

## A case for hierarchical rings with deflection routing: An energy-efficient on-chip communication substrate



Rachata Ausavarungnirun<sup>a,\*</sup>, Chris Fallin<sup>a</sup>, Xiangyao Yu<sup>b</sup>, Kevin Kai-Wei Chang<sup>a</sup>,  
Greg Nazario<sup>a</sup>, Reetuparna Das<sup>c</sup>, Gabriel H. Loh<sup>d</sup>, Onur Mutlu<sup>a</sup>

<sup>a</sup> Carnegie Mellon University, United States

<sup>b</sup> University of Michigan, United States

<sup>c</sup> Massachusetts Institute of Technology, United States

<sup>d</sup> Advanced Micro Devices, United States

### ARTICLE INFO

#### Article history:

Available online 11 February 2016

#### Keywords:

Network on Chip  
Parallelism  
Interconnect

### ABSTRACT

Hierarchical ring networks, which hierarchically connect multiple levels of rings, have been proposed in the past to improve the scalability of ring interconnects, but past hierarchical ring designs sacrifice some of the key benefits of rings by reintroducing more complex in-ring buffering and buffered flow control. Our goal in this paper is to design a new hierarchical ring interconnect that can maintain most of the simplicity of traditional ring designs (i.e., no in-ring buffering or buffered flow control) while achieving high scalability as more complex buffered hierarchical ring designs.


To this end, we revisit the concept of a hierarchical-ring network-on-chip. Our design, called **HiRD** (Hierarchical Rings with Deflection), includes critical features that enable us to mostly maintain the simplicity of traditional simple ring topologies while providing higher energy efficiency and scalability. *First*, HiRD does not have any buffering or buffered flow control within individual rings, and requires only a small amount of buffering between the ring hierarchy levels. When inter-ring buffers are full, our design simply *deflects* flits so that they circle the ring and try again, which eliminates the need for in-ring buffering. *Second*, we introduce two simple mechanisms that together provide an end-to-end delivery guarantee within the entire network (despite any deflections that occur) without impacting the critical path or latency of the vast majority of network traffic.

Our experimental evaluations on a wide variety of multiprogrammed and multithreaded workloads and synthetic traffic patterns show that HiRD attains equal or better performance at better energy efficiency than multiple versions of both a previous hierarchical ring design and a traditional single ring design. We also extensively analyze our design's characteristics and injection and delivery guarantees. We conclude that HiRD can be a compelling design point that allows higher energy efficiency and scalability while retaining the simplicity and appeal of conventional ring-based designs.

© 2016 Elsevier B.V. All rights reserved.

\* Corresponding author. Tel.: +1 412-589-9198.

E-mail address: [rachata@cmu.edu](mailto:rachata@cmu.edu) (R. Ausavarungnirun).



**Fig. 1.** A traditional hierarchical ring design [22,24,53,54,64] allows “local rings” with simple node routers to scale by connecting to a “global ring” via bridge routers.

## 1. Introduction

Interconnect scalability, performance, and energy efficiency are first-order concerns in the design of future CMPs (chip multiprocessors). As CMPs are built with greater numbers of cores, centralized interconnects (such as crossbars or shared buses) are no longer scalable. The Network-on-Chip (NoC) is the most commonly-proposed solution [12]: cores exchange packets over a network consisting of network switches and links arranged in some topology.

Mainstream commercial CMPs today most commonly use *ring*-based interconnects. Rings are a well-known network topology [11], and the idea behind a ring topology is very simple: all routers (also called “ring stops”) are connected by a loop that carries network traffic. At each router, new traffic can be injected into the ring, and traffic in the ring can be removed from the ring when it reaches its destination. When traffic is traveling on the ring, it continues uninterrupted until it reaches its destination. A ring router thus *needs no in-ring buffering or flow control* because it prioritizes on-ring traffic. In addition, the router’s datapath is very simple compared to a mesh router, because the router has fewer inputs and requires no large, power-inefficient crossbars; typically it consists only of several MUXes to allow traffic to enter and leave, and one pipeline register. Its latency is typically only one cycle, because no routing decisions or output port allocations are necessary (other than removing traffic from the ring when it arrives). Because of these advantages, several prototype and commercial multicore processors have utilized ring interconnects: the Intel Larrabee [55], IBM Cell [51], and more recently, the Intel Sandy Bridge [27].

Unfortunately, rings suffer from a fundamental scaling problem because a ring’s bisection bandwidth does not scale with the number of nodes in the network. Building more rings, or a wider ring, serves as a stopgap measure but increases the cost of every router on the ring in proportion to the bandwidth increase. As commercial CMPs continue to increase core counts, a new network design will be needed that balances the simplicity and low overhead of rings with the scalability of more complex topologies.

A hybrid design is possible: rings can be constructed in a *hierarchy* such that groups of nodes share a simple ring interconnect, and these “local” rings are joined by one or more “global” rings. Fig. 1 shows an example of such a *hierarchical ring* design. Past works [22,24,53,54,64] proposed hierarchical rings as a scalable network. These proposals join rings with *bridge routers*, which reside on multiple rings and transfer traffic between rings. This design was shown to yield good performance and scalability [53]. The state-of-the-art design [53] requires *flow control and buffering* at every node router (ring stop), because a ring transfer can make one ring back up and stall when another ring is congested. While this previously proposed hierarchical ring is much more scalable than a single ring [53], the reintroduction of in-ring buffering and flow control nullifies one of the primary advantages of using ring networks in the first place (i.e., the lack of buffering and buffered flow control within each ring).

**Our goal** in this work is to design a ring-based topology that is simpler and more efficient than prior ring-based topologies. To this end, our design uses simple ring networks that do not introduce any in-ring buffering or flow control. Like past proposals, we utilize a hierarchy-of-rings topology to achieve higher scalability. However, beyond the topological similarities, our design is very different in how traffic is handled within individual rings and between different levels of rings. We introduce a new *bridge router* microarchitecture that facilitates the transfer of packets from one ring to another. It is in these, and *only* these, limited number of bridge routers where we require any buffering.

**Our key idea** is to allow a bridge router with a full buffer to *deflect* packets. Rather than requiring buffering and flow control in the ring, packets simply cycle through the network and try again. While deflection-based, bufferless networks have been proposed and evaluated in the past [2,5,20,26,46,56], our approach is effectively an elegant hybridization of bufferless (rings) and buffered (bridge routers) styles. To prevent packets from potentially deflecting around a ring arbitrarily many times (i.e., to prevent livelock), we introduce two new mechanisms, the *injection guarantee* and the *transfer guarantee*, that ensure packet delivery even for adversarial/pathological conditions (as we discuss in [3] and evaluate with worst-case traffic in Section 4.3).

This simple hierarchical ring design, which we call *HiRD* (for Hierarchical Rings with Deflection), provides a more scalable network architecture while retaining the key simplicities of ring networks (no buffering or flow control within each ring). We show in our evaluations that HiRD provides better performance, lower power, and better energy efficiency with respect to the buffered hierarchical ring design [53].

Download English Version:

<https://daneshyari.com/en/article/523831>

Download Persian Version:

<https://daneshyari.com/article/523831>

[Daneshyari.com](https://daneshyari.com)