionally buffered in a transmitter register. The capacitive transmitters transmit the data over the link. Parallel to the data-bus, a gated half-rate clock is also transmitted (or in other words, data transfer is “double-pumped” or at “double-data rate”). The sense-amplifier at the receiver consists of two interleaved parts which act on the opposite edges of the clock to enable proper




Fig. 10. Cascade of direct forwarding transceivers.

sampling with a half-rate clock. Simple NOR-gates are used to combine the two outputs and create a static output signal.

The clock is transmitted at half-rate because a full-rate clock would be more heavily attenuated by the wire transfer. Full-swing drivers are used for the transmission of the clock to provide as much voltage-swing as possible. Attenuation of the clock can not be compensated by clocked sense amplifiers and conventional amplifiers (cascades of inverters) are used at the receiving end.

The clock is also gated to stop transmission when there is no data (e.g., in between packets). Both halves of the transmitter are also set high during absence of data, to eliminate static current as mentioned earlier. When the clock is stopped, both the halves of the differential clock signals will become low, to signal the receiver that there is no data. When this happens, both halves of the sense amplifier are also reset low, which enables automatic elimination of static current in following transceiver stages in case transceivers are cascaded, as discussed in the following.

B. Cascaded Transceivers

The synchronizing FIFO that is shown in Fig. 9 is normally present to realign the data with the local clock, and is often combined with queues to buffer the incoming data [8]. However, in certain router schemes, one can also omit the realignment at intermediate routers and directly forward the data to the next link, which can reduce the latency of the hops significantly. Direct forwarding—also known as wave-pipelining—can for example be useful in a circuit-switched network [26], where the crossbars that connect the links are pre-configured and there is no need to realign the data to the local router-clock at each hop, but only at the destination. Source-synchronous transceivers with direct-forwarding can also be interesting for more fine-grained systems that use static routing, such as field-programmable gate arrays (FPGAs).

To test the concept of direct-forwarding and its wave-pipelined clock, a number of transceivers are cascaded and simulated (omitting the switch fabric for simplicity), as shown in Fig. 10. Each transceiver in the chain resembles the schematic from Fig. 9, but without the synchronizing FIFO and with the interleaved sense amplifiers also performing the function of input register. Chains of inverters are used in the clock-path to drive the clock-interconnects. The number of inverters is chosen such that the delay of the clock-path is larger than the delay of the data path: $t_d(\text{Clk1 to Clk2}) > t_d(\text{Clk1 to SA}_{\text{Out}1}) +$




Fig. 11. Direct-forwarding transceiver signals at 5 Gb/s.

$t_d(\text{SA}_{\text{Out}1} \text{ to } \text{TX}_{\text{out}1}) + t_d(\text{TX}_{\text{out}1} \text{ to } \text{RX}_{\text{in}2}) + t_{\text{setup}}(\text{SA})$. The closer these two delays match, the shorter the latency will be, but at the cost of a reduced timing margin.

Some simulated time signals are shown in Fig. 11. As can be seen in the figure, the transmission and especially the startup of the clock is in this setup a speed-limiting factor, as the interconnects already cause quite some attenuation of the 2.5-GHz clock. At rates higher than 5 Gb/s/channel, the accumulation of clock disturbances over multiple stages prevents proper reception during the startup-transient. Simulations with clock-wires in a two times larger metal layer (such that they have four times lower resistance) showed that the entire system is capable to run at 9 Gb/s. The purpose of Fig. 11 is to show that even when the clock wires have to fit in the same area as a single-data channel, it is still possible to reach 5 Gb/s.

In the current setup, which uses moderately aggressive timing between data and clock, the latency is 300 ps for a single stage (independent of the data-rate), so it would cost 1500 ps to cross 10 mm of interconnect over five stages, which is only slightly larger than the latency of transceivers that use uninterrupted interconnects of 10 mm [10], [12].

As expected from earlier sections, the energy consumed in a single stage is 129 fJ/transition which amounts to 75 fJ/bit for random data ($24 \text{ fJ} + p_{\text{activity}} \times 105 \text{ fJ}$). In comparison, [4] needs 350 fJ/bit to cross 5 mm at 1.6 GHz, while 2.5 stages from this design can do it for 188 fJ/bit. The pseudo-differential low-swing transceiver from [9] needs 1.92 pJ/transition to drive a wire that has a capacitance of 1 pF, which corresponds to two stages from this design, which only needs 256 fJ/transition. The transceiver in [12] uses a similar data transceiver which is optimized to cross 10 mm of uninterrupted wire. Five stages from this design need 35% more energy per bit, but the multiple stages (clocked repeaters) enable a much higher data-rate (5 versus 2 Gb/s) and a higher yield (with respect to offset and PVT variations).

The power consumed in the clock is left out of the comparison above. In this design, the power needed for transmission of the forwarded clock is shared across all the data channels in the bus. The transmission of the clock consumes 1.3 pJ/transition when its inverter cascade is loaded by 64 sense amplifiers, which amounts to 20 fJ/bit/channel.




Fig. 12. Line output signals of three channels in a twisted bus.

The source-synchronous nature of this transceiver helps to make it resilient towards process spread. Simulations with the slow process corner at a temperature of 100 °C show an increase in delay of 65 ps per stage. At the fast process corner at -25 °C, the delay per stage is 45 ps lower than in the nominal situation. At both corners, the transceiver chain still operates correctly at 5 Gb/s as the change in clock-path delay is equal to the change in data-path delay within 5 ps.

The simulations described above were, for simplicity reasons carried out with only one data channel with simple one-dimensional lumped models for the interconnects. To test the effect of crosstalk, a simulation with a bus with twisted interconnects was also carried out. The interconnects are twisted as shown in Fig. 1 and a 2-D mesh of RC -lumps was used to model its behavior. Simulation results of the interconnect outputs of three neighboring channels are shown in Fig. 12. Hardly any crosstalk is visible in the outputs (compare to the single-ended bus signals in Fig. 3), which illustrates the effectiveness of the twists.

Because part of the wire capacitance is mutual between the wires in the bus, the common-mode transfer of the bus is different from that of a single wire. However, the dip in the common-mode that is visible in Fig. 12 is a startup transient that does not cause any difficulty for the sense amplifiers.

VI. CONCLUSION

In this paper, we have shown that the combination of a low-swing capacitive pre-emphasis transmitter, a bus with properly twisted differential wires, a double-tail sense amplifier and a source-synchronous clocking scheme is very suitable for communication in a NoC.

Compared to other low-swing transceivers, the capacitive transceiver does: 1) not need a second supply; 2) can operate at higher speeds; 3) has a higher power efficiency; and 4) has a better immunity to supply noise. The capacitively coupled transmitter also makes the transceiver suitable to cross different voltage domains.

The transceiver circuits are compatible with standard digital CMOS circuits and are easily scalable to future technologies. Analysis predicts that the power-optimal swing is about 120 mV, also in future technologies.

At this swing, the power consumption of the presented differential transmitter is four times lower than the power consumption of a conventional full-swing single-ended transmitter, while the obtainable data-rate is 80% higher. When we include the power of the sense amplifier and assume (optimistically) that a full-swing transmitter needs no dedicated receiver, then the presented transceiver is still a factor 3.3 more power efficient. For the 25-tile NoC example with 5-GHz clock and 25% average switching activity, this means that the total link power would drop down to 0.8 W, instead of the original 2.7 W.

With multiple transceiver stages cascaded in a wave-pipelined fashion, the transceiver can also compete with global-interconnect transceivers as it enables high data-rates (5 versus 3 Gb/s in [10] or 2 Gb/s in [12]) at a high reliability (6σ for random offset and correct operation over process and temperature corners) and with simple build-in synchronization. As such, the transceiver is also suitable for the long link distances that are for example found in networks with a torus or star topology.

ACKNOWLEDGMENT

The authors would like to thank Philips Research for chip fabrication, P. Wolkotte, G. Smit and the STW user committee for helpful discussions. They would also like to thank G. Wienk and H. de Vries for their technical assistance.

REFERENCES

- [1] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proc. IEEE*, vol. 89, no. 4, pp. 490–504, Apr. 2001.
- [2] L. Benini and G. De Micheli, "Networks on chips: A new SoC paradigm," *IEEE Computer*, vol. 35, no. 1, pp. 70–78, Jan. 2002.
- [3] W. J. Dally and B. Towles, "Route packets, not wires: On-chip interconnection networks," in *Proc. 38th Des. Autom. Conf.*, Jun. 2001, pp. 684–689.
- [4] K. Lee, S.-J. Lee, S.-E. Kim, H.-M. Choi, D. Kim, S. Kim, M.-W. Lee, and H.-J. Yoo, "A 51 mW 1.6 GHz on-chip network for low-power heterogeneous SoC platform," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2004, pp. 152–153.
- [5] S.-J. Lee, K. Lee, S.-J. Song, and H.-J. Yoo, "Packet-switched on-chip interconnection network for system-on-chip applications," *IEEE Trans. Circuits Syst. II, Express Briefs*, vol. 52, no. 6, pp. 308–312, Jun. 2005.
- [6] S. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, P. Iyer, A. Singh, T. Jacob, S. Jain, S. Venkataraman, Y. Hoskote, and N. Borkar, "An 80-tile 1.28TFLOPS network-on-chip in 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 98–99.
- [7] D. Lattard, E. Beigne, C. Bernard, C. Bour, F. Clermidy, Y. Durand, J. Durupt, D. Varreau, P. Vivet, P. Penard, A. Bouttier, and F. Berens, "A telecom baseband circuit based on an asynchronous network-on-chip," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 258–259.
- [8] S. Vangal, A. Singh, J. Howard, S. Dighe, N. Borkar, and A. Alvandpour, "A 5.1 GHz 0.34 mm² router for network-on-chip applications," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2007, pp. 42–43.
- [9] H. Zhang, V. George, and J. M. Rabaey, "Low-swing on-chip signaling techniques: Effectiveness and robustness," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 8, no. 3, pp. 264–272, Jun. 2000.
- [10] 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, no. 1, pp. 297–306, Jan. 2006.
- [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 Electron. Des. (ISLPED)*, Aug. 2005, pp. 186–191.
- [12] 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 ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 414–415.
- [13] E. Mensink, D. Schinkel, E. A. M. Klumperink, E. Van Tuijl, and B. Nauta, "Optimal positions of twists in global on-chip differential interconnects," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 15, no. 4, pp. 438–446, Apr. 2007.
- [14] 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 ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 314–315.
- [15] D. Pamunuwa, L. R. Zheng, and H. Tenhunen, "Maximizing throughput over parallel wire structures in the deep submicrometer regime," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 2, pp. 224–243, Apr. 2003.
- [16] H. Shah, P. Shiu, B. Bell, M. Aldredge, N. Sopory, and J. Davis, "Repeater insertion and wire sizing optimization for throughput-centric VLSI global interconnects," in *Proc. Int. Conf. Comput.-Aided Des.*, Nov. 2002, pp. 280–284.
- [17] H. Bakoglu, *Circuits, Interconnections and Packaging for VLSI*. Reading, MA: Addison-Wesley, 1990.
- [18] E. Mensink, "High-speed global on-chip interconnects and transceivers," Ph.D. dissertation, IC Design Group, Univ. Twente, Enschede, The Netherlands, 2007.
- [19] A. Morgenshtein, I. Cidon, A. Kolodny, and R. Ginosar, "Comparative analysis of serial and parallel links in networks-on-chip," in *Proc. SoC Conf.*, Nov. 2004, pp. 185–188.
- [20] C. Svensson, "Optimum voltage swing on on-chip and off-chip interconnect," *IEEE J. Solid-State Circuits*, vol. 36, no. 7, pp. 1108–1112, Jul. 2001.
- [21] R. Ho, K. Mai, and M. Horowitz, "Efficient on-chip global interconnects," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2003, pp. 271–274.
- [22] F. Worm, P. lenne, P. Thiran, and G. De Micheli, "A robust self-calibrating transmission scheme for on-chip networks," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 13, no. 1, pp. 126–139, Jan. 2005.
- [23] R. Ho, I. Ono, F. Liu, R. Hopkins, A. Chow, J. Schauer, and R. Drost, "High-speed and low-energy capacitively-driven on-chip wires," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 412–413.
- [24] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.
- [25] R. H. Dennard, F. H. Gaensslen, V. L. Rideout, E. Bassous, and A. R. LeBlanc, "Design of ion-implanted MOSFET's with very small physical dimensions," *IEEE J. Solid-State Circuits*, vol. 9, no. 5, pp. 256–268, Oct. 1974.
- [26] P. T. Wolkotte, G. J. M. Smit, G. K. Rauwerda, and L. T. Smit, "An energy-efficient reconfigurable circuit-switched network-on-chip," in *Proc. IEEE Int. Symp. Parallel Distrib. Process.*, Apr. 2005, pp. 155a–155a.



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, Enschede, the Netherlands, in 2003.

From 2003 to 2007, he worked as a Ph.D. student at the University of Twente in 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 about high-speed on-chip communication. 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 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 16 papers.



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) and the Ph.D. degree in high-speed on-chip communication from the University of Twente, Enschede, the Netherlands, in 2003 and 2007, respectively.

He is currently an ASIC Design Engineer with Bruco B.V., Borne, The Netherlands.



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), Enschede, The Netherlands, 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 Ph.D., Eric 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 authored and coauthored more than 80 journal and conference papers. In 2006 and 2007, he served as Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—II: EXPRESS BRIEFS, and since 2008 for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—I: REGULAR PAPERS.

Dr. Klumperink was a corecipient of the ISSCC 2002 “Van Vessel 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. He received the M.Sc. degree (*cum laude*) in electrical engineering and the Ph.D. degree in analog CMOS filters for very high frequencies from the University of Twente, Enschede, The Netherlands, in 1987 and 1991, respectively.

In 1991, he joined the Mixed-Signal Circuits and Systems Department, 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 a Full Professor heading the IC Design Group, which is part of the CTIT Research Institute. He is also part-time consultant in industry and in 2001 he cofounded Chip Design Works. His current research interest is high-speed analog CMOS circuits.

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 the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—II: ANALOG AND DIGITAL SIGNAL PROCESSING. After this, he served as Guest Editor, Associate Editor (2001–2006)—and from 2007 as Editor-in-Chief for the IEEE JOURNAL OF SOLID-STATE CIRCUITS. He is also member of the technical program committees of the International Solid State Circuits Conference (ISSCC), the European Solid State Circuit Conference (ESSCIRC), and the Symposium on VLSI circuits. He is a corecipient of the ISSCC 2002 “Van Vessel Outstanding Paper Award,” is distinguished lecturer of the IEEE, and elected member of IEEE-SSCS AdCom.