CCIT Report #568
December 2005

Curing Hotspots in Wormhole NoCs

Isask'har Walter[1], Israel Cidon[2], Ran Ginosar[2], Avinoam Kolodny[2]

Electrical Engineering Department, Technion, Haifa, Israel

[1]zigi@tx.technion.ac.il [2]{cidon, ran, kolodny}@ee.technion.ac.il

ABSTRACT

Network on Chip (NoC) has been proposed for future SoC interconnect. Hotspots are SoC modules which occasionally receive traffic that exceeds the rate at which they can absorb data. Hotspots are common in real-life SoCs, such as external DRAM or internal components (caches, CAMs, special purpose processors) that are bandwidth limited and in high demand by other units. In this paper we demonstrate that hotspot modules on wormhole-based NoCs dramatically reduce network efficiency and unfairly allocate system resources. A single hotspot may ruin the performance of the entire NoC. In order to resolve these problems, we introduce a novel low-cost end-to-end credit based resource allocation technique that regulates access to the hotspot module. Using simulation, we show the effectiveness of the suggested mechanism.

Categories and Subject Descriptors

System-Level Design and Co-Design: Network-on-Chip (NoC)

General Terms

Algorithms, Performance, Design

Keywords

Network on-Chip, wormhole, flow-control, hotspot, SoC

1. INTRODUCTION

Wormhole [1] switching is commonly employed in NoC (e.g. [2], [3], [4], [5]), thanks to its small buffer requirements and low latencies at light load. Each packet is divided into small fixed size parts called flits, which are transmitted to the next hop without waiting for the entire packet to be received. This causes the transmitted packet to be "spread" along the path between the source and destination nodes in a pipeline fashion. The main drawback of wormhole interconnect is the high sensitivity to packet blocking, calling for high performance networks to operate at relatively low utilization and to include multiple virtual channels [6].

Bandwidth can be allocated to NoC links for providing adequate performance [7], but speeding up the operation of IP modules is impossible or has an unacceptable cost. Furthermore, at certain times the aggregated traffic demand might exceed the destination module's bandwidth capacity. A bandwidth-limited SoC module working close to its capacity is termed a *hotspot*. In this situation, when the module is unable to consume incoming packets fast enough, the entire network may be affected. Hop-by-hop backpressure causes buffers at the router adjacent to the hotspot to be filled up and become stalled blocking new arrivals to this router. This creates a domino effect, by which the delivery of packets to ports of more distant routers is slowed down, forming a *saturation tree* [8] with the hotspot module as its root, as illustrated in Figure 1. The NoC suffers increased delays in packet delivery and unfair network utilization (modules near the hotspot get larger portion of its capacity).

While this effect may also exist in packet based store-and-forward networks, the threat is particularly troublesome in wormhole based architectures due to packet "stretching" across several hops. As a result, hotspot effects may extend network wide instantly. It is important to note that this hotspot phenomenon is independent of links bandwidth, as a saturation tree may build up in a system with infinite capacity links and a single heavily loaded module. Consequently, even carefully designed, largely over-provisioned NoCs may suffer of poor performance if potential hotspots are left unhandled.

We propose a novel credit-based resource allocation mechanism for solving hotspot congestion problems in wormhole-based NoCs. A hotspot allocation controller is introduced to arbitrate short, high priority credit requests. The controller regulates hotspot access according to the quality of service requirements of the specific system. Credit requests and grants are transmitted as high-priority packets (grants may be piggybacked on other messages). Auto-refresh or pre-allocation can be used for selected modules to eliminate light load latencies. One side of the mechanism is implemented in modules' interfaces and the other in a strategic location, while NoC routers remain unchanged. This prevents the accumulation of packets destined at a hotspot module in the network. Consequently, other traffic remains unaffected even when the hotspot load increases significantly.

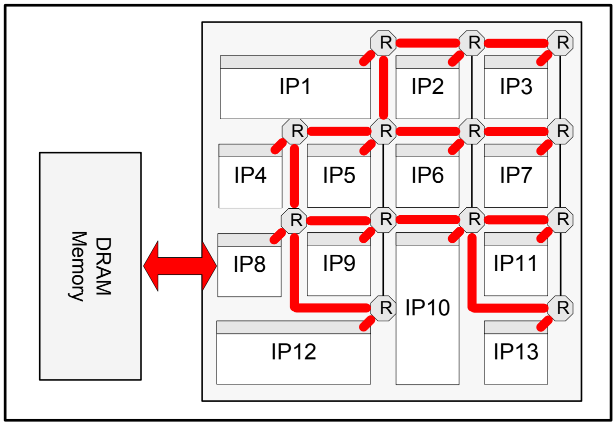

Figure 1: SoC hotspot at external DRAM interface and the resulting NoC saturation tree (highlighted links)

This work was partially supported by the Semiconductor Research Corporation (SRC), Intel Corp., and the iSRC consortium