

Networks-on-Chip With Double-Data-Rate Links

Anastasios Psarras, Savvas Moisisidis, Chrysostomos Nicopoulos, and Giorgos Dimitrakopoulos

Abstract—The need for higher throughput and lower communication latency in modern networks-on-chip (NoC) has led to low- and high-radix topologies that exploit the speed provided by on-chip wires—after appropriate wire engineering—to transfer flits over longer distances in a single clock cycle. In this paper, motivated by the same principle of fast link traversal, we propose the RapidLink NoC architecture, which exploits said speed to rapidly transfer flits between adjacent routers using double-data-rate (DDR) link traversals. RapidLink is enhanced with novel low-cost DDR elastic buffers that pipeline link traversal (when needed) to multiple flow-controlled half-cycle segments, whereby each segment is driven with data on both the positive and negative edges of the clock. DDR link traversal leads to multiple NoC configurations that can markedly increase network performance without increasing the area/power cost of the NoC relative to state-of-the-art single-data-rate architectures. Extensive cycle-accurate network simulations and hardware implementation results demonstrate the efficiency of RapidLink and its potential as a scalable NoC architecture.

Index Terms—Network-on-chip, double data rate, elastic buffers.

I. INTRODUCTION

NETWORKS-on-Chip (NoC) have been established as the dominant communication backbone in multi-core environments, primarily due to their innate scalability attributes. In modern scaled systems that add extra Intellectual Property (IP) cores in each generation, the effective network throughput per core decreases (due to elevated cross traffic), while the average source-to-destination hop count increases (due to increasing network diameter). To sustain scalability into the many-core realm, it is imperative to provide NoC architectures with higher throughput and lower latency, *without* incurring any power/area penalties.

Overcoming NoC scalability limitations has led to either high-radix networks [1]–[3] with long connecting links and complex routers, or to networks that allow flits to traverse multiple network hops of shorter wires in a single clock cycle (through router bypassing) [4], [5]. The effectiveness of the aforementioned solutions relies on two key attributes: (a) on their fundamental property of fast link traversal, i.e., transferring flits over longer distances in a single

clock cycle; and (b) on the adoption of highly efficient pipelined router organizations [6], [7] that also include fine-grained pipeline-stage-bypassing capabilities [8], thus effectively reducing packet latency and increasing packet delivery rate.

Fast link traversal, although possible for reasonable link lengths, cannot be achieved at high clock frequencies, without appropriate wire engineering [9], [10], i.e., (a) promoting NoC links to upper metal layers, (b) increasing the wire spacing, or using wire shielding, and (c) using across-wire repeaters. Such design decisions fit well within the NoC's physical properties, since the NoC links are mostly routed using intermediate metal layers that are neither too resistive, nor too dense. This is a natural choice since metal layers reserved for local routing are used by the processing cores and their caches, while the top metal layers (with significantly lower resistance) are primarily occupied by power and clock signals [1], [11], [12]. As measured in [1], and used in a real prototype at 32 nm, the wire delay of the group of metal layers used for NoC links ranges from 60 to 300 ps/mm, depending on repeater placement and wire spacing.

Similar results have been shown by a variety of real prototypes. For example, IBM has shown that, with appropriate wire spacing and metal layer selection, wires can cross distances of up to 2.7 mm in 210 ps at 45 nm [13], while, at the same technology node, Intel drives a wire of 5.4 mm in 270 ps [7]. Recently, SMART [4], [14] was demonstrated to traverse 16 mm of wire (16 hops of 1 mm each) at 1 GHz, by utilizing 1-mm-spaced repeaters and 3× larger wire spacing than the minimum allowed. This translates into crossing 4 mm in less than 250 ps. Similarly, NoCs designed recently with high-radix routers assume repeated wire delays of 66 ps/mm, which are used to cross 5.4 mm long links in a single cycle [2].

All aforementioned approaches are mostly applicable to *tile-based* Chip Multi-Processors (CMP), whereby the system comprises a number of identically-sized logic blocks, called tiles, organized in a regular 2D layout [15]. Since the overall size of a CMP die tends to remain constant across different processor generations (due to yield and cost issues), any increase in the number of on-chip tiles is inevitably accompanied by a corresponding decrease in the size of each tile, which translates to scaled-down inter-tile link distances. This effect is illustrated in Fig. 1, which highlights the decreasing length of the network links, as CMP cores scale from 16 to 64. Hence, achieving fast wire traversal speeds in *scaled-down* links (i.e., wire lengths that scale *down* as a result of decreasing tile dimensions) is more easily attainable, after also taking into account the fact that clock frequencies are only modestly increasing to keep power consumption under control.

Manuscript received April 1, 2017; revised June 25, 2017; accepted July 26, 2017. Date of publication August 14, 2017; date of current version November 22, 2017. This paper was recommended by Associate Editor D. Zito. (Corresponding author: Giorgos Dimitrakopoulos.)

A. Psarras, S. Moisisidis, and G. Dimitrakopoulos are with the Electrical and Computer Engineering Department, Democritus University of Thrace, 67100 Xanthi, Greece (e-mail: apsarra@ee.duth.gr; smoysidi@ee.duth.gr; dimitrak@ee.duth.gr).

C. Nicopoulos is with the Electrical and Computer Engineering Department, University of Cyprus, 1678 Nicosia, Cyprus (e-mail: nicopoulos@ucy.ac.cy).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSI.2017.2734689




Fig. 1. Under equal die size, as the number of tiles increases, the inter-tile link length decreases. This is illustrated here by comparing a 16-core CMP (left) to a 64-core CMP (right).

On the contrary, it becomes increasingly difficult to sustain single-cycle Link Traversal (LT) for general-purpose – and primarily *heterogeneous* – Multi-Processor Systems-on-Chip (MPSoC), with irregular physical layout and longer links.

Hence, we need a method that can leverage fast wire traversal for increased performance on both (a) *scaled-down link lengths*, when dealing with tiled CMPs, and (b) on *longer links*, where the link delay *cannot* fit inside one clock cycle.

To achieve this goal, we exploit an inherent idiosyncrasy of NoC architectures to expose a new design opportunity that can embrace both scalability flavors. Specifically, we harness the intrinsic asymmetry between the *intra-* and *inter-*router delays encountered in modern multi-core systems. For reasonable inter-router distances of up to a few millimeters, the delay of LT is substantially shorter than the delay of the routers. Currently, the fastest state-of-the-art NoC routers for 2D meshes exhibit intra-router delays ranging from 600 ps to 1000 ps [1], [4], [7], when measured at voltages of around 0.8 V, i.e., more than $2\times$ longer than the typical inter-tile link delay for scaled tile dimensions.

Rather than using the fast link traversal to cover longer distances in a given cycle, the proposed *RapidLink* architecture exploits said speed to *rapidly transfer flits* between adjacent routers connected with links of reasonable length in *half a clock cycle* and *utilizes both edges of the clock* during the sending and receiving of flits. As a result, each upstream/downstream router pair benefits from Double-Data-Rate (DDR) transfers, whereby two flits can be sent/received per clock cycle. Half-cycle single-data-rate link traversal has been exploited in the past for low power in [16], and for maximizing timing safety [17], [18].

In RapidLink, the original clock frequency of the NoC is unaffected, and the NoC routers do *not* need to run any faster than their normal operating frequency. The only constraint is that the Link Traversal (LT) delay cannot exceed one half of the delay of the router, which – as previously mentioned – is feasible for small/medium wire lengths, after appropriate wire engineering. However, RapidLink can also handle the cases where the link delay *cannot* fit within half of a clock cycle. In those cases, RapidLink “fragments” the link into multiple DDR half-cycle segments using newly introduced novel DDR **Dual-Stream Elastic Buffers (DS-EBs)**. Said elastic buffers act as flow-controlled DDR pipeline registers. In this way, the benefits of DDR link traversal can be reaped by *any* NoC design, even those with long wires and increased link delays.

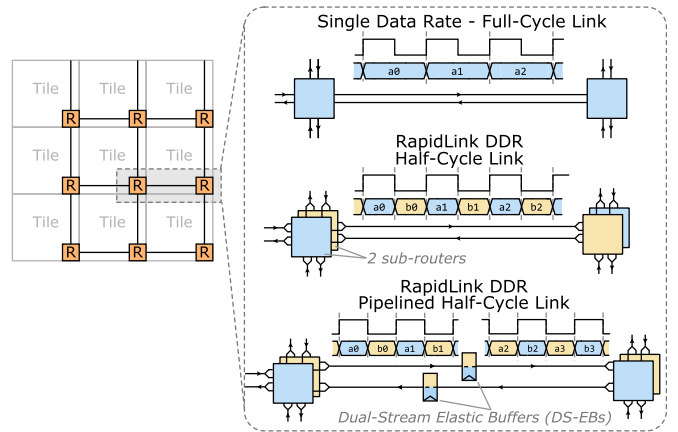


Fig. 2. A high-level overview of the organization of a Network-on-Chip with Single- and Double-Data-Rate links. The proposed RapidLink architecture enables DDR inter-router transfers on links of any length.

The RapidLink architecture and the new DDR DS-EBs are exploited to provide multiple NoC configurations that can be applied to NoC designs supporting Virtual Channels (VC), or to simplified single-stream (aka wormhole) NoC designs. In all cases, the adoption of RapidLink can markedly increase throughput with substantially lower power consumption as verified by extensive cycle-accurate network simulations, and detailed hardware analysis using placed-and-routed designs at 45 nm technology.

The rest of the paper is organized as follows: Section II presents RapidLink, and the design of the novel low-cost DDR DS-EBs. Section III describes how RapidLink (an inherently dual-stream architecture) can handle a single-stream NoC configuration, while Section IV presents the experimental results. Finally, conclusions are drawn in Section V.

II. RAPIDLINK NOCS WITH DDR LINKS

Traditional NoC design assumes – at each node of the network – a single router that operates in a single cycle, or in multiple cycles (when its operation is pipelined) [19]. Typically, the transfer of flits across routers takes one full cycle, as depicted at the top right of Fig. 2. In the presence of longer links, link traversal may remain at one full cycle, after aggressive wire engineering, or it may be split into multiple cycles, using pipelining [20], for maximum safety.

RapidLink aims at providing Double-Data-Rate transmissions on the NoC links *without* constraining the delay of router traversal, which should remain a full-cycle operation. When the NoC links operate in DDR mode, it means that the sender and the receiver on each link should be able to send and receive flits at both the positive and negative edges of the clock. To enable this operation, the NoC routers can provide two separate send/receive paths to each inter-router link, where each path carries a **separate stream (flow) of data**. With RapidLink, these two separate streams (data flows) can be transferred in a *time-multiplexed* manner across the same inter-router link in DDR mode; one stream would “ride” the positive phase of the clock, while the other stream would “ride” the negative phase. This organization is abstractly

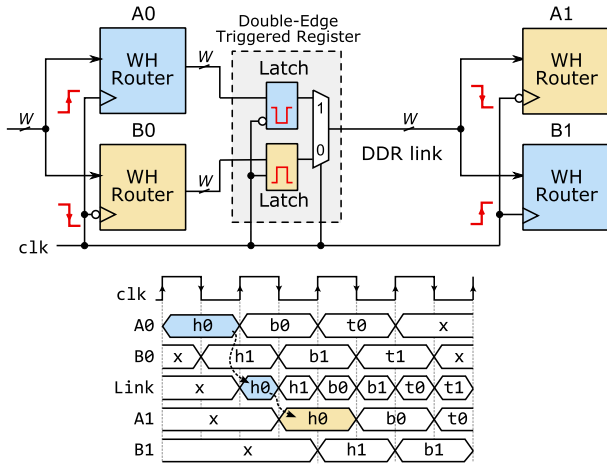


Fig. 3. RapidLink employs two full-cycle routers and DDR half-cycle links. Each node consists of two W -bit wormhole (i.e., single-stream) sub-routers that serve separate network streams (e.g., request and response message classes), and they time-share a W -bit link using both edges of the clock.

illustrated in Fig. 2, where a pair of sub-routers per node drive a DDR link. Data may reach the next node in one half-cycle of the clock, or after 2 (or more) half-cycles, when the link needs to be pipelined for timing closure. The detailed operations of RapidLink and all its constituent parts are explained below.

A. RapidLink With DDR Half-Cycle Links

Each router in the baseline RapidLink architecture consists of *two full-cycle wormhole sub-routers* of W -bits each, as depicted in Fig. 3. The output of the two sub-routers is time-shared on a W -bit inter-router link using both edges of the clock. For half a clock period (e.g., during the positive phase of the clock), the link is traversed by flits of sub-router 0, while during the other half of the clock period (e.g., the negative phase), flits of sub-router 1 use the link. The two sub-routers operate locally on opposite clock edges and they connect to sub-routers with opposite clock edges downstream, due to the half-cycle link traversal.

The output data of each sub-router is fed into a shared Double-Edge Triggered (DET) register. Each DET register consists of two latches placed in parallel, which are enabled on opposite phases of the clock, and an output multiplexer driven by the clock signal [21]. Thus, DET registers incur marginal overhead, as compared to generic single-edge-triggered registers, which are also built using two latches (master and slave latches placed in series). The clock signal driving the multiplexer of the DET register is appropriately gated when no new valid flits arrive from any of the sub-routers, thus preventing unnecessary switching activity on the DDR link.

Fig. 3 also illustrates the activity on the link, as flits flow from routers A0/B0 to routers A1/B1. On the positive edge of cycle 0, flit ‘h0’ enters router A0, while, on the negative edge of the same cycle, the ‘h1’ flit is written in the input buffer of router B0. On the next positive edge (cycle 1), the A0 latch stores flit ‘h0,’ which has completed a whole cycle inside the router and moves directly to the link. Half a cycle later, the flit reaches router A1 and is captured on the positive clock edge.

At the same time, ‘h1’ appears on the link, and the same pattern continues in the following cycles.

By construction, RapidLink serves two streams of data that traverse the network in a time-multiplexed fashion using distinct sub-routers of alternating clock edges. Thus, effectively, RapidLink implements two distinct sub-networks that remain in isolation, i.e., moving data from the first sub-network to the second sub-network is impossible, due to the physical separation imposed by clock-edge interleaving. Equivalently, this organization resembles a NoC with two Virtual Channels (VCs), where the multiplexing of the flits of each VC is statically determined by the phases of the clock, and exchanging traffic across VCs (i.e., allowing the packet of one VC to transfer to another VC) is not possible.¹

Therefore, the effectiveness and the hardware complexity of RapidLink should be judged relative to a network that hosts two VCs in one network and comprises VC-based routers (supporting 2 VCs) with inputs and outputs of W bits. The combined area/power budget of the two sub-networks of RapidLink (i.e., two sub-routers per node) is expected to be lower – or equal to – the area/power of a network that supports two VCs. This property is attributed to the fact that the two sub-routers of RapidLink are both faster – due to the simplification of parts of the allocation and multiplexing logic – and more area-efficient under equal delay.

B. RapidLink With Pipelined Half-Cycle links Using Dual-Stream Elastic Buffers (DS-EB)

The data that travels on the DDR link should reach the next router in one half of a clock cycle. This requirement can, perhaps, be satisfied in scaled *tilde* CMPs, after appropriate wire engineering. However, in the general case, fast wire traversal may *not* always be possible. In such cases, the link should be pipelined using dual-edge triggered registers [22], thereby splitting the link into multiple half-cycle LT steps. Instead of adding mere pipeline registers, we propose a new approach. For the first time – to the best of our knowledge – we propose the use of DDR *elastic* buffers, similar in vein to traditional single-stream, single-data-rate EBs [23]–[25]. The newly proposed **DDR Dual-Stream Elastic Buffers (DS-EB)** can both (a) split the timing paths on the link, and (b) enable the flow-controlled transfer of data across routers, effectively converting the pipelined channel into a distributed DDR FIFO queue.

The proposed DDR DS-EB replaces the DET flip-flop at the sub-routers’ outputs, in order to control the data flowing through the two streams. It operates under a minimal elastic protocol and can be used in a plug-and-play manner to pipeline the DDR link (e.g., to meet timing constraints) in any number of *buffered* stages. As shown in Fig. 4, each dual-stream DDR link consists of a data bus and two pairs of handshaking signals (one for each stream), which control data transmissions occurring at the positive and negative levels of the clock. For each stream, a forward valid signal indicates whether the data

¹Completely separating the traffic of different VCs is a useful property in NoCs when VCs are used to implement different message classes that need to remain separate (e.g., request/response traffic).

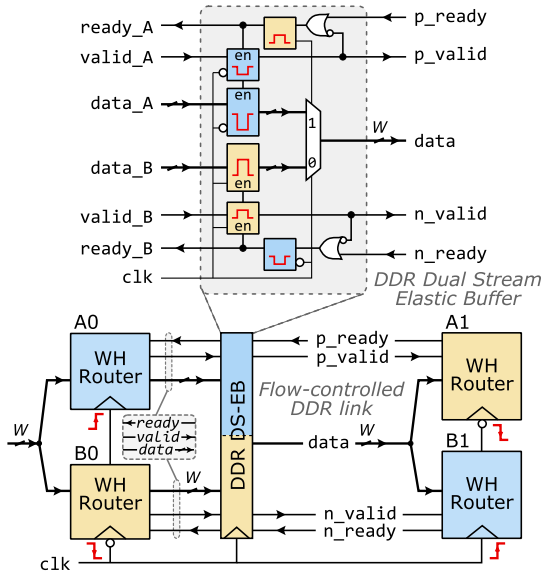


Fig. 4. The organization of the proposed DDR DS-EB (upper part) and its integration at the output of a RapidLink pair of sub-routers (lower part).

bus contains valid data, while the backward ready is used by the receiver to indicate when it is ready to accept new data. Whenever valid and ready are asserted, data transmission was successful during that phase of the clock; in all other cases, no transfer occurred.

The organization of the elastic buffer, seen in the upper part of Fig. 4, uses two independently-enabled data latches, latches for the valid bits, and a multiplexer controlled by the clock signal. Each latch captures the data sent by the two opposite-clock-edged sub-routers. The two latches of each stream are controlled by one more latch that captures the backward ready signal, generated by the input buffer of the downstream router to indicate buffer availability. The latched ready bit acts as an enable signal for the data and valid latches, retaining data whenever the receiving buffer could not capture the last transfer (or, if no transfer occurred in the previous cycle). Note that the receiving sub-router will capture the arriving data on the *opposite* clock edge and, thus, the backward ready signal is captured by a latch enabled by an opposite clock level relative to the corresponding valid and data bits.

The operation of the DDR DS-EB is demonstrated in the running example of Figure 5, which illustrates the cycle-by-cycle activity of a DDR flow-controlled link between two nodes. Two streams P and N are multiplexed on the DDR link. During the first cycle, neither of the two streams is utilizing the DDR link, as both valid signals are low. Flit ‘P1’ appears in the first cycle at the output of the positive-edged sub-router and reaches the link during the following positive clock phase (p_valid is asserted). The transfer of flit ‘P1’ was successful, since p_ready was sampled high on the previous falling edge of the clock. During the negative clock phase, flit ‘N1’ is sent through the link and flit ‘P2’ is transferred during the following positive clock phase. However, in the next negative phase, the link remains idle, since n_valid is de-asserted. Next, the transfer for the positive stream fails (flit ‘P3’), since the receiver’s p_ready was de-asserted on

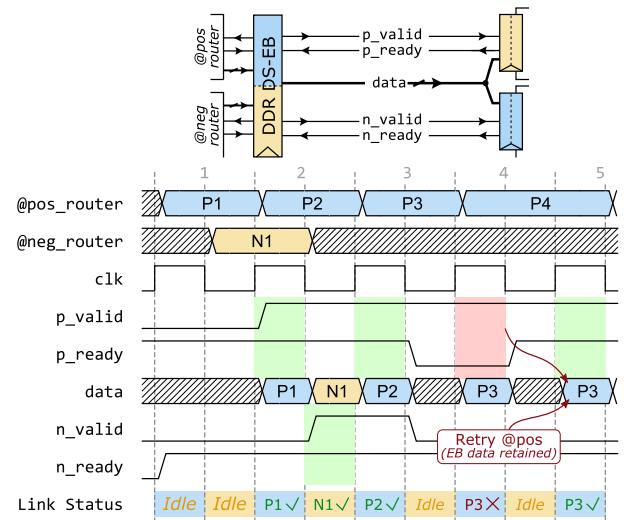


Fig. 5. Cycle-by-cycle activity of the flow-controlled DDR link. The two data streams P and N are sharing the DDR link in alternating clock phases, and their transfers are controlled by the corresponding ready and valid handshake.




Fig. 6. Multiple DDR DS-EBs may be placed in series – acting as a distributed DDR FIFO – to pipeline the DDR link. Neighboring DS-EBs should be driven by alternating clock phases.

the previous falling clock edge. Therefore, the latched ready signal of the DS-EB is preventing the valid and data latches of the P stream to be updated; otherwise, the stored flit (‘P3’) would be overwritten by the incoming flit (‘P4’). In the negative clock phase, flit ‘P3’ cannot retry its transmission, since this phase is reserved for stream N, and the link remains idle again. Finally, as the clock phase switches, stream P is allowed to retry (p_ready is high) and ‘P3’ is actually transferred successfully.

The DDR link can be pipelined by placing multiple DS-EBs in series, as depicted in Fig. 6. To achieve half-cycle traversal across each link segment, and full throughput DDR transmission, consecutive DS-EBs should use alternate clock phases.

C. Integrating RapidLink With Single-Data-Rate Network Interfaces

In RapidLink, the DDR operation of the NoC links does *not* affect the full-cycle (single-edge) operation of the NoC routers, nor the Network Interfaces (NIs). Each NI can safely assume an injection/ejection throughput of at most 1 flit/full-cycle/NI, as in any single-edge-triggered baseline NoC. Although this feature simplifies RapidLink’s integration, since no NI modifications are required, it introduces a data rate mismatch between the NIs and the DDR operation of the NoC links.

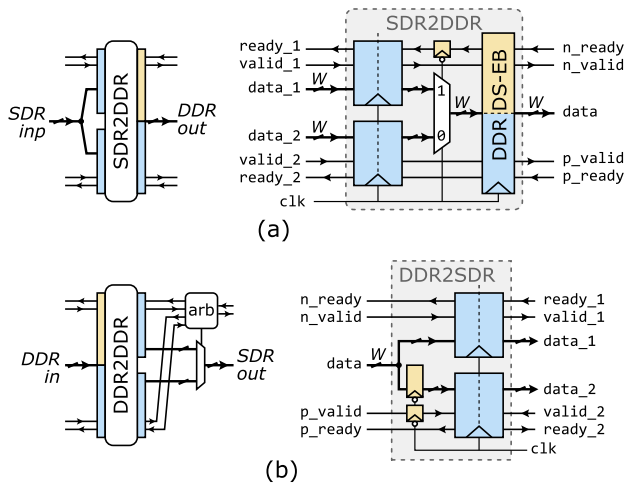


Fig. 7. (a) The SDR2DDR bridge for stream merging (at the injection point), and (b) the DDR2SDR bridge for stream splitting (at the ejection point). Both bridges connect RapidLink’s DDR inputs and outputs to single-rate network sources and sinks, respectively, which can handle messages from two distinct streams (flows).

To interface between the two domains with different data rates, while preserving RapidLink’s ease of integration, we provide two lightweight *bridge modules* that act as “glue logic” between the injection and ejection points of the network.

Fig. 7 illustrates the architecture of the bridges used at the injection and ejection points of the network. At the injection points, we should *merge* two independently flow-controlled data streams that operate on their own ready/valid handshake – driven on the positive clock edge – to a dual-stream DDR interface. Merging is performed by the SDR2DDR bridge, as illustrated in Fig. 7(a). The injected flits are first buffered into one of the positive-edge-triggered elastic buffers, depending on the stream they belong to. From there, they move to the DDR DS-EB that will guide them to the appropriate sub-router of the network, through a multiplexer that is controlled by the phase of clock. In order for the handshake signals to be sampled properly, the *n_ready* signal of the DDR DS-EB is re-timed before being transferred to the positive-edge-triggered EB of the injection interface.

At the ejection points, we need to perform the opposite operation and *split* the two data streams coming out of RapidLink on opposite clock edges, to two independently flow-controlled streams that operate on a single clock edge. This splitting is performed by the DDR2SDR bridge, shown in Fig. 7(b), which de-multiplexes the DDR stream back to two separate single-data-rate streams by placing the flits ejected from the network into two parallel positive-edge-triggered EBs. The positive-edge-triggered EBs capture data only once per clock cycle, while the DDR input provides new data twice per clock cycle. Therefore, one data item per clock cycle will be lost and not captured by any of the two EBs. To resolve this issue, we add a re-timing stage for *data* and *p_valid* bits, in order for the data of the P stream to wait safely for the next positive clock edge, and not be overwritten by the new data of the N stream that appears in the middle of the cycle. Finally, once the two streams are separated at the output of

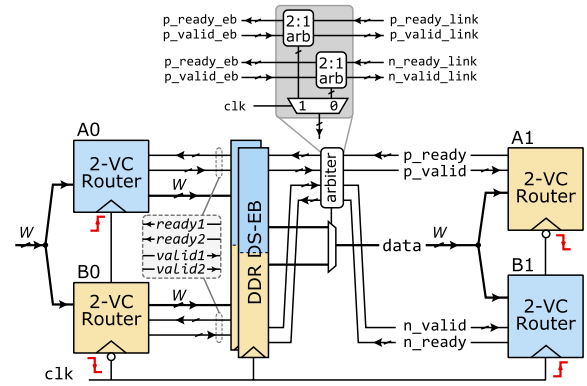


Fig. 8. The RapidLink organization supporting a total of $V = 4$ VCs. Multiple DS-EBs are placed in parallel at the output of the RapidLink sub-routers to support $V/2$ VCs per clock-phase.

the DDR2SDR bridge, they need to arbitrate to gain access to the exit of the network.

D. Supporting Multiple VCs

The baseline RapidLink design supports the transfer of two independent data streams flowing in separate sub-networks, each one synthesized using simplified wormhole sub-routers, i.e., routers without VC support. Nevertheless, the baseline architecture can be extended to support more than two flows, either for implementing more complex protocols that require more than two separate message classes, or to increase performance.

RapidLink can support V VCs in total, by assigning half of them to each sub-network. A network node consists of two sub-routers, each one hosting $V/2$ VCs, as shown in Fig. 8, for the case of a 4-VC configuration. In each clock cycle, at most one flit may appear at each sub-router output, according to typical VC-based router architectures [19]. When a flit tries to reach the output, it is stored into one of the $V/2$ DS-EBs that are placed in parallel at every output port. A one-hot valid vector encodes the flit’s VC id, indicating the DS-EB where the flit will be stored. For the 4-VC configuration of Fig. 8, each sub-router serves 2 VCs, providing 2-bit valid outputs to the 2 DDR DS-EBs.

Driving the flow-controlled DDR link with only one flit per clock phase is accomplished by arbitrating and multiplexing among the DS-EBs that currently host valid flits [26]. Two $V/2$ -input arbiters serve the two streams and determine which one of the $V/2$ VCs can use the link during the positive and negative clock phases, respectively. Arbiter requests are only made by the VCs whose associated DS-EB valid bits are asserted. As shown in Fig. 8, the select signal of the data multiplexer used to drive the link is the corresponding arbiter’s grant, depending on the current phase of the clock. Data reaches the receiver, along with a one-hot valid vector per stream that indicates the input VC buffer where the flit must be stored within the downstream router. Similarly, each downstream sub-router generates a $V/2$ ready vector to indicate buffer availability of each VC buffer. At the end of the active clock phase, if the granted VC’s incoming ready signal

was asserted, the transmission was successful. Otherwise, the stream can retry in its next clock phase.

Supporting multiple VCs (say V VCs) inevitably incurs the cost of $V/2$ DS-EBs at the output of each RapidLink node, and $V/2$ DS-EBs for each stage of a pipelined inter-router link (if link pipelining is required). However, the cost of a DS-EB is quite low, since its design is latch-based. Thus, the cost of the multiple DS-EBs needed to support multiple VCs is by no means prohibitive. In fact, it will be shown in Section IV-B that the hardware cost of a RapidLink NoC supporting 4 VCs in the presence of long pipelined links is lower than that of existing state-of-the-art NoC architectures.

III. SINGLE-STREAM RAPIDLINK

The baseline RapidLink architecture supports, by construction, the transfers of two independent data streams on the NoC links. In most systems, supporting two independent data streams is a necessity, since the system operation is based on a higher-level transaction protocol that requires the use of separate message classes. These message classes can range from simple request/response traffic to ones encountered in more complex cache-coherence protocols. However, if multiple message classes, or VCs, are *not* required, we need to identify a way to enable RapidLink to operate on a *single message stream per network source*, even if it is inherently constructed to support two message streams per source.

A. Concentrated RapidLink

Each network source of RapidLink, as shown in Fig. 7(a), injects in the network two independent message flows that are multiplexed inside RapidLink on opposite clock edges. Equivalently, we can replace the two message flows of one source with the traffic originating from two single-message-flow sources. To achieve this, we connect to the input of an SDR2DDR bridge two network sources, as shown in Fig. 9(a). This network-source merging in the SDR2DDR bridges is equivalent to network concentration [27], where two source/sink nodes of the original network are mapped to a single RapidLink node.

In this way, by time-multiplexing the two concentrated flows on DDR links, the network diameter is effectively reduced, without any implementation overhead, since neither the router population, nor the router radix are increased. An example of the method's applicability is presented in Fig. 9(c) in a 4×4 2D mesh serving 16 Processing Elements (PEs; e.g., CPUs), with each one connected to the network through a single in/ejection port. In a concentrated RapidLink (right-hand side of Fig. 9(c)), each vertical pair of nodes is merged into one that operates in DDR mode. This leads to a "folding" of the original single-data-rate 4×4 2D mesh onto a 2×4 mesh with DDR links.

Equivalently, at the ejection points, shown in Fig. 9(b), the DDR RapidLink stream is split in two flows using the DDR2SDR bridge. Since we cannot guarantee that the two split flows are not destined to the same output, an extra arbitration and switching stage is employed prior to ejection,

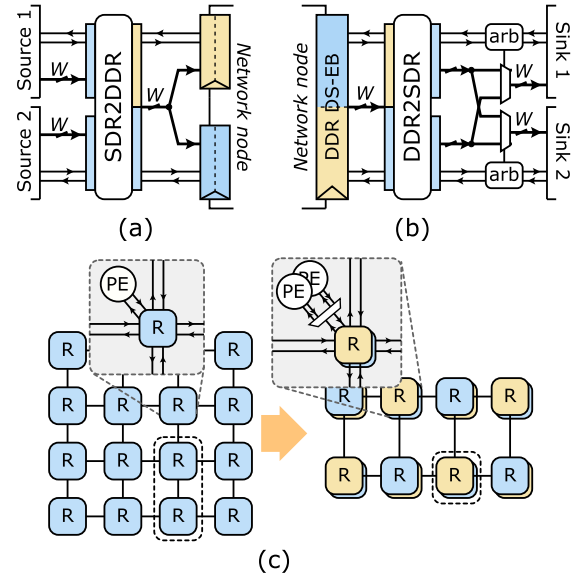


Fig. 9. (a) The injection and (b) ejection network interfaces that allow two single-stream sources and sinks to be attached to a *concentrated* RapidLink node. (c) A 4×4 2D mesh (left) is folded onto a 2×4 mesh (right) with DDR links, implemented with 5-port RapidLink routers, where each sub-router serves two Processing Elements (PEs).




Fig. 10. (a) The organization of a RapidLink node that employs master and slave sub-routers operating in *lockstep mode*. The organization of the (b) injection and (c) ejection interfaces that split the packets of a single stream to two independently flow-controlled streams driven on opposite clock edges.

in order to guarantee that only one stream can have access to each of the two connected sink points.

B. Lockstep-Mode RapidLink

Instead of merging two single-stream network sources to produce the dual DDR stream required by the baseline RapidLink (as described in the previous sub-section), we can produce the two streams by *splitting the packets of a single-stream in half*. In this configuration, flits entering the network

are split in two, and the two halves (of $W/2$ bits each) are injected in consecutive opposite clock edges into the two narrower sub-networks.

This organization is shown in Fig. 10. The two halves of the same packet travel “tied” together, appearing on the DDR links one after the other on consecutive (alternating) clock levels. To achieve this, the two sub-routers make sure that only one of the two halves is competing with other flows, arbitrating and allocating resources; the other one, always follows. To achieve this behavior, the two sub-routers are not treated equally. The “master” sub-router preserves the original switch architecture, including routing and allocation logic, and is responsible to perform arbitration among contending flows. The “slave” sub-router operates in *lockstep mode* with the master, and routes flits in *exactly the same way* as the master, by blindly copying its arbitration decisions. The switch configuration is transferred to the slave through a re-timing stage. In this way, a whole cycle is provided to the full-fledged router (the master), while a *half-cycle* path is enforced to the slave router, which only involves simpler multiplexing circuits.

The master and slave sub-routers must be set to alternating clock edges on consecutive nodes, so that the pre-pending flit always appears first on each node. Interfacing with the network requires re-wiring the W -bit input of the packet sources to the two $W/2$ -bit streams (on injection) and back (on ejection), as shown in Figs. 10(b) and (c), respectively. Note that the injection wiring must make sure that the half containing the header (e.g., flit type, destination, etc.) “rides” the clock phase that will follow the master sub-network. Between the source and the SDR2DDR bridge, extra control logic (“fork”) is inserted in the flow control signals, to make sure that both flit halves are injected simultaneously, without any idle cycles in-between. At the ejection points, similar logic (“join”) guarantees that the valid signal towards the sink is only asserted when both $W/2$ -wide outputs are valid.

IV. EVALUATION

In this section, we evaluate the performance of RapidLink and compare it under different topologies, in terms of network performance and hardware complexity, with state-of-the-art NoC architectures. In order to contain the number of possible configurations, we follow the design presented in the Scorpio chip [8]. Scorpio follows a tile-based chip floor-plan, as the one shown in Fig. 1, with a tile size of approximately 2×2 mm. Scorpio was built at 45 nm technology, which matches the technology library we have available for our implementation. The CMP consists of 64 tiles (nodes), the NoC is required to support 4 VCs, and the link width is set to 64 bits.

All designs were implemented in SystemVerilog. Latency and throughput measurements were derived from cycle-accurate network simulations, while the hardware cost evaluation is conducted after synthesizing the RTL models using a commercial 45 nm standard-cell library under worst-case conditions (0.8 V, 125°C), and performing placement-and-routing of the resulting designs using the Cadence digital implementation flow.

The networking performance evaluation involves four synthetic traffic patterns: Uniform Random (UR), non-uniform

Localized (LC) traffic, and two versions of permutation traffic: Bit-Complement (BC) and transpose (TS) traffic patterns. Under UR traffic, every node sends its packets to all other nodes of the network with equal probability. For LC traffic, we assume that 75% of the overall traffic is local (i.e., the destination is one hop away from the source), while the remaining 25% of the overall traffic is uniform-randomly distributed to the non-local nodes. The injected traffic consists of two types of packets to mimic realistic system scenarios: 50% of the packets are 1-flit short packets (just like request packets in a CMP), and 50% longer 5-flit packets (just like response packets carrying a cache line).

In addition to using purely synthetic traffic patterns, we also employ traffic patterns that are derived from real application workloads. Specifically, we employ the hot-spot traffic model from [28], which synthesizes traffic that closely resembles the traffic behavior of PARSEC application benchmarks [29] running on a CMP. Under this PARSEC-derived Hot-Spot (HS) traffic pattern, 20% of the nodes receive $50\times$ more traffic than the rest, while the remaining injected traffic is uniformly distributed to all other destinations. In order to mimic the behavior of real applications, the first two VCs (out of the 4 total VCs per input port) receive 77% of the traffic while the other two get 22% and 1% of the injected traffic, respectively [28]. Packet distribution is skewed, with 1- and 5-flit packets being 70% and 30% of the total packets injected, respectively.

For power measurements of the NoCs under comparison, we guaranteed that each architecture is driven by the same arriving packet sequence (uniform random traffic of 1-flit and 5-flit packets) under the same injection load. The power analysis is reported after taking into account all layout parasitics, while the switching activity has been computed using delay-accurate simulations of the derived logic-level netlists. The payload of each packet was produced using a uniform random generator, and the average data switching activity observed in all cases was 15%.

A. Networks With Medium-Length Inter-Router Links

In the first set of experiments, we assume that the 64 tiles of the CMP are connected using an 8×8 2D mesh network supporting XY routing; in this case the NoC links follow exactly the tile size, and they are 2 mm long.

RapidLink is compared with three state-of-the-art architectures, each one having different characteristics, thereby covering all the design space of possible NoC architectures. Table I summarizes the key hardware attributes of all architectures under evaluation.

The first design used in the comparisons corresponds to a NoC that employs single-cycle (“SC”) 4-VC routers employing combined allocation [19], [30] and full-cycle link traversals. In this case, the router’s delay limits the operating frequency of the NoC, which can operate at a frequency of 1.1 GHz, as shown in Table I. The single-cycle router uses 3 buffer slots/VC, as needed to cover the credit Round-Trip Time (RTT) in single-cycle routers with full-cycle links.

The second design – that aims at higher throughput by elevating the clock frequency – is an optimized 3-stage pipelined

TABLE I
HARDWARE CHARACTERISTICS OF ALL ARCHITECTURES UNDER EVALUATION FOR
A 64-NODE CMP USING MEDIUM-LENGTH LINKS IN AN 8×8 2D MESH NoC

Router Architecture (4 VCs – 64-bit links)	Clock Frequency	Area	Minimum Required Input Buffers	Link Latency	Router Power	Link Power	Total Power
Singe Cycle (SC)	1.1 GHz	3.07 mm ²	3 slots/VC	1 full cycle	260 mW	70 mW	330 mW
3-stage Pipelined (Pipe)	1.9 GHz	4.22 mm ²	5 slots/VC	1 full cycle	487 mW	68 mW	555 mW
3-stage Pipelined with Bypass (Bypass)	1.5 GHz	4.76 mm ²	5 slots/VC	1 full cycle	427 mW	69 mW	496 mW
RapidLink-Half-LT	1.0 GHz	3.08 mm ²	3 slots/VC	1 half cycle	186 mW	74 mW	260 mW
RapidLink-Full-LT	1.1 GHz	3.21 mm ²	3 slots/VC + 1 DS-EB (4VCs)	2 half cycles	204 mW	72 mW	276 mW

implementation [6] (“Pipe”), where each flit spends 3 cycles traversing a router, and 1 cycle on the link. In this case, the NoC can operate at 1.9 GHz, and is, again, limited by the router’s critical path and not the 2 mm NoC links. The pipelined implementation of the NoC routers increases the RTT across two nodes, and, thus, each router is obligated to have 5 buffer slots/VC, in order not to limit the link-level transmission throughput.

The last evaluated design matches exactly the design used in the Scorpio chip [8]. It uses 3-stage pipelined routers that employ a fine-grained pipeline-bypassing mechanism (“Bypass”). When flits do not experience any contention inside the router, they can bypass the corresponding pipeline stages and move to their requested output port within a single cycle. The departing flits then reach the downstream routers after one additional cycle. On the contrary, when a flit experiences contention within the router, it spends 3 full cycles in the router, and one cycle on the link. This architecture approximates the latency of a single-cycle router at low traffic loads, and the high throughput of a pipelined design at high traffic loads. The bypass paths increase the delay of the allocation/multiplexing logic, and, therefore, the best operating frequency for this design is 1.5 GHz. Again, in this case, the RTT is increased, and the routers employ 5 buffer slots/VC.

Note that the clock frequencies reported in this paper correspond to a low voltage of 0.8 V and roughly match the clock frequency of ultra-fast, 3-stage commercial routers [7], [31] when operated at 0.8 V.

For RapidLink, we evaluate two design options: “RapidLink-Half-LT” assumes half-cycle link traversals, while “RapidLink-Full-LT” assumes full-cycle link traversal that is split in two half-cycle DDR segments using DS-EBs in the middle of the link (as described in Section II-B). In both cases, the NoC routers of RapidLink (two pairs of 2-VC routers) are assumed to be single-cycle routers. This is possible, since RapidLink with DDR links creates two independent streams that can be interleaved on the link using the opposite edges of the clock. This allows for DDR operation on the links, while spending a full cycle in each sub-router. Since router traversal costs 1 cycle, the RTT is covered with 3 buffer slots/VC.

In RapidLink-Half-LT, the operating frequency of the NoC is limited by the wire delay and not the NoC routers. From our implementations, the repeated 2 mm wires impose a critical path of roughly 500 ps when routed on metal 6, including the




Fig. 11. Latency vs. load curves for a 8×8 2D mesh NoC under UR, LC, BC, and TS traffic patterns. (a) Uniform Random (UR). (b) Localized (LC). (c) Bit-Complement (BC). (d) Transpose (TS).

overhead of the registers on the two sides of the link, and an additional overhead due to the extra clocking uncertainty imposed when operating on both edges of the clock. Thus, in RapidLink-Half-LT, half a clock cycle cannot be less than 500ps which means that the NoC’s clock frequency cannot exceed 1 GHz.

Conversely, in RapidLink-Full-LT, the wire delay is not an issue, since it is distributed across the two half-cycle pipeline stages of the link. In this case, the NoC’s operating frequency is limited by the delay of the single-cycle routers. Even if each sub-router of RapidLink serves only 2 VCs and can run faster than the 1.1 GHz of the 4-VC single-cycle routers, we pessimistically assume that RapidLink-Full-LT also operates at 1.1 GHz to account for any additional clock uncertainty, due to the dual-edge clock operation of the elastic buffers.

The DDR link operation is expected to offer higher saturation throughput, with a minimal overhead in zero-load latency, relative to the fast pipelined NoC implementations. This behavior is, indeed, verified by the network performance results shown in Fig. 11. RapidLink configurations

TABLE II
HARDWARE CHARACTERISTICS OF ALL ARCHITECTURES UNDER EVALUATION FOR
A 64-NODE CMP USING LONG LINKS IN A 4×4 2D CONCENTRATED MESH NoC

Router Architecture (4 VCs – 64-bit links)	Clock Frequency	Area	Minimum Required Input Buffers	Link Latency	Router Power	Link Power	Total Power
Single Cycle (SC)	1.0 GHz	1.57 mm ²	5 slots/VC	2 full cycles	117 mW	47 mW	164 mW
3-stage Pipelined (Pipe)	1.6 GHz	2.28 mm ²	7 slots/VC	2 full cycles	195 mW	44 mW	239 mW
3-stage Pipelined with Bypass (Bypass)	1.3 GHz	2.61 mm ²	7 slots/VC	2 full cycles	180 mW	51 mW	231 mW
RapidLink-Full-LT	1.0 GHz	1.69 mm ²	3 slots/VC + 2 DS-EBs (4 VCs)	3 half cycles	104 mW	49 mW	153 mW

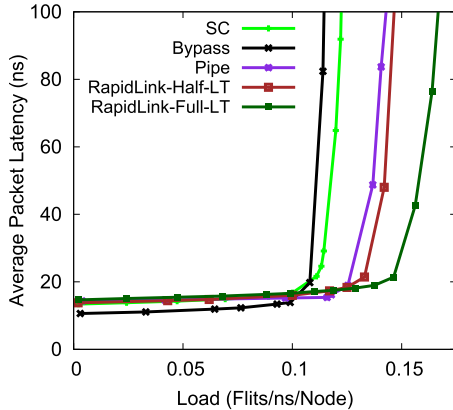


Fig. 12. Latency vs. load curves for a 8 × 8 2D mesh NoC under PARSEC-derived Hot-Spot (HS) traffic patterns.

achieve 31%, on average, increase in the NoC’s saturation throughput under UR, BC, and TS traffic. More importantly, RapidLink variants offer higher throughput than the fast pipelined NoC router implementations. Even if those designs operate at a much higher clock frequency, their throughput is limited by their inherent speculative operation in terms of allocation, i.e., as a flit moves deeper in the pipeline, it may have to repeat previous stages if it loses in arbitration. Single-cycle and RapidLink designs do not have such problems, since their single-cycle operation is inherently non-speculative. In the case of LC traffic, network contention is very low, which favors the very fast 3-stage pipelined implementations.

In all examined scenarios, the Scorpio-based design (“Bypass”) offers the lowest zero-load latency, since it effectively combines the benefits of a single-cycle design and the speed of a pipelined organization. The latency of RapidLink follows the same trend as the latency of the other designs that lead to higher throughput implementations.

Similar conclusions are drawn in the case of HS traffic as depicted in Fig. 12. In this scenario, RapidLink-Full-LT offers the highest saturation throughput, which is 14% higher than the saturation throughput of the second-best architecture, i.e., the 3-stage pipelined design (“Pipe”), and 43% higher than “Bypass,” which, again, exhibits the lowest zero-load latency.

The significant throughput increase offered by RapidLink is achieved without dedicating more resources to the NoC than the baseline NoC with full-cycle and single-data-rate

links (neither within the routers, nor on the links). Table I reports the layout area occupied by a complete NoC, for all designs under comparison. The area of RapidLink also includes any bridge modules put in front of the network interfaces, as well as the DS-EBs put on the links in the case of RapidLink-Full-LT. RapidLink – that is built using 2 routers of 2-VCs each – and the 4-VC single-cycle designs require the least area, due to their simplified buffering and allocation logic.

As far as the power is concerned, which is computed assuming an injection load of 0.15 flits/ns/node for all designs, both RapidLink variants consume substantially lower power than the 3-stage pipelined (“Pipe”) and the Scorpio-based (“Bypass”) designs. Specifically, the power savings of RapidLink – compared to these two designs – range from 44% to 53%. Compared to the single-cycle (“SC”) 4-VC baseline, RapidLink saves 16% to 21% in power consumption. Increasing the injection load would also proportionally increase the observed power consumption.

B. Networks With Long Inter-Router Links

The effectiveness of RapidLink in increasing the NoC’s throughput in a power-efficient manner is also observed when the NoC adopts higher-radix topologies. We experimented with the same 64-node system shown in Fig. 1, but, instead on relying on an 8×8 2D mesh for connecting the tiles, we employed a 4×4 high-radix 2D mesh architecture. In this case, each 4-VC router is placed in the middle of 4 tiles (cores) and connects directly to all of them. The concentration of 4 tiles to a single router leads to a NoC router with 8 input/output ports and *long 64-bit NoC links that are 4 mm long*. This increased link length causes an equivalent increase in the delay of the repeated wires, which now reaches 1050 ps. With this wire delay, all architectures under comparison need some of form of link pipelining, in order for the wires not to limit the NoC’s clock frequency.

The derived clock frequencies and link latencies for all designs are shown in Table II. The clock frequency of every design is inevitably reduced relative to the low-radix 2D mesh (Table I), due to the increased number of input/output ports in each higher radix router, which increases the logic depth of arbitration and the multiplexing circuitry inside the router.

For RapidLink, we only consider the RapidLink-Full-LT configuration, which can pipeline the link using the

TABLE III

HARDWARE CHARACTERISTICS OF ALL ARCHITECTURES UNDER EVALUATION FOR A 64-NODE CMP WITH SINGLE-STREAM WORMHOLE ROUTERS

NoC Architecture (Wormhole - 64-bit links)	Clock Frequency	Area	Minimum Required Input Buffers	Link Latency	Router Power	Link Power	Total Power
Base	1.67 GHz	0.877 mm ²	3 slots	1 full cycle	207 mW	72 mW	279 mW
RapidLink-Lockstep	1.51 GHz	0.803 mm ²	3 slots + 1 DS-EB	2 half cycles	159 mW	47 mW	206 mW
RapidLink-Concentrated	1.51 GHz	0.766 mm ²	3 slots + 1 & 2 DS-EBs	2 & 4 half cycles	109 mW	83 mW	192 mW




Fig. 13. Latency vs. load curves for a 4×4 concentrated 2D mesh of 64 nodes under UR and HS traffic. (a) Uniform Random (UR). (d) Hot-Spot (HS).

DS-EB structures described in Section II-B. Under this configuration with long 4 mm links, RapidLink-Full-LT requires 3 pipeline stages of DS-EBs (i.e., 3 half cycles) on the link. As in the low-radix NoCs, even if RapidLink-Full-LT routers (two 2-VC 8-port sub-routers per node) are faster than the 4-VC 8-port single-cycle routers, we assume that both operate at the same clock frequency to account for the extra clocking uncertainty arising from the dual-clock-edge operation.

RapidLink-Half-LT requires half-cycle link traversal which, for a 4 mm link, means that the NoC routers should operate below 500 MHz. Even if this configuration is acceptable for low-cost systems [32], it is not appropriate when comparing fast NoC designs, as done in this case.

The networking performance results, which are shown in Fig. 13, include UR and HS traffic. Permutation traffic patterns show exactly the same trend as in the case of low-radix NoC, while LC traffic is not interesting in this case (all architectures under comparison behave equivalently), since the majority of the traffic is produced and consumed within each router without ever using the NoC’s links, due to the concentration of 4 local cores per router. In both reported scenarios, RapidLink-Full-LT offers the highest saturation throughput, while its zero-load latency is relatively close to the zero-load latency of the “Bypass” design. RapidLink-Full-LT, irrespective of its lower clock frequency, achieves 35% higher saturation throughput, on average, relative to all designs under comparison.

Again this benefit comes at a minimal hardware cost. As shown in Table II, the area of RapidLink is significantly lower than the pipelined designs, and it is slightly higher than the “SC” configuration. The same trend is followed regarding the average power consumption. In this case, the average power shown in Table II was measured for UR traffic at an injection load of 0.15 flits/ns/node, i.e., a throughput that is below saturation for all designs. The RapidLink architecture

offers the best power profile, consuming 36% and 34% less power than “Pipe” and “Bypass” designs, respectively, and 7% less power than the “SC” design. The link power of all designs is increased, due to the longer links and the pipeline registers (or the DS-EBs, in the case of RapidLink) added on the link to enable higher clock frequencies.

C. Single-Stream Networks

As a final step, we compare the performance and hardware complexity of single-stream RapidLink configurations (as described in Section III) versus a simple wormhole network that supports – by default – single-stream sources. In these experiments, we, again, assume the same 64-node CMP following the floorplan of the Scorpio chip [8] at 45 nm with 64-bit links. Specifically, we evaluate a baseline wormhole network, a *concentrated* RapidLink (Section III-A), and a RapidLink NoC with routers operating in *lockstep mode* (Section III-B).

In the case of the baseline wormhole network (“Base”) and the lockstep-mode RapidLink NoC (“RapidLink-Lockstep”), the NoC’s topology is an 8×8 2D mesh with 2 mm links. Wormhole routers are simpler designs that can achieve higher clock frequencies. In our case, the baseline wormhole NoC can operate at 1.67 GHz, limited by the speed of the router; a flit spends one cycle inside the router and one full cycle on the links. Lockstep-mode RapidLink routers can operate at the same clock frequency. However, the dual-clock operation imposes an additional clock uncertainty. Therefore, we assume a degraded operating frequency of 1.51 GHz, which approximately corresponds to a 10% additional clock uncertainty. At this clock speed, a half-cycle link traversal of a 2 mm wire is not possible. Thus, RapidLink-Lockstep employs full-cycle link traversals (2 half cycles), through pipelining with DS-EBs.

The concentrated RapidLink (“RapidLink-Concentrated”) merges two single-stream sources in one DDR stream through network concentration. This 2-way concentration inevitably creates an unbalance in the NoC’s topology, which becomes an 8×4 2D mesh. In this case, the links in one dimension are 2 mm long, while, in the other dimension, they are 4 mm long. The concentration in this case does not increase the radix of the NoC routers that can still operate safely at 1.51 GHz; stream multiplexing occurs outside the NoC using appropriate bridge modules as the one shown in Figs. 9(a) and (b). In order not to let the link delay limit the NoC’s clock frequency of 1.51 GHz, we split the 2 mm links in two half-cycle pipeline stages, and the 4 mm links in four half-cycle pipeline stages, which suffice to cover the 500 ps and 1050 ps delays of the corresponding links, respectively.




Fig. 14. Latency vs. load curves for *single-stream* networks (i.e., with no VCs) under UR, LC, BC, and TS traffic patterns. (a) Uniform Random (UR). (b) Localized (LC) (c) Bit-Complement (BC). (d) Transpose (TS).

The hardware characteristics of the single-stream (wormhole) designs under evaluation are summarized in Table III.

Figs. 14(a)–(d) depict the latencies imposed by each NoC as a function of the node injection rate, for the four examined synthetic traffic patterns. Application-derived HS traffic is omitted here, since it refers to NoCs with VCs. Single-stream synthetic HS traffic was also tested and the results follow the same trend as the other synthetic traffic patterns. In all cases, the two RapidLink alternatives achieve the highest throughput: The concentrated architecture excels in UR, BC, and TS traffic, while the lockstep-mode alternative in LC traffic. The throughput increase achieved by RapidLink ranges between 5% and 26% when compared to “Base”, even if RapidLink NoCs are assumed to operate at a lower clock frequency for higher safety with respect to on-chip variations and clock jitter. This lower clock frequency, and the slightly increased latency in the network interfaces of RapidLink, are the two main reasons for the better zero-load latency of “Base,” relative to RapidLink architectures.

The areas of the two single-stream RapidLink alternatives (i.e., RapidLink-Concentrated and RapidLink-Lockstep) and the baseline wormhole network are depicted in Table III. For the two single-stream RapidLink configurations, the area numbers also include the area of the bridge modules at the network interfaces, and the area of the DS-EBs placed on the links. The reported area analysis indicates that both RapidLink configurations require the least area. RapidLink-Concentrated and RapidLink-Lockstep are 12% and 8% smaller – in terms of occupied area – than the baseline NoC, respectively.

The same trend is observed when considering the average power consumption of all designs under comparison,

except that the savings reaped by RapidLink are even higher. As reported in Table III, RapidLink-Concentrated and RapidLink-Lockstep consume 31% and 26% less power than the baseline NoC, respectively.

Hence, if one takes into account the high network performance and the low area/power consumption of RapidLink-Concentrated, one can safely conclude that it is the best *overall* architecture for single-stream networks. However, the RapidLink-LockStep design also emerges as the best performer – in terms of throughput – under localized traffic. Thus, if a designer knows in advance that the system will be dominated by localized traffic (e.g., in an application-specific embedded SoC), then RapidLink-LockStep would offer the best performance at a low hardware cost.

V. CONCLUSIONS

The asymmetry between the NoC’s intra- and inter-router delays has been exploited in many forms in the past, primarily aiming to allow flits to traverse longer distances within a single clock cycle. NoCs with high-radix routers constitute such an example; they allow flits to reach their destinations using fewer hops, albeit through the use of longer links and fairly complex routers. The increased number of ports complicates allocation and switching logic, and requires custom design to achieve acceptable operating frequencies. The design of high-radix networks typically leads to complicated layouts and wire-routing congestion, which also necessitate custom design effort. Additionally, high-radix NoCs incur higher latencies and power consumption when handling local traffic (e.g., near-neighbor), because of unnecessary data movement over longer distances.

Other alternatives employ single-cycle multi-hop link traversal by relying on complicated flow control rules and router bypassing to cross multiple hops “asynchronously” in one cycle. Once again, this philosophy does not offer a true benefit under localized traffic, which involves packets traversing one, or at most two, hops and requires significant redesign both at the micro-architectural and physical design levels.

In both above-mentioned philosophies, fast wire traversal is achieved assuming tile-based homogeneous systems that are characterized by a regularity in their physical layout. When longer links exist in the NoC, link pipelining should be adopted, which ruins the main property of delivering flits over a long distance in a single clock cycle.

On the other hand, the proposed RapidLink NoC architecture complements previous state-of-the-art proposals by following a distinct and more scalable design path, which improves network performance without increasing design cost. RapidLink is minimally intrusive to both the router’s micro architecture and the flow-control policies, and it can be applied to any low- or medium-radix topology. RapidLink does not lose the benefits of local connectivity, and by using the proposed DDR dual-stream elastic buffers, one can split even longer links into multiple half-cycle DDR segments. RapidLink can be equally applied to single- or multiple-VC NoC configurations, while still offering significant network performance improvements and without increasing the

hardware cost. Furthermore, RapidLink is shown to outperform fast state-of-the-art pipelined router organizations that also include fine-grained pipeline-stage-bypassing capabilities.

REFERENCES

- [1] K. Sewell *et al.*, “Swizzle-switch networks for many-core systems,” *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 2, pp. 278–294, Jun. 2012.
- [2] N. Abeyratne *et al.*, “Scaling towards kilo-core processors with asymmetric high-radix topologies,” in *Proc. IEEE Int. Symp. High Perform. Comput. Archit. (HPCA)*, Feb. 2013, pp. 496–507.
- [3] J. Kim, J. Balfour, and W. J. Dally, “Flattened butterfly topology for on-chip networks,” in *Proc. 40th Annu. IEEE/ACM Int. Symp. Microarchitecture MICRO*, Washington, DC, USA, Dec. 2007, pp. 172–182.
- [4] T. Krishna, C.-H. O. Chen, W. C. Kwon, and L.-S. Peh, “Smart: Single-cycle multi-hop traversals over a shared network-on-chip,” *IEEE Micro*, vol. 34, no. 3, pp. 43–56, May/June 2014.
- [5] B. K. Daya, L. S. Peh, and A. P. Chandrakasan, “Towards high-performance bufferless nocs with sceptor,” *IEEE Comput. Archit. Lett.*, vol. 15, no. 1, pp. 62–65, Jan. 2016.
- [6] M. Azimi, D. Dai, A. Kumar, and A. S. Vaidya, “On-chip interconnect trade-offs for tera-scale many-core processors,” in *Designing Network On-Chip Architectures in the Nanoscale Era*, J. Flich and D. Bertozzi, Eds. Boca Raton, FL, USA: CRC Press, 2011.
- [7] P. Salihundam *et al.*, “A 2 Tb/s 6×4 mesh network with DVFS and 2.3 Tb/s/W router in 45 nm CMOS,” in *Proc. VLSI Circuits*, 2010, pp. 79–80.
- [8] B. Daya *et al.*, “SCORPIO: A 36-core research chip demonstrating snoopy coherence on a scalable mesh noc with in-network ordering,” in *Proc. Int. Symp. Comput. Archit.*, Jun. 2014, pp. 25–36.
- [9] R. Manevich, L. Polishuk, I. Cidon, and A. Kolodny, “Designing single-cycle long links in hierarchical NoCs,” *Microprocessors Microsyst.*, vol. 38, no. 8, pp. 814–825, 2014.
- [10] N. Weste and D. Harris, *CMOS VLSI Design: A Circuits and Systems Perspective*. 3rd ed. Reading, MA, USA: Addison Wesley 2010.
- [11] W. J. Dally, C. Malachowsky, and S. W. Keckler, “21st century digital design tools,” in *Proc. ACM Des. Autom. Conf. (DAC)*, 2013, p. 94.
- [12] P.-H. Ho, “Interesting problems in physical synthesis,” in *Proc. Int. Symp. Phys. Des.*, 2017, p. 131.
- [13] A. Golander, N. Levison, O. Heymann, A. Briskman, M. J. Wolski, and E. F. Robinson, “A cost-efficient L1–L2 multicore interconnect: Performance, power, and area considerations,” *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 58, no. 3, pp. 529–538, Mar. 2011.
- [14] T. Krishna *et al.*, “Single-cycle multihop asynchronous repeated traversal: A smart future for reconfigurable on-chip networks,” *IEEE Comput.*, vol. 46, no. 10, pp. 48–55, Oct. 2013.
- [15] J. Balfour and W. J. Dally, “Design tradeoffs for tiled cmp on-chip networks,” in *Proc. 20th Annu. Int. Conf. Supercomput.*, 2006, pp. 187–198.
- [16] A. Psarras, J. Lee, P. Mattheakis, C. Nicopoulos, and G. Dimitrakopoulos, “A low-power network-on-chip architecture for tile-based chip multi-processors,” in *Proc. ACM Great Lakes Symp-VLSI*, 2016, pp. 335–340.
- [17] I. Miro-Panades, F. Clermidy, P. Vivet, and A. Greiner, “Physical implementation of the DSPIN network-on-chip in the FAUST architecture,” in *Proc. Int. Symp. Netw.-Chip*, 2008, pp. 139–148.
- [18] T. Bjerregaard, M. B. Stensgaard, and J. Sparso, “A scalable, timing-safe, network-on-chip architecture with an integrated clock distribution method,” in *Proc. Des. Autom. Test Europe Conf. Exhibit. (DATE)*, 2007, pp. 1–6.
- [19] G. Dimitrakopoulos, A. Psarras, and I. Seitanidis, *Microarchitecture of Network-on-Chip Routers: A Designer’s Perspective*. New York, NY, USA: Springer, 2015.
- [20] T. Halfhill, “Automating front-end SoC design with NetSpeed’s on-chip network IP,” The Linley Group, Mountain View, CA, USA, Mar. 2015.
- [21] Q. Wu, M. Pedram, and X. Wu, “A new design of double edge triggered flip-flops,” in *Proc. Asia South Pacific Des. Autom. Conf. (ASP-DAC)*, 1998, pp. 417–421.
- [22] J.-S. Seo, D. Sylvester, D. Blaauw, H. Kaul, and R. Krishnamurthy, “A robust edge encoding technique for energy-efficient multi-cycle interconnect,” in *Proc. ACM/IEEE Int. Symp. Low-Power Electron. Des. (ISLPED)*, Aug. 2007, pp. 68–73.
- [23] H. M. Jacobson *et al.*, “Synchronous interlocked pipelines,” in *Proc. Int. Symp. Asynchron. Circuits Syst.*, Apr. 2002, pp. 3–12.
- [24] J. Cortadella, M. Kishinevsky, and B. Grundmann, “Synthesis of synchronous elastic architectures,” in *Proc. ACM Des. Autom. Conf. (DAC)*, 2006, pp. 657–662.
- [25] G. Michelogiannakis and W. J. Dally, “Elastic buffer flow control for on-chip networks,” *IEEE Trans. Comput.*, vol. 62, no. 2, pp. 295–309, Feb. 2013.
- [26] G. Dimitrakopoulos, E. Kalligeros, and K. Galanopoulos, “Merged switch allocation and traversal in network-on-chip switches,” *IEEE Trans. Comput.*, vol. 62, no. 10, pp. 2001–2012, Oct. 2013.
- [27] P. Kumar, Y. Pan, J. Kim, G. Memik, and A. Choudhary, “Exploring concentration and channel slicing in on-chip network router,” in *Proc. ACM/IEEE Int. Symp. Netw.-Chip*, May 2009, pp. 276–285.
- [28] J. Lee, C. Nicopoulos, H. G. Lee, and J. Kim, “TornadoNoC: A lightweight and scalable on-chip network architecture for the many-core era,” *ACM Trans. Archit. Code Optim. (TACO)*, vol. 10, no. 4, p. 56, Dec. 2013.
- [29] C. Bienia, S. Kumar, J. P. Singh, and K. Li, “The PARSEC benchmark suite: Characterization and architectural implications,” in *Proc. 17th PACT*, Oct. 2008, pp. 72–81.
- [30] Y. Lu, C. Chen, J. V. McCanny, and S. Sezer, “Design of interlock-free combined allocators for networks-on-chip,” in *Proc. IEEE 25th Int. SoC Conf. (SoCC)*, Sep. 2012, pp. 358–363.
- [31] Y. Hoskote, S. Vangal, A. Singh, N. Borkar, and S. Borkar, “A 5-GHz mesh interconnect for a Teraflops processor,” *IEEE Micro*, vol. 27, no. 5, pp. 51–61, Sep./Oct. 2007.
- [32] C. Tan, A. Kulkarni, V. Venkataramani, M. Karunaratne, T. Mitra, and L.-S. Peh, “LOCUS: Low-power customizable many-core architecture for wearables,” in *Proc. Int. Conf. Compil., Archit. Synth. Embedded Syst. (CASES)*, 2016, pp. 11–11–11–10.



Anastasios Psarras received the Diploma and master’s degrees in electrical and computer engineering from the Democritus University of Thrace, Xanthi, Greece, in 2012 and 2013, respectively, where he is currently pursuing the Ph.D. degree.

His current research interests include system-on-a-chip design, and in particular, on-chip interconnection networks.



Savvas Moisisdis received the Diploma in electrical and computer engineering from the Democritus University of Thrace, Xanthi, Greece, in 2016, where he is currently pursuing the Ph.D. degree.

His research interests include nanoelectronics with emphasis in circuit design for emerging nanodevices.



Chrysostomos Nicopoulos received the B.S. and Ph.D. degrees in electrical engineering with a specialization in computer engineering from Pennsylvania State University, State College, PA, USA, in 2003 and 2007, respectively.

He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, University of Cyprus, Nicosia, Cyprus. His current research interests include networks-on-chip, computer architecture, multi/many-core microprocessor, and computer system design.



Giorgos Dimitrakopoulos received the B.S., M.Sc., and Ph.D. degrees from the University of Patras, Patras, Greece, in 2001, 2003 and 2007, respectively, all in computer engineering.

He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi, Greece. He is interested in the design of digital integrated circuits, electronic design automation, and computer architecture, with emphasis in low-power systems design.