MIT Libraries | DSpace@MIT

MIT Open Access Articles

40.4fJ/bit/mm Low-Swing On-Chip Signaling with Self-Resetting Logic Repeaters Embedded within a Mesh NoC in 45nm SOI CMOS

MIT Massachusetts Institute of Technology

40.4fJ/bit/mm Low-Swing On-Chip Signaling with Self-Resetting Logic Repeaters Embedded within a Mesh NoC in 45nm SOI CMOS

Sunghyun Park[†], Masood Qazi[‡], Li-Shiuan Peh[†] and Anantha P. Chandrakasan[†]

[†]EECS Department, Massachusetts Institute of Technology, Cambridge, MA, USA

[‡]Cypress Semiconductor Corporation, San Jose, CA, USA

Email: pshking@mit.edu, qazi@cypress.com, peh@csail.mit.edu, anantha@mtl.mit.edu

Abstract—**Mesh NoCs are the most widely-used fabric in high-performance many-core chips today. They are, however, becoming increasingly power-constrained with the higher on-chip bandwidth requirements of high-performance SoCs. In particular, the physical datapath of a mesh NoC consumes significant energy. Low-swing signaling circuit techniques can substantially reduce the NoC datapath energy, but existing low-swing circuits involve huge area footprints, unreliable signaling or considerable system overheads such as an additional supply voltage, so embedding them into a mesh datapath is not attractive. In this paper, we propose a novel low-swing signaling circuit, a self-resetting logic repeater, to meet these design challenges. The SRLR enables single-ended low-swing pulses to be asynchronously repeated, and hence, consumes less energy than differential, clocked low-swing signaling. To mitigate global process variations while delivering high energy efficiency, three circuit techniques are incorporated. Fabricated in 45nm SOI CMOS, our 10mm SRLR-based low-swing datapath achieves 6.83Gb/s/μm bandwidth density with 40.4fJ/bit/mm energy at 4.1Gb/s data rate at 0.8V.**

I. INTRODUCTION

Multicore architectures have become mainstream in high-performance SoCs. As core counts grow, Networks-on-Chips (NoCs) have emerged as a scalable, high-bandwidth communication backbone in such chips [1]. A mesh is the most widely-used NoC topology for high-performance many-core chips as it is scalable and maps readily to on-die layout [2]–[6]. Unlike indirect, multi-stage NoC topologies such as Clos or Butterflies [7], meshes support the locality present in many applications, allowing nearby traffic to be transported at lower delay and energy.

As high-performance SoCs with many cores demand increasing on-die bandwidth, NoCs are becoming increasingly power-constrained. When we inspect the power breakdown of a mesh NoC, it comprises three components: links (39% in RAW, 31% in TRIPS, 17% in TeraFLOPS), crossbar switches (30% in RAW, 33% in TRIPS, 15% in TeraFLOPS) and buffers (31% in RAW, 35% in TRIPS, 22% in TeraFLOPS) [2], [3], [6]. While buffer power can be reduced by virtual bypassing flow control [8]–[10] or bufferless routing algorithms [11]–[13], links and crossbar switches form the unavoidable portion of mesh NoC power, being responsible for physical data transmission through metal wires. Furthermore, this physical datapath power will increase in percentage relative to control and storage circuitry power as CMOS process technology

scales down [14], [15]. Therefore, it is critical to reduce the power consumption of the physical datapath composed of links and crossbar switches.

Low-swing signaling is now one well-known low-power design technique that can significantly improve the energy efficiency of the NoC datapath [16]. The low-swing circuit technique is based on the dependence of dynamic energy on swing voltage. Reducing the voltage swing across a datapath leads to decreased charging and discharging of the wire capacitance in comparison with the full-swing signaling, thereby making the datapath more energy-efficient. Low-swing drivers have been embedded within mesh NoC routers and shown to substantially reduce NoC energy [17], [18], but existing low-swing circuits face key NoC design challenges. First, the area overhead imposed by low-swing drivers is of prime concern, since a NoC shares precious on-die real estate with processor cores, caches, memory controllers, etc. Second, low-swing signaling comes at the cost of reduced noise margin, which is crucial as packet losses are not tolerated in NoCs. Thirdly, existing low-swing circuits impose a considerable system overhead such as an additional dedicated power supply voltage or clocking circuitry in an entire NoC datapath, or provide energy-optimal design of only one-to-one signaling, making their adoption in a mesh fabric infeasible. We will next explain in detail why prior circuits come up short in area, robustness and energy-efficient application to a mesh.

Apart from traditional low-swing circuits which use a lower supply voltage or inherent threshold voltage drop [16], [19], there have been a number of more sophisticated low-swing circuits proposed, based on linear-mode transistors [17], [18], charge sharing [20]–[22], cut-off drivers [19], [23], [24] and channel attenuation [25]–[27]. The low-swing drivers exploiting linear-mode transistors [17], [18] are composed of PMOS pullups and pulldowns only (or NMOS pullups and pulldowns only) to obtain lower linear drive resistance even at small V_{ds}. While such designs enable better energy efficiency and higher bandwidth than the traditional low-swing signaling generated by simply lowering power supply voltage, they require differential wiring, clocked sense amplifiers and an additional power supply voltage. In particular, the additional power supply dedicated only to the NoC low-swing datapath can be a substantial system overhead in multicore processor design. The

charge sharing-based low-swing drivers [20]–[22] limit voltage swing without a second power supply voltage, but they require fixed data patterns for reliable operation, which is infeasible in NoCs. The voltage swing of the cut-off drivers [19], [23], [24] is directly affected by threshold voltage variation of drive transistors, thus requiring complicated receivers to sense and calibrate the threshold voltage variation, resulting in significant area overhead.

Equalized on-chip interconnects [25]–[27] can generate low-swing signaling by leveraging the inherent channel attenuation of RC-dominant wires and have successfully provided high-bandwidth low-power global links that transmit data through long wires (5-10mm). These long equalized links can be used as point-to-point wires between pairs of cores, but as there is insufficient on-die wiring to support dedicated links between all pairs of cores, equalized links map more readily to indirect, multi-stage NoC topologies with long global links. However, as we mentioned earlier, such topologies do not leverage application locality, turning all traffic into cross-die global traversals, which leads to high NoC latency and energy overheads. Meshes, on the other hand, are dominated by short local core-to-core links. Adopting equalizers as parallel links in a mesh NoC will lead to huge area overhead (*e.g.,* the 10mm 1-bit driver of [26] occupies $1760\mu m^2$). Yet another way of incorporating long equalized links in meshes is to use them as express links between far-away cores [28], [29]. That increases router port count though, leading to high NoC area overhead. Besides, direct transmission on a long global wire makes equalized interconnects vulnerable to wire capacitance/resistance variation and crosstalk coupling noise.

In this paper, we seek to tackle the above-mentioned design challenges of incorporating low-swing signaling in a mesh NoC datapath. We propose a novel low-swing signaling circuit: self-resetting logic repeaters (SRLR) embedded within each router of a mesh NoC. Our proposed low-swing signaling circuit has the following features:

- The SRLR enables single-ended low-swing signaling, and hence, consumes less energy than differential low-swing signaling at the same wire density (*i.e.,* the SRLR can have higher wire density at the same energy budget).
- The SRLR achieves low-swing signaling mainly through the inherent wire channel attenuation so it does not require additional power supplies and works across all data patterns.
- The SRLR enables low-swing signaling to be repeated with a single repeater length (*i.e.,* ∼1mm), the wire length of local core-to-core links in a mesh. A single optimized SRLR design can thus be used for energy-efficient signaling between any pair of nodes in a mesh. As a side benefit, the SRLR enables 1-to-N multicasts for free since inherent full-swing signals are available at every intermediate repeater node. This multicast capability is a significant benefit as multicast traffic forms a sizable portion of NoC traffic [10].
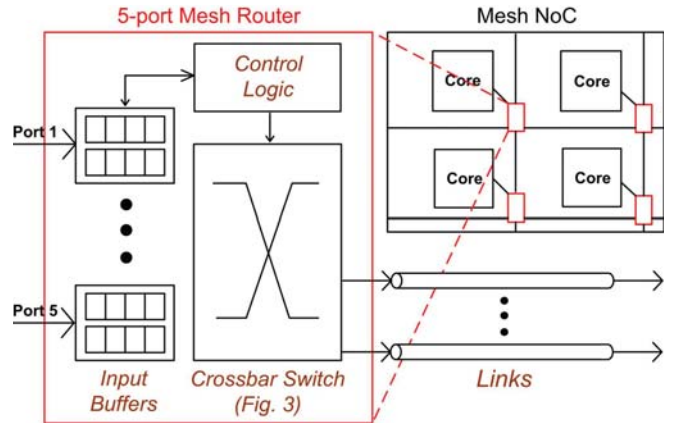- The SRLR incorporates circuit techniques to mitigate global process variation and ensure robustness.

Fig. 1.   Typical 5-port mesh router microarchitecture.

Fabricated in 45nm SOI CMOS, our SRLR low-swing circuit achieves 40.4fJ/bit/mm energy at 6.83Gb/s/μm bandwidth density at 0.8V.[1] Each 1mm SRLR takes 10.2x4.7=47.9μm^2 of active silicon area, so the SRLR-based low-swing datapath of a 64-bit 5-port mesh router (as will be shown in Fig. 3) will occupy 47.9x64x5x4=0.061mm^2. This area overhead is reasonable when compared to 0.34mm^2 of the overall router area, for a 3-stage mesh NoC router with 4 VCs and 16 buffers [15] (*i.e.,* the proposed low-swing datapath will take about 18% of the entire router footprint). Robustness wise, measurement results show that our link delivers at BER less than 10^{-9} and 1000-run Monte Carlo simulations attest to its robustness against process variations.

The rest of this paper is organized as follows: Section 2 explains the proposed SRLR circuit and its transistor sizing methodology for energy-efficient, reliable signaling. Section 3 focuses on the circuit techniques to make the SRLR variation-tolerant. Section 4 shows the measurement results of our SRLR test chip fabricated in 45nm SOI CMOS, and finally, we conclude in Section 5.

II. SELF-RESETTING LOGIC REPEATER

Fig. 1 shows a typical 5-port mesh router microarchitecture composed of 4 main components: input buffers, a control logic, a crossbar switch and links. The input buffers store incoming packets till they are sent to the next router. The control logic determines when specific packets proceed through the router pipeline and sets up the crossbar switch. The crossbar switch physically moves data from input ports to output ports, followed by links that forward output port data to the next router. As mentioned earlier, the crossbar switch and links form the NoC datapath whose power consumption is unavoidable [18] and comprises significant portion of NoC power (69% in RAW [2], 64% in TRIPS [3], 32% in TeraFLOPS [6]). Our proposed SRLR enables low-swing signaling throughout such NoC datapath, resulting in low-power mesh NoCs.

Fig. 2 shows the overall 10mm link with SRLRs located at the end of each 1mm wire segment connecting adjacent

---

[1]The energy unit, [fJ/bit/mm], is normalized by wire length while the bandwidth density unit, [Gb/s/μm], is normalized by wire density given by wire width and space.
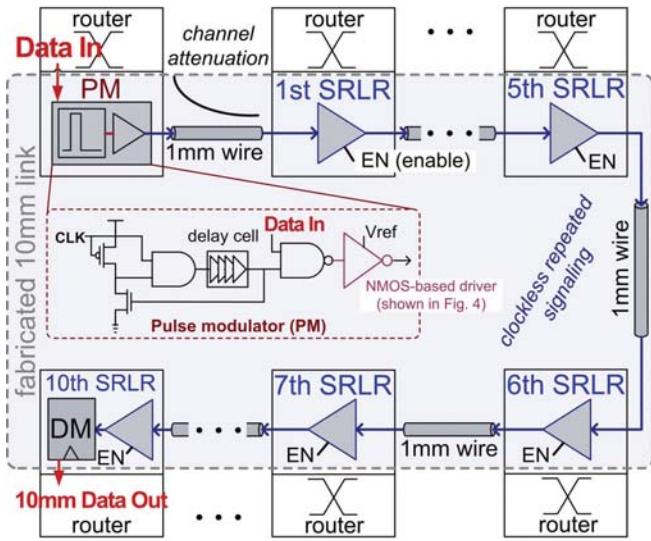
Fig. 2. 10mm SRLR-based link for the mesh NoC where the local router-to-router distance is 1mm.
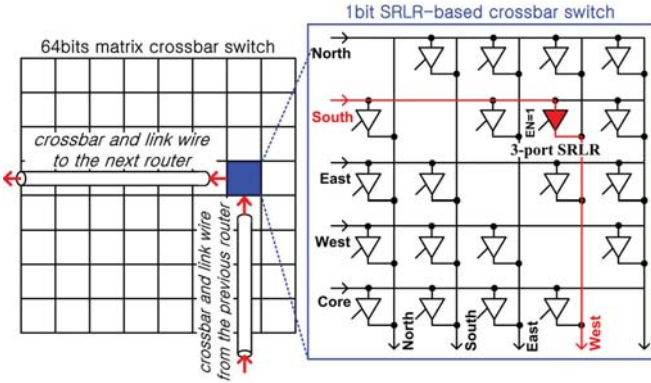

Fig. 3. 64bit low-swing crossbar switch and links composed of 3-port SRLRs (IN, OUT, EN) inserted at each of the 20 crosspoints of the crossbar switch (*i.e.,* 64x20 SRLRs in total).

routers in a mesh NoC while Fig. 3 describes how such SRLR-based low-swing signaling can be integrated into a crossbar switch. Typically, embedding repeaters within the crosspoints of a crossbar can lead to increased layout complexity due to the active silicon region in the midst of wires. The SRLR-based datapath, however, averts that by ensuring that the SRLR insertion length is equal to the router-to-router distance in a mesh NoC. This work assumes that the local router-to-router distance is 1mm, and accordingly, the SRLR transistors are optimally-sized to directly drive the 1mm wire in order to offer low-swing repeated signaling without adding to layout complexity. The only implementation overheads of the proposed low-swing signaling are thus a pulse modulator (PM) and a demodulator (DM) required for pulse-based data communication. With the PMs and DMs at every router, our proposed circuit can send low-swing pulses to a far-away node in a mesh without energy overheads since each SRLR drives only a 1mm wire segment and the low-swing pulses are repeated without clocking.

In addition, our SRLR-based datapath provides low-swing 1-to-N multicast capability for free while equalized links [25]–
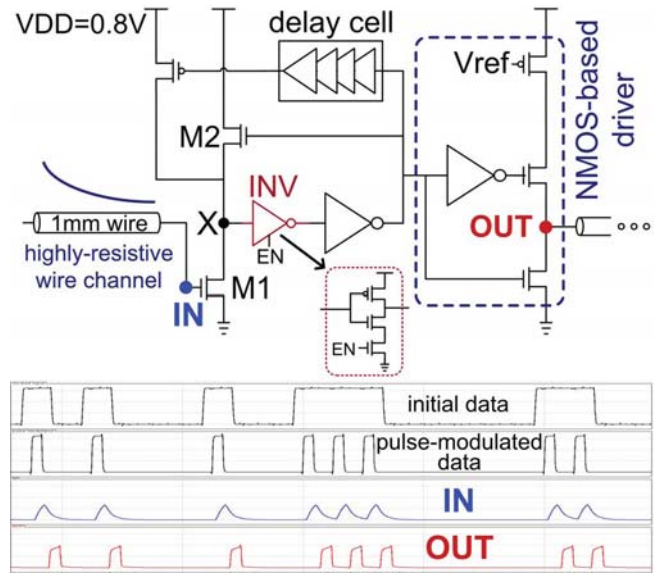

Fig. 4. Proposed SRLR circuit with simulated waveforms.

[27] offer only 1-to-1 unicasts. For instance, in Fig. 2, the data sent from the 1st SRLR to the 10th SRLR can be directly sampled at all the intermediate SRLRs (5th SRLR, 6th SRLR, 7th SRLR, ...). This inherent multicast capability can result in substantial benefits in NoCs that see significant multicast traffic [10].

Fig. 4 shows the proposed SRLR circuitry along with its simulation waveforms. When a pulse (whose low-swing is obtained by wire channel attenuation) arrives at an input NMOS (M1), the node X is discharged and output voltage of the SRLR (OUT) becomes high. The node X is again charged when a reset signal comes back through a delay cell, generating another pulse at the output. As a last step, a keeper NMOS (M2) lowers the node X voltage down to VDD-V_{th} after the pulse is repeated. The reduced standby voltage at the node X increases amplification gain of the current-starved inverter (INV) but this standby voltage should stay above the threshold voltage of INV across process variation. Also, the size ratio of M1/M2 should be designed to allow enough SRLR input sensitivity at a given low-swing voltage level.

The current-starved inverter (INV) amplifier becomes activated when enable signal (EN) is high, and this 3-port (IN, OUT and EN) circuit design allows SRLRs to be directly integrated into a crossbar switch as shown in Fig. 3.

III. PROCESS VARIATION ROBUST SRLR CIRCUIT

While single-ended low-swing signaling has higher energy efficiency than differential low-swing signaling, this comes at the expense of global (die-to-die) process variation immunity. To mitigate such variation effects on the proposed on-chip signaling, the SRLR-based link employs three circuit techniques: an alternating delay cell design, an NMOS-based driver and an adaptive swing voltage scheme.

A. Alternating Delay Cell Design

First, we propose an alternating delay cell design where odd SRLRs and even SRLRs incorporate different delay cells.
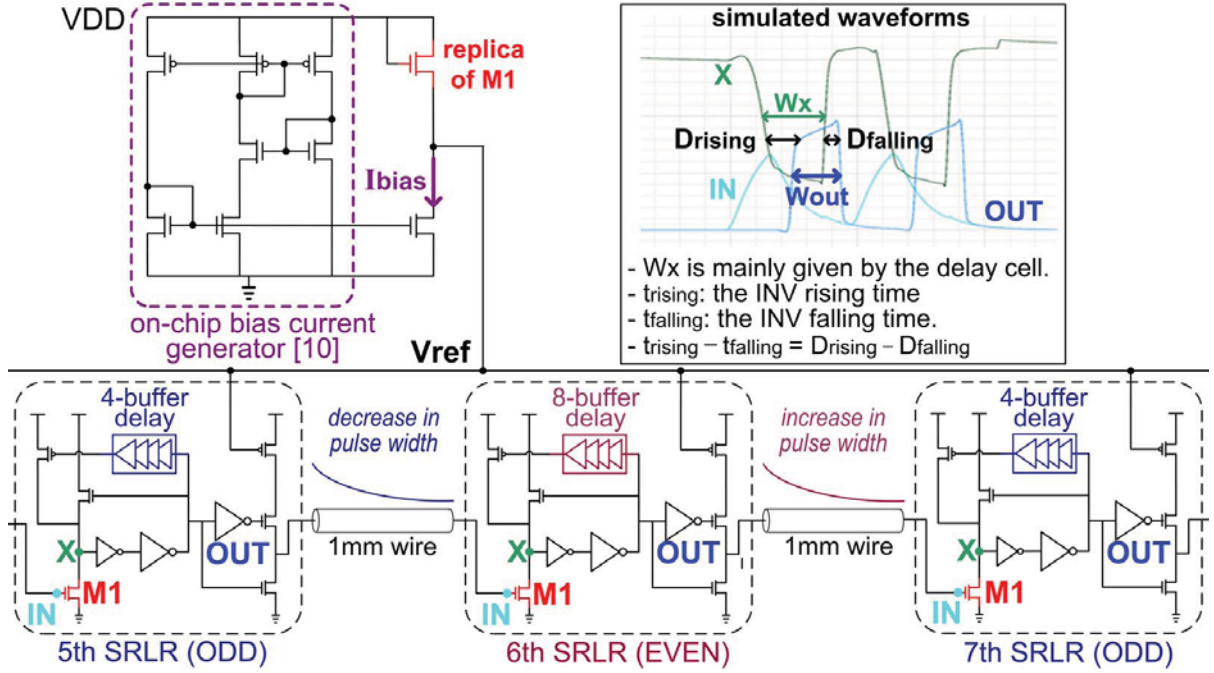
Fig. 5. Process variation robust SRLR circuit with (1) an alternating delay cell design, (2) NMOS-based drivers and (3) an adaptive swing voltage scheme.

As shown in Fig. 5, the SRLR output pulse width (W_{out}) is a function of node X's pulse width (W_x), which is mainly given by the delay of the delay cell, and the difference between rising time (t_{rising}) and falling time ($t_{falling}$) of the INV amplifier. At an $n - th$ SRLR, the output pulse width ($W_{out,n}$) can be expressed as

$$
\begin{aligned}
W_{out,n} &= W_{x,n} - D_{rising,n} + D_{falling,n} \\
&= W_{x,n} - (t_{rising,n} - t_{falling,n}).
\end{aligned}
$$

The rising time becomes longer *(or shorter)* as input pulse swing gets smaller *(or bigger)*; whereas, the falling time experiences little change with the input pulse swing change. With a single delay cell design (*e.g.,* 6-buffer)[2], this influence of the input pulse swing on the rising time of INV accumulates over several SRLR stages, and hence, the rising time gradually becomes longer *(or shorter)* at the smaller *(or bigger)* initial pulse swing caused by the process variation. The increasing *(or decreasing)* rising time causes a shrinking *(or widening)* output pulse width, resulting in a failure at the end of the 10mm link. In other words, the output pulse widths obtained from process corner simulations of the single delay cell design are

$$W_{out,0} > W_{out,1} > W_{out,2} > \ldots > W_{out,10} \quad (1)$$
$$(bit\ 1\ transmission\ failure)$$
$$or$$
$$W_{out,0} < W_{out,1} < W_{out,2} < \ldots < W_{out,10}. \quad (2)$$
$$(bit\ 0\ transmission\ failure)$$

[2]This 6-buffer delay enables the single delay cell design to offer the most reliable repeated signaling at a typical process condition (*i.e.,* no-variation simulation environment).

The proposed alternating delay cell design, on the other hand, enables output pulse widths to increase *(or decrease)* even with the longer *(or shorter)* rising time of the INV amplifier through the intentionally-increased *(or intentionally-decreased)* delay of the delay cell. The alternating design can still saturate, but because of the non-linearity of the feedback (where larger input pulse width causes even larger change in output pulse width) the alternating design takes more stages to saturate. Therefore, the alternating design improves the probability of correct operation for a fixed link length.

B. NMOS-based Driver

Global process variation influences the output stage of the SRLR as well. Under a straightforward implementation, an inverter driver at the output exhibits two distinct failure modes. In one mode, a weak PMOS will generate insufficient voltage swing at the input of the following stage. In the other mode, a strong PMOS generates too much voltage swing for a weak NMOS to fully discharge the node at the end of a wire channel prior to the arrival of the next bit. Accordingly, the worst-case sequence of '11110' will eventually saturate the voltage and prevent transmission of several 1s followed by a 0. The NMOS-based driver in this work (Fig. 2) supplies both pull-up and pull-down currents through NMOS devices, so the strong PMOS condition no longer applies. The resulting circuit is more robust since it is optimized for only one failure mode at a weak NMOS corner, instead of two distinct failure modes across a weak PMOS or a strong PMOS with weak NMOS.

C. Adaptive Swing Voltage Scheme

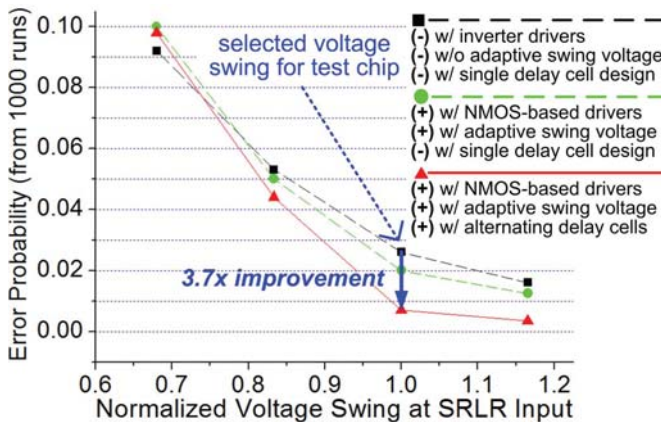Having a robust NMOS-based circuit also allows the optimization of transmission energy. At a strong NMOS corner,

Fig. 6. Monte-Carlo simulation results with various swing voltages.

Fig. 7. Die photograph of the test chip in 45nm SOI CMOS.

the output pulse swing tends to be excessively high, especially for the lower V_{th} of the input NMOS (M1) of the next stage. Therefore, the adaptive voltage swing scheme (Fig. 5) with an on-chip bias current generator[3] tracks the M1 threshold voltage to reduce swing voltage, avoiding the needless waste of energy. In other words, When M1 is fabricated with higher *(or lower)* threshold voltage than the nominal value, the lower *(or higher)* V_{ref} is applied to the NMOS-based drivers to increase *(or decrease)* voltage swing.

Fig. 6 shows the error probability obtained from 1000-run Monte-Carlo simulations on different SRLR designs with various swing voltages. At the voltage swing selected for test chip fabrication, the proposed process variation robust SRLR design achieves about 3.7 times higher process variation immunity than the straightforward SRLR design that incorporates inverter drivers (instead of NMOS-based drivers) and 6-buffer delay cells only (instead of an alternating delay cell design) without the adaptive swing voltage scheme.

IV. EXPERIMENT RESULTS

To explore the energy efficiency and performance of the proposed low-swing on-chip signaling, a proof-of-concept chip of a 1bit 10mm SRLR-based link (described in Fig. 2) is implemented using a 45nm SOI CMOS process. Fig. 7 shows its die photograph overlaid with a design layout where each SRLR occupies $47.9\mu m^2$ active silicon area.

The fabricated link is fed by pseudo-random binary sequence data generated on-chip and a test circuit performs data comparison and error counting. This on-chip measurement circuit shows that the 1bit 10mm SRLR-based on-chip interconnect can deliver up to 4.1Gb/s data with the BER that is less than 10^{-9}. Measurement results show that the SRLR-based on-chip signaling achieves 6.83Gb/s/μm bandwidth density at its maximum data rate of 4.1Gb/s, consuming 1.66mW (*i.e.,* 404fJ/bit/cm or 40.4fJ/b/mm) at a power supply voltage of 0.8V. Fig. 8 shows 10mm link traversal (LT) energy versus

[3]This bias current, which does not contain any threshold voltage-related terms for the first order analysis [30], is tolerant of process and temperature variations so that V_{ref} is mainly given by the threshold voltage and technology parameters of M1, a primary determinant transistor of the SRLR input sensitivity.
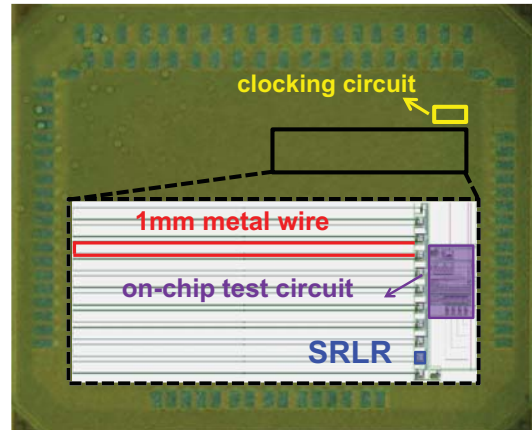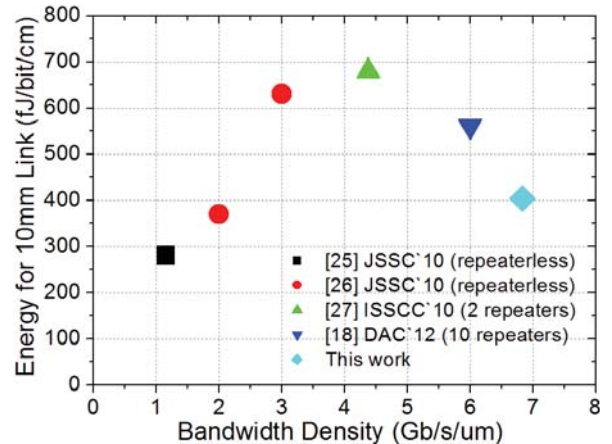

Fig. 8. 1cm LT energy versus bandwidth density of the proposed on-chip signaling and previous works.

bandwidth density characteristics of the SRLR-based link and other silicon-proven on-chip interconnects [18], [25]–[27]. Details of the fabricated test link are summarized in Table I together with the previous works.

The on-chip bias circuit for an adaptive swing voltage scheme consumes 587μW and it can be shared by all parallel links at a NoC router. When considering a 64bit 10mm link implementation, the bias circuit dissipates just 0.6% of total link power.

To compare the power consumption and area of our SRLR-based datapath with those of an entire router, we synthesized a typical mesh router (64bits, 5ports, 4VCs, and 16 buffers) in the same process, 45nm SOI CMOS. Extracted simulation results showed that input buffers and control logic consume 38.8mW and 5.2mW respectively, while our low-swing datapath consumes 12.9mW. Area wise, as discussed in Section I, our low-swing datapath occupies 18% of the overall router footprint.

V. CONCLUSION

This work proposes a self-resetting logic repeater (SRLR) for a low-overhead, robust low-swing signaling datapath of a mesh NoC. The SRLR optimized for the router-to-router dis-

TABLE I
COMPARISON OF SILICON-PROVEN ON-CHIP INTERCONNECTS.

	[25]	[26]	[27]	[18]	This Work
Signaling Type	fully differential	fully differential	fully differential	fully differential	single-ended
Data Rate	2Gb/s	(4Gb/s), 6Gb/s	4.9Gb/s	5.4Gb/s	4.1Gb/s
Bandwidth Density	1.163Gb/s/μm	(2Gb/s/μm), 3Gb/s/μm	4.375Gb/s/μm	6.0Gb/s/μm	6.83Gb/s/μm
Energy for 10mm Link Traversal (LT)	340fJ/bit/cm (repeaterless)	(370fJ/bit/cm), 630fJ/bit/cm (repeaterless)	340 X 2 = 680fJ/bit/cm (2 repeaters)	56.1 X 10 = 561fJ/bit/cm (10 repeaters)	404fJ/bit/cm (10 repeaters)
Process Technology	90nm bulk CMOS	90nm bulk CMOS	90nm bulk CMOS	45nm SOI CMOS	45nm SOI CMOS

⋆ Higher bandwidth density (*i.e.,* smaller wire spacing) incurs larger wire coupling capacitance, resulting in higher energy consumption. Thus, the energy consumption of on-chip interconnects should be considered along with their bandwidth density as shown in Fig. 8.

⋆ CMOS process scaling does not provide much energy benefit for on-chip signaling circuits since the load capacitance of on-chip interconnects is mostly given by their long wire capacitance (not by the gate capacitance) [15].

⋆ [18] requires an additional power supply and its energy is evaluated assuming that the additional power supply has no charge-recycling circuits.

tance in a mesh NoC (*e.g.,* 1mm in this work) provides scalable on-chip signaling without the increased layout complexity. Since the SRLR enables single-ended low-swing pulses to be repeated without a reference clock, the SRLR-based on-chip signaling achieves higher energy efficiency than differential, clocked low-swing signaling circuits. This paper also presents circuit techniques to improve process variation immunity of the SRLR-based on-chip signaling.

ACKNOWLEDGMENT

REFERENCES

[1] W. J. Dally and B. Towles, "Route packets, not wires: On-chip interconnect network," *ACM/IEEE Design Automation Conference (DAC)*, pp. 684–689, June 2001.

[2] M. B. Taylor *et al.*, "The raw microprocessor: A computational fabric for software circuits and general-purpose programs," *IEEE Micro*, vol. 22, no. 2, pp. 25–35, 2002.

[3] P. Gratz, C. Kim, K. Sankaralingam, H. Hanson, P. Shivakumar, S. W. Kecker, and D. Burger, "On-chip interconnection networks of the trips chip," *IEEE Micro*, vol. 27, no. 5, pp. 41–50, 2007.

[4] S. Bell *et al.*, "TILE 64-Processor: a 64-core SoC with mesh interconnect," *IEEE Int'l Solid-State Circ. Conf. Dig. Tech. Papers (ISSCC)*, pp. 88–89, Feb. 2008.

[5] J. Howard *et al.*, "A 48-core IA-32 message-passing processor with DVFS in 45nm CMOS," *IEEE Int'l Solid-State Circ. Conf. Dig. Tech. Papers (ISSCC)*, pp. 108–109, Feb. 2010.

[6] Y. Hoskote, S. Vangal, A. Singh, N. Borkar, and S. Borkar, "A 5-GHz mesh interconnect for a teraflops processor," *IEEE Micro*, vol. 27, pp. 51–61, Sep.-Oct. 2007.

[7] W. J. Dally and B. Towles, *Principles and Practices of Interconnection Networks.* Morgan Kaufmann Publishers, 2004.

[8] A. Kumar, L.-S. Peh, P. Kundu, and N. K. Jha, "Express virtual channels: Towards the ideal interconnection fabric," in *Int'l Symp. on Computer Architecture (ISCA)*, Jun. 2007.

[9] A. Kumar, L.-S. Peh, and N. K. Jha, "Token flow control," in *Int'l Symp. on Microarchitecture (MICRO)*, Nov. 2008.

[10] T. Krishna, L.-S. Peh, B. M. Beckmann, and S. K. Reinhardt, "Towards the ideal on-chip fabric for 1-to-many and many-to-1 communication," in *Int'l Symp. on Microarchitecture (MICRO)*, Dec. 2011.

[11] M. Hayenga, N. E. Jerger, and M. Lipasti, "Scarab: A single cycle adaptive routing and bufferless network," in *Int'l Symp. on Microarchitecture (MICRO)*, Dec. 2009.

[12] T. Moscibroda and O. Mutlu, "A case for bufferless routing in on-chip networks," in *Int'l Symp. on Computer Architecture (ISCA)*, Jun. 2009.

[13] N. Salehi, A. Khademzadeh, and A. Dana, "Minimal fully adaptive fuzzybased routing algorithm for networks-on-chip," *IEICE Electronics Express*, vol. 8, no. 13, pp. 1102–1108, 2011.

[14] H. Zhang, V. George, and J. M. Rabaey, "Low-swing on-chip signaling techniques: Effectiveness and robustness," *IEEE Transactions on Very Large Scale Integration Systems (T-VLSI)*, pp. 264–272, June 2000.

[15] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic, "DSENT - a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling," *International Symposium on Networks-on-Chip (NOCS)*, May 2010.

[16] Jan M. Rabaey, Anantha P. Chandrakasan, and Borivoje Nikolic, "Digital Integrated Circuits: A design perspective," *Prentice Hall, 2nd Edition*, 1998.

[17] T. Krishna, J. Postman, C. Edmonds, L.-S. Peh, and P. Chiang, "SWIFT: A SWing-reduced Interconnect For a Token-based Network-on-Chip in 90nm CMOS," *IEEE International Conference on Computer Design (ICCD)*, pp. 439–446, 2010.

[18] S. Park, T. Krishna, C.-H. O. Chen, B. Daya, A. P. Chandrakasan, and L.-S. Peh, "Approaching the theoretical limits of a mesh NoC with a 16-node chip prototype in 45nm SOI," *ACM/IEEE Design Automation Conference (DAC)*, June 2012.

[19] H. Zhang, V. George, and Jan M. Rabaey, "Low-Swing On-Chip Signaling Techniques: Effectiveness and Robustness," *IEEE Transactions on Very Large Scale Integration Systems (T-VLSI)*, vol. 8, pp. 264–272, 2010.

[20] E.D. Kyriakis-Bitzaros, "Design of low power CMOS drivers based on charge recycling," *IEEE International Symposium on Circuits and Systems*, pp. 1924–1927, June 1997.

[21] M. Hiraki et al, "Data-Dependent Logic Swing Internal Bus Architecture for Ultralow-Power LSIs," *IEEE Journal of Solid-State Circuits (JSSC)*, pp. 397–402, April 1995.

[22] H. Yamauchi et al, "An Asymptotically Zero Power Charge-Recycling Bus Architecture for Battery-Operated Ultrahigh Data Rate ULSIs," *IEEE Journal of Solid-State Circuits (JSSC)*, pp. 423–431, April 1995.

[23] R. Golshan et al, "A novel reduced swing CMOS BUS interface circuit for high speed low power VLSI systems," *IEEE International Symposium on Circuits and Systems*, pp. 351–354, May 1994.

[24] B.-D. Yang et al, "High-Speed and Low-Swing On-Chip Bus Interface Using Threshold Voltage Swing Driver and Dual Sense Amplifier Receiver," *European Solid-State Circuit Conference (ESSCIRC)*, pp. 144–147, September 2000.

[25] E. Mensink, D. Schinkel, E. A. M. Klumperink, E. van Tuijl, and B. Nauta, "Power efficient gigabit communication over capacitively driven RC-limited on-chip interconnects," *IEEE Journal of Solid-State Circuits (JSSC)*, vol. 45, pp. 447–457, Apr. 2010.

[26] B. Kim and V. Stojanovic, "An energy-efficient equalized transceiver for RC-dominant channels," *IEEE Journal of Solid-State Circuits (JSSC)*, vol. 45, pp. 1186–1197, June 2010.

[27] J. sun Seo, R. Ho, J. Lexau, M. Dayringer, D. Sylvester, and D. Blaauw, "High-bandwidth and low-energy on-chip signaling with adaptive pre-emphasis in 90nm CMOS," *IEEE Int'l Solid-State Circ. Conf. Dig. Tech. Papers (ISSCC)*, pp. 182–183, Feb. 2010.

[28] Y.-H Kao and M. Yang and N. S. Artan and H. J. Chao, "CNoC: High-Radix Clos Network-on-Chip," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (T-CAD)*, vol. 30, pp. 1897–1910, 2011.

[29] W. J. Dally, "Express cubes: improving the performance of k-ary n-cube interconnection networks," *IEEE Transactions on Computers*, vol. 40, pp. 1016–1023, 1991.

[30] H. J. Oguey and D. Aebischer, "CMOS current reference without resistance," *IEEE Journal of Solid-State Circuits (JSSC)*, vol. 32, pp. 1132–1135, July 1997.