

Research Article

High-Performance Long NoC Link Using Delay-Insensitive Current-Mode Signaling

Ethiopia Nigussie,¹ Teijo Lehtonen,^{1,2} Sampo Tuuna,¹ Juha Plosila,^{1,3} and Jouni Isoaho¹

¹ Department of Information Technology, University of Turku, 20014 Turku, Finland

² Turku Centre for Computer Science (TUUS), 20520 Turku, Finland

³ Research Council for Natural Sciences and Engineering, Academy of Finland, 00501 Helsinki, Finland

Received 1 November 2006; Revised 24 January 2007; Accepted 1 March 2007

Recommended by Maurizio Palesi

High-performance long-range NoC link enables efficient implementation of network-on-chip topologies which inherently require high-performance long-distance point-to-point communication such as torus and fat-tree structures. In addition, the performance of other topologies, such as mesh, can be improved by using high-performance link between few selected remote nodes. We presented novel implementation of high-performance long-range NoC link based on multilevel current-mode signaling and delay-insensitive two-phase 1-of-4 encoding. Current-mode signaling reduces the communication latency of long wires significantly compared to voltage-mode signaling, making it possible to achieve high throughput without pipelining and/or using repeaters. The performance of the proposed multilevel current-mode interconnect is analyzed and compared with two reference voltage mode interconnects. These two reference interconnects are designed using two-phase 1-of-4 encoded voltage-mode signaling, one with pipeline stages and the other using optimal repeater insertion. The proposed multilevel current-mode interconnect achieves higher throughput and lower latency than the two reference interconnects. Its throughput at 8 mm wire length is 1.222 GWord/s which is 1.58 and 1.89 times higher than the pipelined and optimal repeater insertion interconnects, respectively. Furthermore, its power consumption is less than the optimal repeater insertion voltage-mode interconnect, at 10 mm wire length its power consumption is 0.75 mW while the reference repeater insertion interconnect is 1.066 mW. The effect of crosstalk is analyzed using four-bit parallel data transfer with the best-case and worst-case switching patterns and a transmission line model which has both capacitive coupling and inductive coupling.

Copyright © 2007 Ethiopia Nigussie et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Network-on-Chip (NoC) is the most viable solution for on-chip communication that provides good scalability and enables gigascale integration in single-chip systems. One of the basic reasons for good scalability is that the length of connections is held constant and the signaling is kept local, from one router to another with the maximum distance of few millimeters. However, when the chip size increases, the latency for messages traversing from a processing unit far away from another becomes large. This is either due to the lack of fast paths between remotely situated nodes or due to the type of topology which has long channels. For example, in regular mesh structure to send a data packet between remotely located nodes, the message has to traverse many hops which increases the probability of a message to be blocked. This leads to unpredictable message latencies and difficulty in

achieving guaranteed service operations. In [1], it is showed that using a few additional long-range links in a mesh network reduces the average packet latency significantly and improves the achievable throughput substantially. From topological point of view, the end-around channels in torus network are long which results in excessive latency. This problem can be avoided by folding a torus network. However, folding the torus eliminates the long end-around channels at the expense of doubling the length of the other channels [2] and increasing the layout complexity. Thus, it is preferable to use torus without folding if the long end-around channels can be implemented using high-performance signaling techniques. Both of these cases, the mesh structure with additional long links and the torus structure, require the use of long high-performance NoC links that pass over more than one processing element, thus the length of channels is 4 mm or more [3]. The structures of mesh network

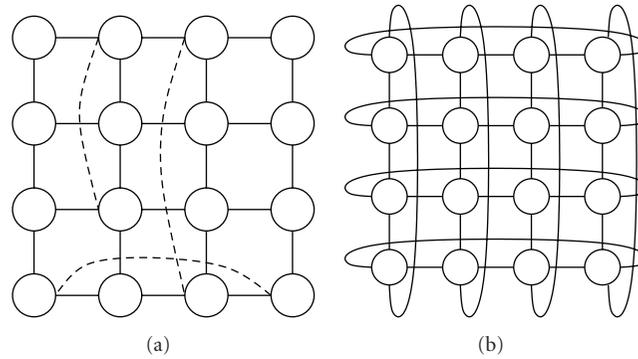


FIGURE 1: NoC architectures using long links. (a) 4×4 mesh with added long links. (b) 4-ary 2-cube torus.

with additional long distance links and torus network are illustrated in Figure 1.

The physical performance of long wires suffers greatly under technology scaling because the length is not scaled instead even longer wires are needed due to the increase of on-chip size. This makes long-range on-chip communication increasingly expensive [4]. The higher wire resistance, increased length, and decreased wire spacing cause the wire delay to increase considerably compared to the gate delay. In order to control this increase, designers scale down the wire cross-sectional area at a slower rate which prevents the dramatical increase of wire resistance. This ongoing trend of controlling the RC delay, combined with the faster rise/fall times and longer wires, results in a situation where the inductive part of the wire impedance can no longer be ignored. Thus, in addition to the capacitive coupling, the inductive coupling also causes crosstalk noise which creates more signal integrity problems.

Furthermore, the impact of process, supply voltage, and temperature variations on the performance and reliability of long on-chip links is expected to increase as technology scales down [5]. These variations cause the signal propagation delay through interconnects to be uncertain which in turn affects the performance and reliability of the system significantly. Moreover, the power dissipation due to global interconnect is increasing compared to the power consumption of the logic.

In order to achieve high-performance on-chip communication, it is necessary to implement efficient signaling technique. Current-mode signaling is faster and has lower dynamic power consumption than voltage-mode signaling. It is also immune to power supply noise and has reduced sensitivity to process-induced variations. Due to these advantages, we use current-mode signaling for the implementation of high-performance long-range NoC links. The delay variations problem can be tackled by using delay-insensitive communication. In this work we combine self-timed 1-of-4 encoded communication protocol with current-mode signaling for achieving a high-performance delay-variation-insensitive long-range on-chip communication.

The paper is organized as follows. We first discuss the principles of self-timed communication and present the delay-insensitive 1-of-4 encoding in Section 2. In Section 3 we discuss the advantages of current-mode signaling compared to voltage mode. In Section 4, brief discussion about multilevel current-mode signaling and its usage in our interconnect design is presented. The implementation of the signaling circuitry for self-timed 2-phase 1-of-4 encoded multilevel current-mode signaling is presented in Section 5 together with the implementations of the two reference 2-phase 1-of-4 encoded voltage-mode signaling circuits. The first reference uses pipelining and the other one uses optimal repeater insertion. In Section 6, first, the wire model used during simulations is presented followed by analysis of the presented current-mode signaling and the reference voltage-mode signaling techniques in latency, throughput, power consumption, and noise tolerance. Section 7 contains discussion about the results and future work, and finally conclusions are presented in Section 8.

2. SELF-TIMED COMMUNICATION

A NoC system consists of many processing blocks which have different timing requirements and can operate at different clock frequencies. Communication between these blocks needs synchronization which is error-prone. Also the clock distribution over a wide chip with low skew and jitter is problematic. A viable solution for this is the use of the globally asynchronous locally synchronous (GALS) design approach, where communication between processing blocks is done asynchronously. Therefore, we base our link on self-timed design principles.

The choice of the handshake protocol affects the throughput of a communication link. The two-phase protocol is often preferred instead of four-phase protocol for long on-chip interconnects to avoid the usage of a time-consuming spacer (return-to-zero phase) between two consecutive data symbols [6]. The use of two-phase protocol also minimizes power consumption since there is less transitions in the control wires.

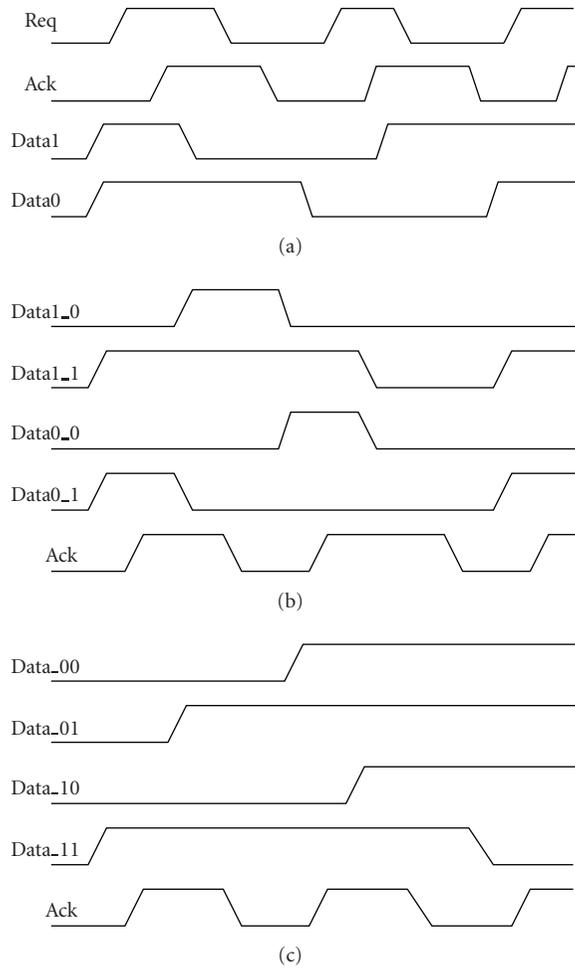


FIGURE 2: Self-timed communication. (a) 2-phase bundled-data (transmitting data “1101001011”). (b) 2-phase dual-rail (transmitting data “1101001011”). (c) 2-phase 1-of-4 (transmitting data “1101001011”).

The communication can be carried out using control wires separately of the data. In this *bundled-data* approach, it is assumed that by the time request arrives, the data have already arrived. 2-phase bundled-data signaling is presented in Figure 2(a). To get rid off the timing constraints, the data validity indicator signal can be included in the data resulting in delay-insensitive communication. The delay-insensitive handshake protocol in which the data validity is transmitted implicitly operates correctly regardless of the delay in the interconnecting wires. The simplest one of the delay-insensitive protocols is the *dual-rail* protocol, which is demonstrated in Figure 2(b). In dual-rail, there are two wires for each bit, one for *zero* and the other for *one*. Either one of these signals is toggled and so at the receiver it can be noticed when all the bits have arrived regardless of their different delays.

In 1-of-4 data encoding, a group of four wires is used to transmit two bits of information per symbol. A symbol is one of the two-bit codes 00, 01, 10, or 11 and it is transmitted through activity on one of the four wires. Since it is

possible to detect the arrival of each symbol at the receiver, 1-of-4 encoding is delay-insensitive, as are all the 1-of- N codes [7]. The 1-of-4 signaling is illustrated in Figure 2(c). Delay-insensitive data communication is a viable method to realize robust on-chip interconnects in future nanoscale technologies in which significant signal propagation delay variations are unavoidable. These delay variations occur due to different reasons, for example, due to crosstalk, temperature, supply voltage, and process variations. Besides being delay-insensitive, 1-of-4 encoding has more immunity against crosstalk effects as compared to single-rail (bundled-data) encoding, because the likelihood of two adjacent wires switching at the same time is much smaller. Furthermore, dynamic power consumption due to wire capacitance is smaller for the 1-of-4 code than for the simpler 1-of-2 (dual-rail) code. This is because the 1-of-4 code conveys two bits of information using only a single transition, while the 1-of-2 code requires two transitions for two bits of information. This effect can be seen in Figure 2. Considering these advantages, 2-phase 1-of-4 encoding is used in the proposed multilevel current-mode interconnect.

3. CURRENT-MODE SIGNALING

The signal transmission systems used in CMOS circuits can be broadly classified into two categories: voltage-mode and current-mode signaling. The important difference between the two transmission systems lies in the type of signal that is forced on the transmission medium. While voltage mode uses voltage as signal, current mode uses current. In voltage mode, the voltage has to swing from rail to rail over the entire length of the wire. This leads to large transient currents consuming more power, larger delay, and it also generates power-supply noise [8]. The optimal repeater insertion technique [9] used in voltage-mode signaling was developed to reduce the wire delay and improve the performance of global interconnections. However, with the increase in number and density of interconnects with technology scaling, the number of repeaters necessary would increase considerably, presenting significant overhead in terms of power and area.

The key to current-mode signal transporting is the low-impedance termination at the receiver which results in reduced signal swings without the need of separate voltage references and increased bandwidth performance. Also this low-impedance termination shifts the dominant pole of the system and leads to a smaller time constant and thus, to a smaller delay. It can operate at a much lower noise margin than the voltage-mode network, and at a much lower swing as well due to its immunity to power supply noise. All these translate into increased bandwidth performance [10], decreased delay and dynamic power dissipation and higher noise immunity. For these reasons, current-mode signaling technique becomes a better alternative than voltage mode for contemporary and future high-speed noise-prone single-chip systems.

Current-mode signaling has already been proven to provide drastic speed enhancements for on-chip signaling [11–13]. It is also shown theoretically in [11] that current-mode

signaling can be three times faster than voltage-mode signaling.

There are three primary sources of power dissipation in current-mode circuits: static, dynamic, and short-circuit power dissipation. In current-mode signaling, static power dissipation is the major component of the total power dissipation that arises from the constant current path from V_{dd} to ground via the termination. Static power dissipation can be minimized using different circuit techniques which reduce leakage currents. Dynamic power is dissipated when the parasitic capacitance of the wire is charged and discharged. Since current-mode signaling operates at low-voltage swing, dynamic power consumption is not as significant source of power dissipation like in voltage-mode signaling. The third source of power dissipation arises from the finite input signal edge rates that result in short-circuit current. Generally, careful control of input edge rates can minimize the short-circuit current component to within 20% of the total dynamic power dissipation [14].

The other important feature of current-mode signaling is its reduced delay sensitivity due to process-induced variations [15]. Inspired by the advantages explained above, we investigate here the use of current-mode signaling for implementing high-performance delay-insensitive links for NoC long-range communication.

4. MULTILEVEL CURRENT-MODE SIGNALING

In delay-insensitive transmission, the data validity indicator is the transmitted data itself. Due to this, the transmission of every new data needs to be seen in the transmitting wire usually in the form of voltage level or transition depending on the type of handshake protocols. Using transition in current-mode signaling may cause unnecessary power consumption due to the constant current flow in some of the wires which have been previously made a transition to high state. In order to save this power, the presented current-mode interconnect allows current flow in the wires only during the respective symbol transmission. In this power-saving transmission scheme, it is not possible to see the arrival of new data during consecutive same symbol transmission using binary current-mode signaling. Due to this, three current levels are required in the proposed current-mode interconnect, two nonzero current levels to differentiate between consecutive same symbol transmissions and the third current level (zero current) to indicate the wire is idle, that is, there is no data transmission through that wire. The transmitted multilevel current is first detected at the receiver by a detecting circuit based on a current comparator. Then, the encoded voltages are estimated using decoding circuitry.

Multilevel current-mode signaling has been demonstrated to be robust and power-efficient in interchip signaling [16, 17]. In addition, using an analogy between digital communication over a band-limited channel and on-chip signaling, it is shown that for a given bit error rate and data rate, four-level current-mode signaling is the most power-efficient compared to binary voltage and current-mode signaling [18]. In this type of signaling the acceptable number of

current levels and the step size between them are limited by the noise margin. In [19], it is shown that radix-8 full-adder design using eight current levels and 10 μ A step size gets large enough noise margin.

Due to mismatch, parameter variations, noise, and other nonidealities, the current levels at the receiver input may deviate from the one predefined in the driver. This may lead to decoding error if the steps between different current levels are not enough and the current comparator has low noise margin. In addition, it is necessary to decode out the data in voltage form as fast as possible to fulfill the requirement of high data transmission rate. The key to achieve these is a large current comparator gain which provides sharp transition and greater noise margins which can accommodate all current levels. Lower threshold current values will increase the gain at the expense of greater comparator delay times. The usual approach to counteract the delay penalty is scaling the input current of the comparator lower than the input using current mirror division and then comparing these scaled input currents to reference currents.

Some fast and robust on-chip links based on multilevel current-mode signaling have been proposed [20–22]. In this paper, we present a high-performance interconnect which uses 1-of-4 data encoding and three distinct current levels per data wire. This interconnect has superiority and uses different approaches compared to [20–22]. In [20], 2-color 1-phase dual-rail encoding using four current levels is presented. As stated in Section 2, the dynamic power consumption due to wire capacitance of dual-rail is larger than 1-of-4 encoding since it requires two transitions rather than one to transmit two-bit data. In addition, it is more susceptible to crosstalk effects because there is a larger probability of adjacent wires switching at the same time than that of 1-of-4 encoding. Furthermore, using three current levels instead of four allows our proposed interconnect to have a larger noise margin than in [20]. The presented multilevel current-mode interconnect is also superior to [21] in terms of performance, power consumption, crosstalk, and noise margin. The interconnect in [21] is designed using 2-phase single-rail encoding with delay-insensitive feature and supporting simultaneous bidirectional data transmission. Although this approach decreases the required number of wires by half, it makes the signaling circuitry more complex, dissipates much more power, and has a significant decrease in signal transmission speed. Also due to the requirement of 7 current levels per wire, the noise margin of this interconnect decreases considerably compared to the presented three current-level interconnect. Moreover, the proposed three-current level interconnect has more immunity to crosstalk than [21] since it uses 1-of-4 encoding rather than single-rail. The interconnect proposed in [22] is designed using synchronous approach and allows to transmit two-bit data per wire. Since the reported performance and power consumption result is using 100 nm technology it is difficult to compare with our interconnect results. However, asynchronous interconnect has many advantages over synchronous one especially in the nanotechnology design of on-chip interconnect. The most relevant ones are the avoidance of clock and

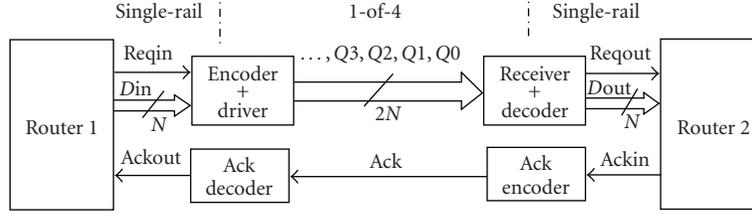


FIGURE 3: Conversion of single-rail to 1-of-4 and back to single-rail encoding.

clock-related problems and allowing delay-insensitive data transfer.

5. IMPLEMENTATIONS

In the subsequent sections, we present two different on-chip link implementations based on 1-of-4 data encoding. Both of them use two-phase protocol, the difference being that one is implemented using voltage-mode signaling and the other using multilevel current-mode signaling. The most common data encoding in GALS design is single-rail (bundled-data) encoding which uses N wires to transfer N -bit information and two additional handshake wires indicating data validity and acceptance. Since this encoding has a timing constraint between control (data validity) and data wires, communication through long on-chip interconnect becomes sensitive to delay variations. Therefore, converting single-rail encoding to delay-variation-insensitive encoding is mandatory for long on-chip communication where delay variations are unavoidable. The general block diagram of the considered signaling system is shown in Figure 3. We assume that the communicating parties, routers 1 and 2, have voltage-mode bundled-data (i.e., single-rail encoded) interfaces. The bundled-data protocol is then converted into the appropriate delay-insensitive 1-of-4 protocol and back to bundled-data protocol by the encoder/decoder units attached to the routers 1 and 2.

5.1. Two-phase 1-of-4 encoded voltage-mode interconnect (TPVm)

In the TPVm scheme, which serves as the reference for the current-mode implementation, one of the four wires makes transition to indicate the presence of a new two-bit symbol. When this new symbol arrives to the receiving module, the receiver accepts the symbol and sends an acknowledgment to the sender module by changing the state of the acknowledge signal. Since voltage-mode signaling is used, the voltage on the interconnect swings from rail to rail over its entire length. This leads to large dynamic power consumption, large delay, and generation of power supply noise. The usual approach to improve the performance of a voltage-mode interconnect is to insert repeaters or pipeline latches. Inserting repeaters decreases the signal propagation delay at

the cost of increasing power consumption and chip area. A higher throughput can be obtained by using pipeline latches instead of repeaters to both amplify the signal and spread the link delay over multiple pipeline stages. This further increases power consumption and area costs compared to the simple repeater approach. We consider here both schemes for the reference voltage-mode 1-of-4 encoded interconnect. The pipelined and repeater-based implementations are called TPVmP and TPVmRep, respectively. In the TPVmP implementation pipeline stages are inserted in every 2 mm along the link wire. This is based on the assumption that the typical distance between two neighbouring (adjacent) routers in the mesh structure is 2 mm [3] and that the local link length can be considered an upper limit for pipeline-free signal transmission [1]. In TPVmRep implementation optimal repeater insertions are used for both data and acknowledgment transmissions. The required optimal number of repeaters and optimal size of the repeater are calculated using [23, equation (36)]. Using this equation, the required number of optimal repeaters becomes $2.22 \cdot L$ and the optimum size of the repeater becomes $76.5 \cdot \text{minimum size inverter}$, where L is the wire length.

The straightforward gate level implementations of the encoder which converts the two-phase single-rail input to the delay-insensitive two-phase 1-of-4 protocol, the pipeline stage, and the decoder and completion detector which converts the delay-insensitive code back to the two-phase single-rail form at the receiver side are shown in Figure 4. The encoder consists of NOR gates which generate the select inputs for the multiplexers depending on the two-bit input codes, double-edge triggered flip-flops which are used to sample the symbol value at both edges of the request signal, and multiplexers each of which allows transition on the corresponding flip-flop output only when the appropriate input symbol is present. The decoder and completion detector circuit consists of XNOR gates which detect the transitions on the wires, NAND gates and an SR latch to decode the data back into the single-rail form, and a four-input XOR gate together with an $N/2$ -input C-element for detecting completion. A C-element is a basic building block of self-timed logic. It is a state-holding element, a special kind of latch. When all of its inputs are 0 or 1, the output is set to 0 or 1, respectively. For other input combinations, it preserves its state. Its truth table is shown in Table 1 where t and $t-1$ indicate the current and previous values, respectively, and $-$ indicate, do not care.

TABLE 1: The truth tables of 2- and 3-input C-elements.

| $a_{0,t}$ | $a_{1,t}$ | c_t | $a_{0,t}$ | $a_{1,t}$ | $a_{2,t}$ | c_t |
|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 1 | c_{t-1} | 0 | 0 | 1 | c_{t-1} |
| 1 | 0 | c_{t-1} | 0 | 1 | – | c_{t-1} |
| 1 | 1 | 1 | 1 | 0 | – | c_{t-1} |
| | | | 1 | 1 | 0 | c_{t-1} |
| | | | 1 | 1 | 1 | 1 |

An inverter is used as both driver and receiver for the transmission of the two-phase acknowledgment signal between the pipeline stages in the TPVMP implementation.

5.2. Pulsed 1-of-4 encoded multilevel current-mode interconnect (PMCm)

The PMCm scheme converts two-phase single-rail voltage-mode signaling into pulsed 1-of-4 multilevel current-mode signaling at the transmitter side. At the receiver side, delay-insensitive current-mode signaling is turned back into single-rail voltage-mode communication. The PMCm scheme is logically equivalent to the TPVm scheme described above, but now information is presented as current rather than voltage transitions. Hence, one of the four data wires draws current to indicate the presence of a new two-bit data symbol. Similarly, an acknowledgement is signaled as current on the acknowledgement wire. As explained in section 3, such current-mode implementation is inherently much faster and more immune against power supply noise and delay variations compared to the voltage-mode implementation. The communication protocol is shown in Figure 5 (from the receiver’s perspective) and the signaling circuits are depicted in Figures 6 and 7. The advantage of this link implementation is that high throughput and low latency can be achieved without using pipelining or repeaters.

The multilevel and pulsed nature of the PMCm scheme can be seen in Figure 5. The current detected at the receiver has three different values: 0, I , and $2I$. The values I and $2I$ are used when the voltage-mode request signal $Reqin$ at the transmitter side is low and high, respectively, reflecting the adopted two-phase communication protocol. The value 0, in turn, means that there is no symbol on a wire. It is used as the initial value of the data wires and for switching off current on a wire when the 2-bit symbol to be transmitted changes, making current on a wire pulse shaped. This feature reduces the overall power consumption of the current-mode interconnect. The values of I and $2I$ are determined by considering the speed, power consumption, and noise margin of the interconnect. In the following consecutive sections, the implementations of the encoder, decoder, and completion detector are separately discussed.

5.2.1. Encoder and driver

The encoder takes the request and two data bits in the voltage-mode single-rail form and converts this information

into multilevel current-mode 1-of-4 signaling. The double-edge triggered flip-flops shown in Figure 6 are used to sample the value of the 2-bit data symbol at each transition of the two-phase request signal $Reqin$. For instance, consider the encoder circuit of the wire $Q3$. Depending on the value of the signal $Reqin$, either transistor $Mn1$ or $Mn2$ conducts making either current I or $2I$ to flow through the wire $Q3$ when the symbol “11” has arrived from the sender module. To prevent the line from drawing current continuously, the transistor $Mn4$ is used to ground the line when other than the symbol “11” is sent. The reset signal rst is controlled by the transmitting module. When a data burst is about to begin, rst is set to high enabling the sampling flip-flops. When the burst has been completed, rst is initialized back to low, meaning that all the data wires become grounded. This is necessary to prevent data wires of the link from drawing current (consuming power) during possibly long idle periods between bursts. In nanometer technology, where NoC is one of the promising candidates, process variation effects are one of the major concerns. Due to process variation effects, the driver output currents may vary from their expected values. In order to minimize this variation, transistors $Mp1$ and $Mp2$ which operate in the linear region form resistive path from the supply voltage to $Mn1$ and $Mn2$ which in turn keeps the switching threshold of $Mn1$ and $Mn2$ transistors constant.

5.2.2. Receiver and current comparator

At the receiver side, consider the current comparator circuit of $Q3$, as depicted in Figure 7. It is composed of the diode-connected input NMOS transistor $Mn2$, the NMOS transistors $Mn3$ and $Mn4$ connected to replicate this input current, the reference or threshold current generating pair of transistors $Mn1$ and $Mp1$, and the PMOS transistors $Mp2$ and $Mp3$ that replicate the threshold current. In addition to serving as an input transistor, $Mn2$ acts also as a termination load. The drains of the PMOS reference current replicating transistors and line current replicating NMOS transistors are connected together to generate the comparator circuit’s output voltages. This comparator provides a logical high output voltage when its input current $I(Q3)$ is less than the threshold current and a logical low output voltage when the input current $I(Q3)$ is greater than the threshold current. Here the current comparator compares current on the wire $Q3$ with two different threshold currents, $0.5I$ and $1.5I$, in order to distinguish the three current levels. To be more specific, if $I(Q3) < 0.5I$, both comparator outputs $V(30)$ and $V(31)$ are high (initial state). If $0.5I < I(Q3) < 1.5I$, $V(30)$ and $V(31)$ are low and high, respectively. If $I(Q3) > 1.5I$, both $V(30)$ and $V(31)$ are low.

In nanometer technology, the line current at the input of the receiver may vary from the nominal value due to crosstalk, process variation effects, and other noise sources. However, this does not affect the reliability as long as the current levels are within the specified margins. Since there are only three current levels, it is easily possible to meet the required noise margins at minimal power consumption cost. In addition, the reference current may also vary from its nominal value due to process variations. This affects only

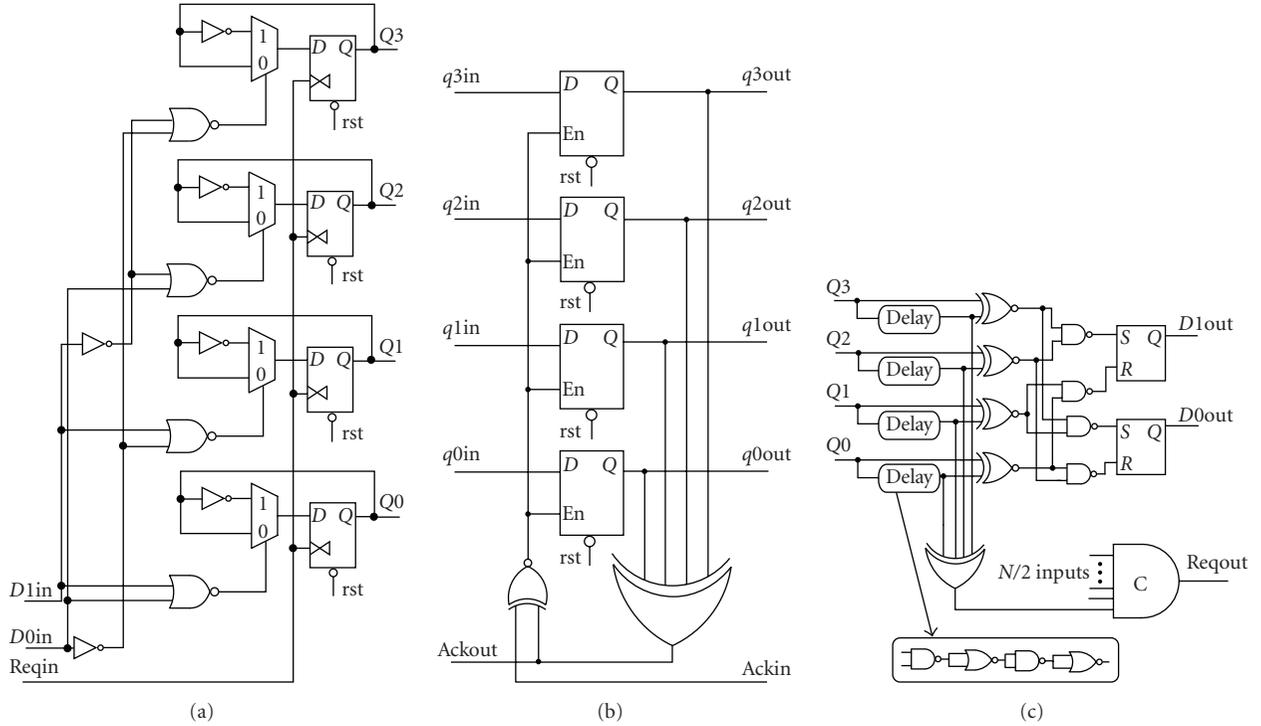


FIGURE 4: The reference 2-phase 1-of-4 encoded voltage-mode interconnect components. (a) Encoder. (b) Pipeline stage. (c) Decoder and completion detector.

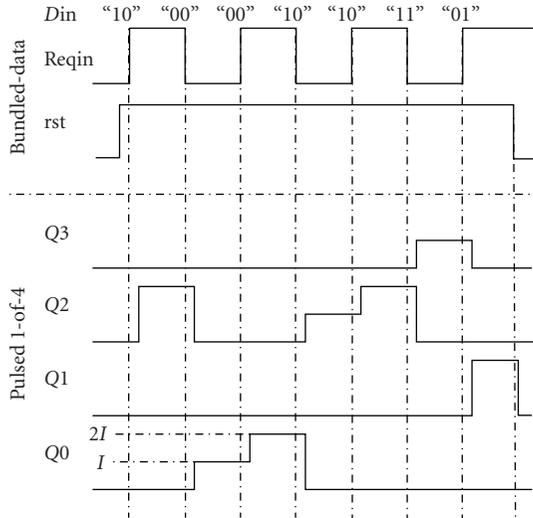


FIGURE 5: Communication protocol of 1-of-4 encoding in pulsed multilevel current-mode interconnect.

the speed but not the reliability of the communication since delay-insensitive data transfer mechanism is used. For example, if the reference current decreases from its nominal value, the comparison takes place ahead. This shifts the data output point as well as the data validity indicator to the left. Thus there is no threat to the reliability of the communication.

5.2.3. Decoder and completion detector

As shown in Figure 7, the data decoder, composed of three inverters and two OR gates, needs as inputs the outputs of the current comparators of the wires Q3, Q2, and Q1 to reconstruct the two bits ($D1out$, $D0out$) sent from the transmitter module. Only the comparator outputs of the threshold current $0.5I$ (i.e., $V(10)$, $V(20)$, and $V(30)$) are needed for this purpose. Formally, the logic is as follows:

$$\begin{aligned}
 (V(30) = '0') \wedge (V(20) = '1') \wedge (V(10) = '1') & \\
 \Rightarrow (D1out = '1') \wedge (D0out = '1'), & \\
 (V(30) = '1') \wedge (V(20) = '0') \wedge (V(10) = '1') & \\
 \Rightarrow (D1out = '1') \wedge (D0out = '0'), & \\
 (V(30) = '1') \wedge (V(20) = '1') \wedge (V(10) = '0') & \\
 \Rightarrow (D1out = '0') \wedge (D0out = '1'), & \\
 (V(30) = '1') \wedge (V(20) = '1') \wedge (V(10) = '1') & \\
 \Rightarrow (D1out = '0') \wedge (D0out = '0'). &
 \end{aligned} \tag{1}$$

The completion detector reads all current comparator outputs as illustrated in Figure 7. For each 4-wire block, the completion detection circuit includes two 4-input NAND gates ($N0$ and $N1$), a 2-input NAND gate ($N2$), and a resettable 2-input C-element ($C1$). To produce the receiver-side request signal $Reqout$, the completion signals of the $N/2$ 4-wire blocks are combined with an $N/2$ -input C-element, where N is the bit-width of the transmitted data. The completion detection process is started by sensing the current

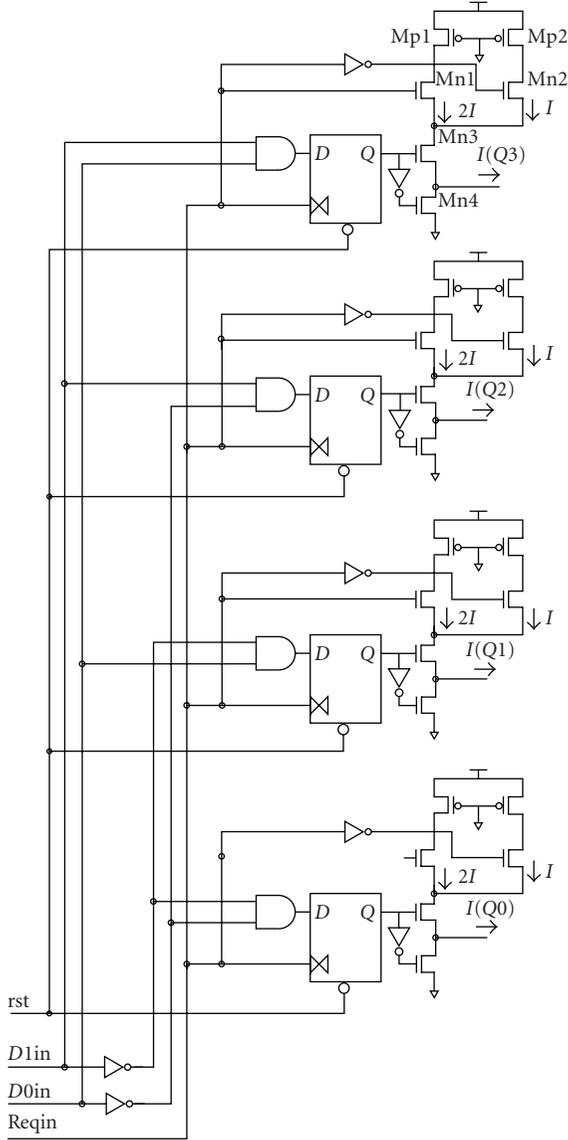


FIGURE 6: Encoder of pulsed multilevel current-mode interconnect.

values on the four wires. In our pulsed implementation of 1-of-4 encoding, current flows only in one of the four wires. Current through the wire becomes I or $2I$ when the transmitter-side request signal $Reqin$ is low or high, respectively. Hence, if the input current of the comparator is greater than the threshold $1.5I$, then the output of the C-element $C1$ and subsequently the receiver-side request signal $Reqout$ go high. Correspondingly, if the comparator input current is between the thresholds $0.5I$ and $1.5I$, the output of $C1$ and the signal $Reqout$ go low. The completion detection logic uses as inputs the current comparator outputs $V(30)$ and $V(31)$ of $Q3$, $V(20)$ and $V(21)$ of $Q2$, $V(10)$ and $V(11)$ of $Q1$, and $V(00)$ and $V(01)$ of $Q0$. For instance, consider again the receipt of the symbol “11” through the wire $Q3$. Assuming that the transmitter-side request signal $Reqin$ is high, the current on the wire $Q3$ is $2I$. Consequently, the comparator outputs

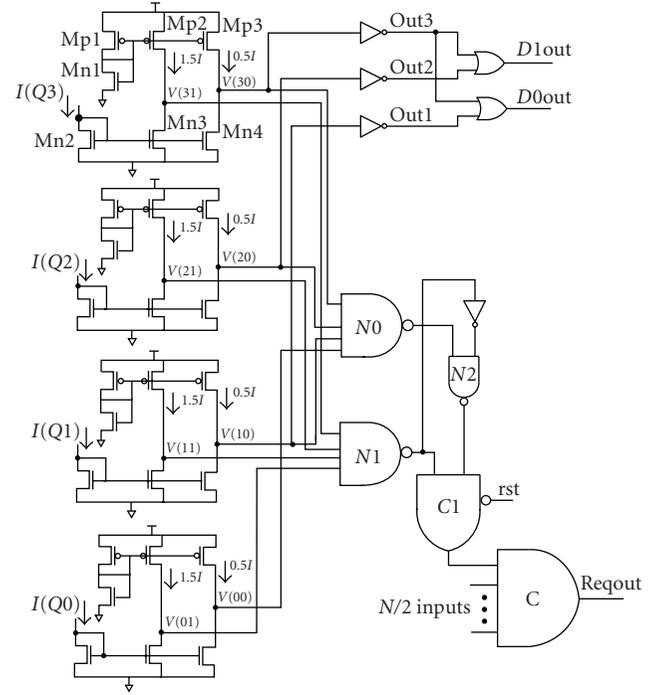


FIGURE 7: Decoder and completion detector circuits of pulsed multilevel current-mode signaling.

$V(30)$ and $V(31)$ become low, and all the other comparator outputs remain high since no current flows through the wires $Q2$, $Q1$, and $Q0$. This makes the outputs of the NAND gates $N1$ and $N2$ high, causing an up-going transition on the output of the C-element $C1$. Formally, the completion detection logic for the symbol “11” is as follows (we denote the output of a gate X by $O(X)$):

$$\begin{aligned}
 & (V(30) = '0') \wedge (V(31) = '0') \quad (\text{current is } 2I) \\
 & \Rightarrow (O(N0) = '1') \wedge (O(N1) = '1') \\
 & \Rightarrow (O(N2) = '1') \\
 & \Rightarrow (O(C1) = '1'), \\
 & (V(30) = '0') \wedge (V(31) = '1') \quad (\text{current is } I) \\
 & \Rightarrow (O(N0) = '1') \wedge (O(N1) = '0') \\
 & \Rightarrow (O(N2) = '0') \\
 & \Rightarrow (O(C1) = '0').
 \end{aligned} \tag{2}$$

The waveforms of $V(30)$ and $V(31)$ are shown in Figure 8.

5.2.4. Acknowledgment transmission

The voltage-mode bundled-data acknowledge signal ($Ackin$), sent by the receiver module, is converted into a current-mode signal during transmission and back into a voltage-mode signal ($Ackout$) at the transmitter side. In this interconnect design transmission of acknowledgment signal also uses multilevel current-mode signaling. The current through

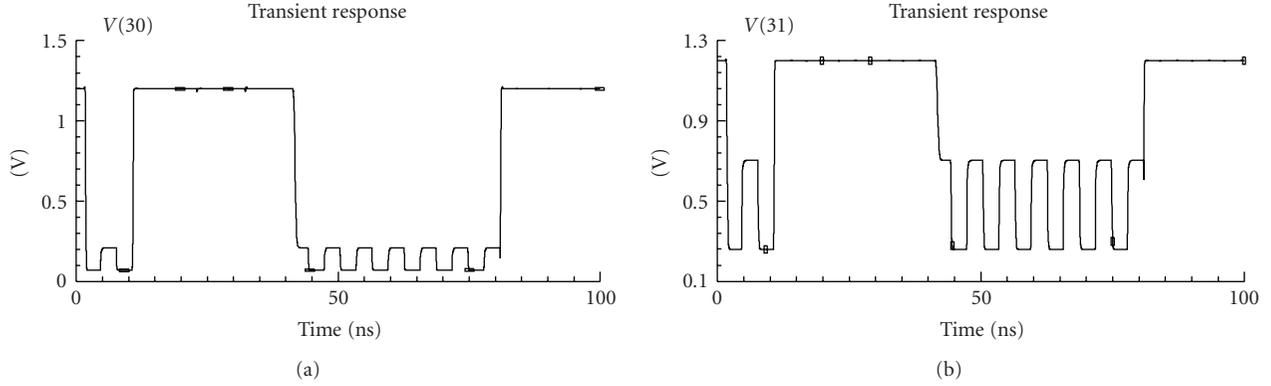


FIGURE 8: Outputs of current comparator.

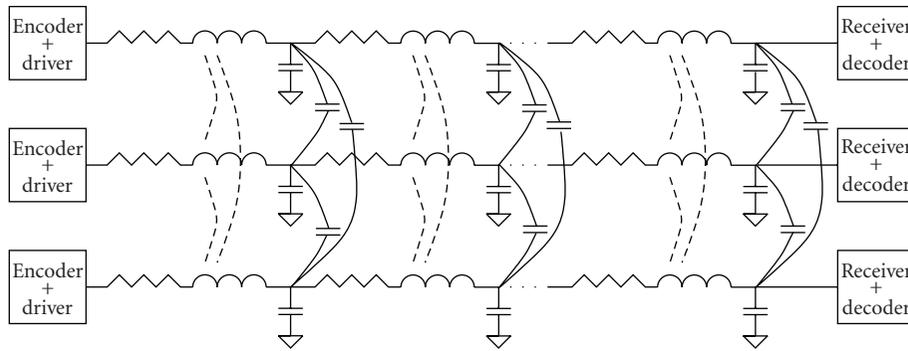


FIGURE 9: Distributed RLC model for capacitively and inductively coupled wires.

acknowledgment wire becomes I and $2I$ when acknowledgment signal from the receiving module is low and high, respectively. The same current comparator circuit is used to detect the value of the current through acknowledgment wire and output the result in voltage form. Inverter is used as a decoding logic.

6. ANALYSIS

6.1. Wire model

Due to the scaling of technology and increasing operating speeds, accurate modeling of wires has become a necessity. Wires have traditionally been modeled as lumped RC segments, but for long high-speed wires, transmission line modeling is needed. Transmission line modeling needs to be applied when the time of flight across the wire becomes comparable to the signal rise time. A transmission line can be thought as a large number of lumped segments in series so that they represent the distributed nature of the wire.

The importance of modeling inductive effects in wires is increasing because of faster rise times and longer wires. Wide wires used in upper metal layers can be especially susceptible to inductive effects due to their low resistance [24]. Since we are considering high-performance signaling over

long wires, we modeled the wires using a distributed RLC model, as shown in Figure 9. In order to accurately model crosstalk noise, both capacitive and inductive coupling between all wires was included.

A 130 nm CMOS technology with metal 4 wires was used. The bus consisted of eight parallel wires. The RLC values of the wires were extracted using field solvers. The resistance and inductance matrices were extracted using FastHenry [25], while the capacitance matrices were extracted using Linpar [26]. The wire length was varied in the simulations from 2 mm to 12 mm, which corresponds to 1–6 expected processing unit widths.

6.2. Performance analysis

In this section, we consider latency and throughput as main parameters to analyze the performance of multilevel current-mode on-chip interconnects along with the two reference voltage-mode interconnects. The most common approach to achieve high-performance long-range on-chip communication is using pipelining or inserting repeaters in voltage-mode signaling. Thus in our first reference interconnect, TPVMP, pipeline stages are inserted every 2 mm assuming that the local wire length (between neighbour routers) is

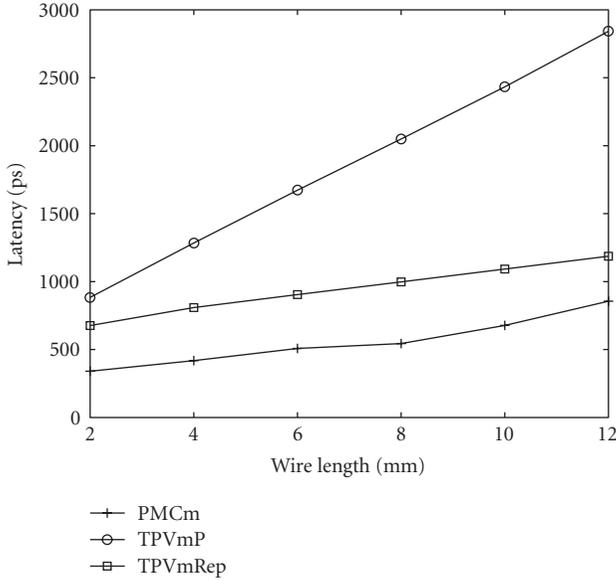


FIGURE 10: Forward latency of the interconnects.

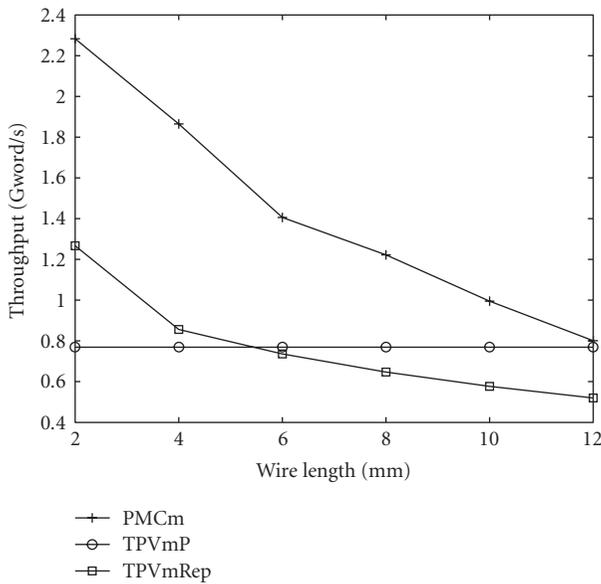


FIGURE 11: Throughput of the interconnects.

2 mm [3]. This improves the throughput at the expense of increased forward latency, power consumption, and chip area. In the second reference interconnect, TPVmRep, optimal size repeaters are inserted at optimal distances.

Here we define forward latency as the delay from a transition on the bundled-data request signal (*Reqin*) at the transmitter side to the corresponding transition on the bundled-data request signal (*Reqout*) at the receiver side (see Figure 3). In other words, the time required for one packet to traverse from the sending router to its receiving router.

The change in the forward latency of the three interconnects when wire length is varied from 2 mm to 12 mm is shown in Figure 10. Since PMCM interconnect uses current-mode signaling, its forward latency is much smaller than the two reference interconnects. At global wire length of 8 mm, PMCM's forward latency was less than one third of TPVmP latency. The latency of pipelined voltage-mode interconnect was much larger than both PMCM and TPVmRep at global lengths of the wire.

The throughput of PMCM, along with the two reference interconnects, is shown in Figure 11 in Gword/s by assuming there is one word packet data transfer between the routers at a time. The throughput of PMCM was greater than the TPVmP and TPVmRep interconnects at all wire lengths (2 to 12 mm) of the interconnect. In case of the reference interconnects, TPVmP achieved a throughput of 769 Mword/s while the throughput of TPVmRep is varied from 1.267 Gword/s to 520 Mword/s when the wire length is varied from 2 to 12 mm. The reported latency and throughput values are for one group of 1-of-4 encoding (2-bit data transfer).

Therefore PMCM interconnect is a better alternative than TPVmP and TPVmRep to realize high-performance long-range NoC links. In addition to achieving high-performance, PMCM circuitry is simpler and takes a smaller chip area compared to pipelined and optimal repeater insertion TPVm. This is because the complexity and required chip area of encoder and decoder of both TPVm and PMCM interconnects are almost the same. However, the number of required pipeline stages and the number of repeaters increase with wire length which makes the two reference TPVm interconnects complex and require larger area for long-range NoC links.

6.3. Power analysis

The average total power consumption for 2-bit data transfer of the proposed current mode and the two reference interconnects when wire length is varied from 2 to 12 mm is shown in Figure 12. The power consumption of PMCM was higher than that of TPVmP at all wire lengths, but its power consumption was lower than that of TPVmRep starting from 6 mm wire length. The power consumption of TPVmP increases at a faster rate with wire length compared to PMCM due to the increase in the number of pipeline stages. Thus at global lengths of wires, the difference in power consumption between these two interconnects decreases considerably. Due to the increase in the number of repeaters inserted at global lengths of the wire, power consumption of TPVmRep is much larger than the other two interconnects.

Here we use a metric called power-throughput ratio which measures the energy consumed per data transmission. This actually corresponds to the power-delay product metric of logic gates. The power-throughput ratio of PMCM is significantly less than that of TPVmRep and slightly greater than that of TPVmP at intermediate and global wire lengths as shown in Figure 13. The voltage-mode interconnect with repeaters has much larger power-throughput ratio than the TPVmP and PMCM interconnects.

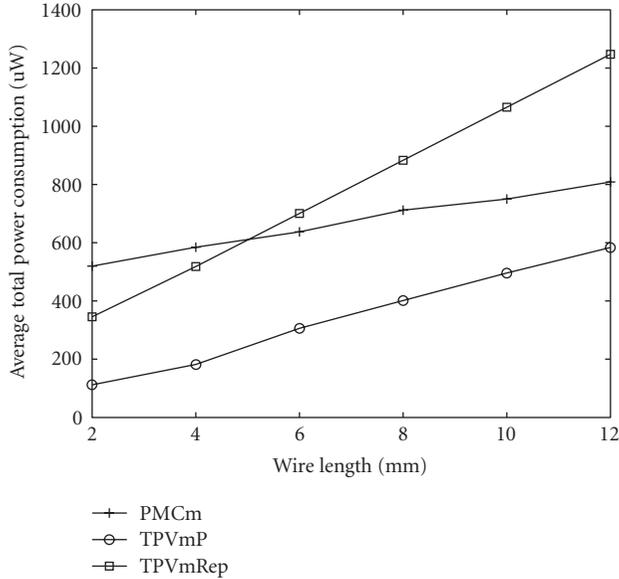


FIGURE 12: Average total power consumption of the interconnects.

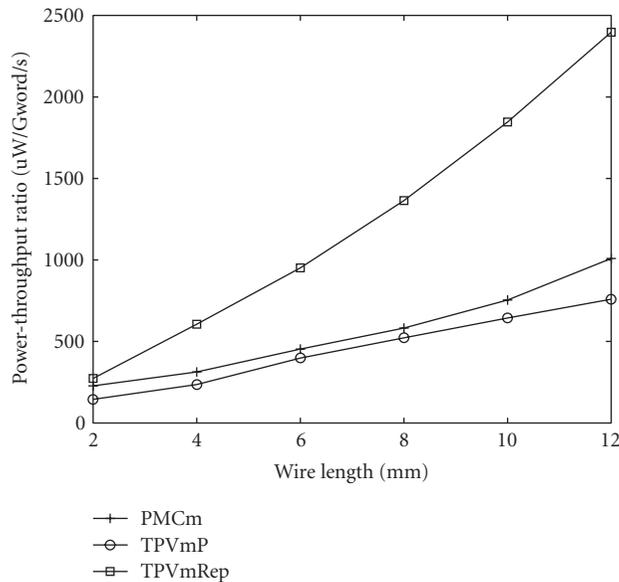


FIGURE 13: Power consumption per throughput of the interconnects.

6.4. Noise analysis

The impact of crosstalk noise on latency and throughput was also studied. In this analysis, 4-bit parallel data transfer was assumed. This requires 9 (8 parallel data transmission + 1 acknowledgment) physical wires since we are using 1-of-4 encoding. The acknowledgment wire was designed as having shielding from the parallel data transmission wires, to counteract the coupling effect. The wires were modeled as transmission lines which have both capacitive and inductive coupling between each other. During this analysis, min-

imum wire separation distance with minimum global pitch specified in 130 nm technology and 1.2 V supply voltage were used. The delay variation due to both capacitive and inductive coupling was simulated by considering the worst-case and best-case switching patterns. These switching patterns depend on the RLC values of the wire. In our case, we assumed that the capacitive coupling dominates the inductive coupling which is the most usual case in on-chip parallel wires. The effect of crosstalk on latency and throughput when the wire length was varied from 2 mm to 12 mm is shown in Figures 14 and 15, respectively.

During best-case and worst-case switching, the latency variation of TPVmP from the latency without crosstalk effect (nominal latency) was slightly less than the PMcM one. For example, at a wire length of 8 mm, the increase in latency due to best-case switching from the nominal latency of TPVmP and PMcM was 59.8% and 62.3%, respectively. In worst-case switching, the TPVmP and PMcM latency variations were 144% and 147%, respectively, at the same wire length. In fact, these percentage values are rather large because in the nominal case shown in Figure 10 the considered capacitive loads were only to ground. In other words, the nominal case capacitive loads do not consider the coupling capacitances loading effect. The decrease in throughput due to crosstalk was greater for TPVmP than for PMcM, specially at long wire length. For example at 12 mm wire length, the throughput of TPVm was decreased by 38% while the PMcM was only by 30%.

7. DISCUSSION

In order to mitigate real-life applications, we can assume there is a 64-bit data transfer using the long links presented in this work. In the pipelined voltage-mode interconnect, there is a need of completion detection circuits at both sides of the communication; in the receiver side to indicate the validity of the arrived data and in the transmitter side to indicate the acceptance of the transmitted data since an acknowledgment is sent per each 1-of-4 group. This requires 32-input C-element at both sides which creates a considerable delay, because its complexity approximately corresponds to that of a 32-input AND gate with multiple logic levels. However, the pulsed multilevel current-mode interconnect requires only receiver side completion detection since it is possible to use one acknowledgment signal per data transfer. Even though the delay due to the completion detection is reduced by half in the current-mode interconnect, it is necessary to have a fast completion detection mechanism. Thus, our future work will be designing of a fast and area efficient completion detection circuit for the pulsed multilevel current-mode interconnect, for example performing completion detection by sensing currents. This can be done by summing up the currents of all the wires and comparing the sum with a threshold current.

Based on *International Technology Roadmap for Semiconductors* [27], the long on-chip wire length can be even longer than 10 mm in future nanoscale technologies. The proposed current-mode interconnect throughput becomes

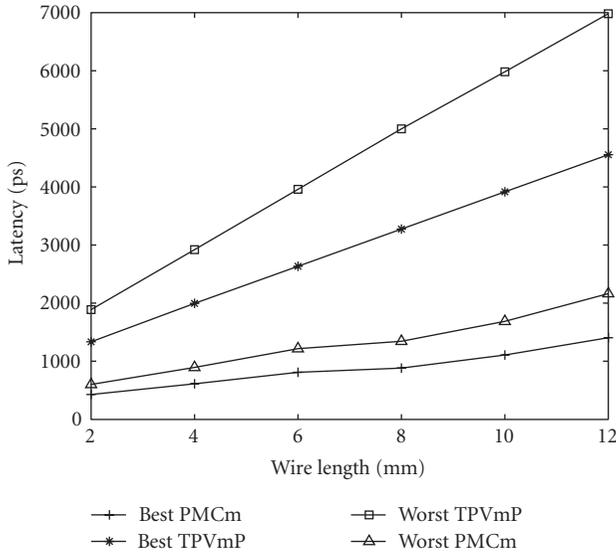


FIGURE 14: Forward latency of the interconnects in the presence of crosstalk.

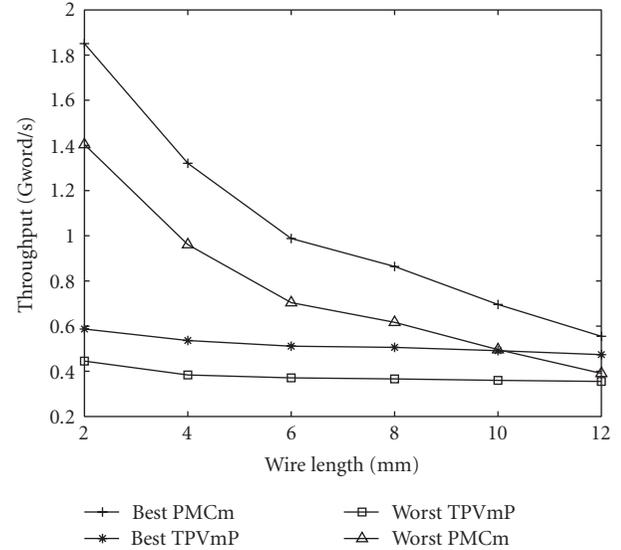


FIGURE 15: Throughput of the interconnects in the presence of crosstalk.

almost equal and even slightly less than the pipelined voltage mode throughput when the wire length exceeds 10 mm. To maintain the high throughput of the current-mode interconnect, an efficient current-mode pipeline stage could be inserted after every 10 mm wire length. This increases the forward latency but it will not be that significant since a pipeline stage is needed only every 10 mm.

Another direction of our future work is examining how much improvement in overall average latency and throughput of the network can be achieved by using our high-performance current-mode link for end-around torus channels and for additional long channels of mesh network and what are the expenses.

8. CONCLUSION

We presented a high-performance delay-variation-insensitive long on-chip interconnect which uses two-phase 1-of-4 encoding and multilevel current-mode signaling. This interconnect is a promising candidate for long-range NoC communication links since it has low latency, high throughput, and low power-throughput ratio. In addition, its delay-insensitive data transfer ability makes it appropriate for future nanoscale long-range NoC interconnects where delay variations are inevitable. Since the usual way of improving the performance of long on-chip interconnects is using voltage mode signaling along with either repeater insertion or using pipeline stages, we designed two-phase 1-of-4 encoded voltage-mode signaling references, one with pipeline stages and the other with optimally inserted repeaters. These voltage-mode interconnects serve as references to our proposed current-mode interconnect.

The performance analysis shows that the current-mode interconnect has higher throughput and lower latency than the two reference interconnects. It achieves a throughput of 1.222 Gword/s at 8 mm wire length which is 1.58 times higher than the throughput of the pipelined voltage-mode interconnect and 1.89 times higher than the one using optimal repeater insertion. From the power consumption analysis, it is seen that the current-mode interconnect consumes less power than the reference voltage-mode interconnect with optimal repeater insertion starting from the wire length of 6 mm. On the other hand, it consumes more power than the voltage-mode interconnect with pipeline stages for 2 to 12 mm wire length. However, the power consumption difference between these two interconnects becomes smaller when the wire length increases. The power-throughput ratio of the proposed current-mode interconnect is much less than the voltage-mode interconnect with optimal repeaters and slightly greater than the pipelined voltage-mode interconnect. The effects of crosstalk on latency and throughput of the interconnects are also analyzed. The variation in forward latency of the current-mode interconnect was a few percents larger than that of the pipelined voltage-mode interconnect. In case of throughput reduction due to crosstalk, the throughput of pipelined voltage mode is more affected than the current mode.

Therefore, using the proposed multilevel current-mode interconnect for long-range NoC links such as the torus end-around channels allows the network to achieve high throughput and low latency along with delay-variation-insensitive communication. The delay insensitivity makes the communication robust and attains average-case performance rather than worst-case performance which is the situation in communication based on timing constraints.

REFERENCES

- [1] U. Y. Ogras and R. Marculescu, "It's a small world after all": NoC performance optimization via long-range link insertion," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 14, no. 7, pp. 693–706, 2006.
- [2] W. J. Dally and B. P. Towles, *Principles and Practices of Interconnection Networks*, Morgan Kaufmann-Elsevier, San Francisco, Calif, USA, 2004.
- [3] W. J. Dally and B. Towles, "Route packets, not wires: on-chip interconnection networks," in *Proceedings of the 38th Design Automation Conference (DAC '01)*, pp. 684–689, Las Vegas, Nev, USA, June 2001.
- [4] D. Sylvester and K. Keutze, "A global wiring paradigm for deep submicron design," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 19, no. 2, pp. 242–252, 2000.
- [5] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi, and V. De, "Parameter variations and impact on circuits and microarchitecture," in *Proceedings of the 40th Design Automation Conference (DAC '03)*, pp. 338–342, Anaheim, Calif, USA, June 2003.
- [6] R. Ho, J. Gainsley, and R. Drost, "Long wires and asynchronous control," in *Proceedings of the 10th International Symposium on Asynchronous Circuits and Systems (ASYNC '04)*, pp. 240–249, Crete, Greece, April 2004.
- [7] T. Verhoeff, "Delay-insensitive codes—an overview," *Distributed Computing*, vol. 3, no. 1, pp. 1–8, 1988.
- [8] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*, Cambridge University Press, Cambridge, UK, 1998.
- [9] D. Pamunuwa and H. Tenhunen, "Repeater insertion to minimize delay in coupled interconnects," in *Proceedings of the 14th IEEE International Conference on VLSI Design*, pp. 513–517, Bangalore, India, January 2001.
- [10] R. Bashirullah, W. Liu, and R. K. Cavin III, "Current-mode signaling in deep submicrometer global interconnects," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 11, no. 3, pp. 406–417, 2003.
- [11] A. Katoch, E. Seevinck, and H. Veendrick, "Fast signal propagation for point to point on-chip long interconnects using current sensing," in *Proceedings of the 28th European Solid-State Circuits Conference (ESSCIRC '02)*, pp. 195–198, Florence, Italy, September 2002.
- [12] A. Katoch, H. Veendrick, and E. Seevinck, "High speed current-mode signaling circuits for on-chip interconnects," in *Proceedings of IEEE International Symposium on Circuits and Systems (ISCAS '05)*, vol. 4, pp. 4138–4141, Kobe, Japan, May 2005.
- [13] A. P. Jose, G. Patounakis, and K. L. Shepard, "Near speed-of-light on-chip interconnects using pulsed current-mode signalling," in *Proceedings of IEEE Symposium on VLSI Circuits Digest of Technical Papers*, pp. 108–111, Kyoto, Japan, June 2005.
- [14] M. K. Gowan, L. L. Biro, and D. B. Jackson, "Power considerations in the design of the Alpha 21264 microprocessor," in *Proceedings of the 35th Design Automation Conference (DAC '98)*, pp. 726–731, San Francisco, Calif, USA, June 1998.
- [15] R. Bashirullah, "Reduced delay sensitivity to process induced variability in current sensing interconnects," *Electronics Letters*, vol. 42, no. 9, pp. 531–532, 2006.
- [16] J. Q. Zhang, S. I. Long, F. H. Ho, and J. K. Madsen, "Low power current mode multi-valued logic interconnect for high speed interchip communications," in *Proceedings of the 17th Annual IEEE Gallium Arsenide Integrated Circuit Symposium (GaAs IC '95)*, pp. 327–330, San Diego, Calif, USA, October–November 1995.
- [17] J.-Y. Sim, Y.-S. Sohn, S.-C. Heo, H.-J. Park, and S.-I. Cho, "A 1-Gb/s bidirectional I/O buffer using the current-mode scheme," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 4, pp. 529–535, 1999.
- [18] I. B. Dhaou, M. Ismail, and H. Tenhunen, "Current mode, low-power, on-chip signaling in deep-submicron CMOS technology," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 50, no. 3, pp. 397–406, 2003.
- [19] T. Temel and A. Morgul, "Implementation of multi-valued logic gates using full current-mode CMOS circuits," *Analog Integrated Circuits and Signal Processing*, vol. 39, no. 2, pp. 191–204, 2004.
- [20] T. Hanyu, T. Takahashi, and M. Kameyama, "Bidirectional data transfer based asynchronous VLSI system using multiple-valued current mode logic," in *Proceedings of the 33rd International Symposium on Multiple-Valued Logic*, pp. 99–104, Tokyo, Japan, May 2003.
- [21] E. Nigussie, J. Plosila, and J. Isoaho, "Delay-insensitive on-chip communication link using low-swing simultaneous bidirectional signaling," in *Proceedings of IEEE Computer Society Annual Symposium on Emerging VLSI Technologies and Architectures*, pp. 217–222, Karlsruhe, Germany, March 2006.
- [22] V. Venkatraman and W. Burlinson, "Robust multi-Level current-mode on-chip interconnect signaling in the presence of process variations," in *Proceedings of the 6th International Symposium on Quality of Electronic Design (ISQED '05)*, pp. 522–527, San Jose, Calif, USA, March 2005.
- [23] R. Venkatesan, J. A. Davis, and J. D. Meindl, "Compact distributed RLC interconnect models—part IV: unified models for time delay, crosstalk, and repeater insertion," *IEEE Transactions on Electron Devices*, vol. 50, no. 4, pp. 1094–1102, 2003.
- [24] Y. I. Ismail, E. G. Friedman, and J. L. Neves, "Figures of merit to characterize the importance of on-chip inductance," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 7, no. 4, pp. 442–449, 1999.
- [25] M. Kamon, M. J. Tsuk, and J. K. White, "FASTHENRY: a multipole-accelerated 3-D inductance extraction program," *IEEE Transactions on Microwave Theory and Techniques*, vol. 42, no. 9, pp. 1750–1758, 1994.
- [26] A. R. Djordjevic, M. B. Bazdar, T. K. Sarkar, and R. F. Harrington, *Linpar for Windows: matrix parameters for multiconductor transmission lines*, Software and User Manual, Version 2.0, Artech House Publisher, Norwood, Mass, USA, 1999.
- [27] International Technology Roadmap for Semiconductors, 2005, <http://public.itrs.net/>.



Hindawi

Submit your manuscripts at
<http://www.hindawi.com>

