Microprocessors and Microsystems 35 (2011) 119–129

Contents lists available at ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

On an efficient NoC multicasting scheme in support of multiple applications running on irregular sub-networks

Xiaohang Wang [a,b], Mei Yang [b], Yingtao Jiang [b], Peng Liu [a,*]

[a] Department of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, Zhejiang 310027, PR China
[b] Department of Electrical and Computer Engineering, University of Nevada, Las Vegas 89154, USA

ARTICLE INFO

Article history:
Available online 17 August 2010

Keywords:
Network-on-chips (NoCs)
Chip multiprocessor (CMP)
Multicast
Routing

ABSTRACT

When a number of applications simultaneously running on a many-core chip multiprocessor (CMP) chip connected through network-on-chip (NoC), significant amount of on-chip traffic is one-to-many (multicast) in nature. As a matter of fact, when multiple applications are mapped onto an NoC architecture with applicable traffic isolation constraints, the corresponding sub-networks of these applications are mapped onto actually tend to be irregular. In the literature, multicasting for irregular topologies is supported through either multiple unicasting or broadcasting, which, unfortunately, results in overly high power consumption and/or long network latency. To address this problem, a simple, yet efficient hardware-based multicasting scheme is proposed in this paper. First, an irregular oriented multicast strategy is proposed. Literally, following this strategy, an irregular oriented multicast routing algorithm can be designed based on any regular mesh based multicast routing algorithm. One such algorithm, namely, Alternative Recursive Partitioning Multicasting (AL + RPM), is proposed based on RPM, which was designed for regular mesh topology originally. The basic idea of AL + RPM is to find the output directions following the basic RPM algorithm and then decide to replicate the packets to the original output directions or the alternative (AL) output directions based on the shape of the sub-network. The experiment results show that the proposed multicast AL + RPM algorithm can consume, on average, 14% and 20% less power than bLBDR (a broadcasting-based routing algorithm) and the multiple unicast scheme, respectively. In addition, AL + RPM has much lower network latency than the above two approaches. To incorporate AL + RPM into a baseline router to support multicasting, the area overhead is fairly modest, less than 5.5%.

1. Introduction

Advance in technology continues to drive the increase of transistor integration capacity. It is estimated that by 2015, there will be 100 billion transistors integrated on a 300 mm^2 die [1]. To exploit this large number transistors and also take into consideration of pressing high power consumption of ever bigger chips, the design paradigm is migrating to many-core architectures [1,2]. Network-on-chip (NoC) [3] has been proposed as the mainstream on-chip network architecture to efficiently interconnect the large number of (16 or more) processing cores integrated on a many-core system. Some most recent, high profile examples include Intel's Teraflop [4] and Tilera [5] chips featuring many-core chip multiprocessors (CMPs) architectures with 2D mesh topologies [13] for on-chip interconnect.

With the development of diverse applications and programming models on CMPs, one-to-many communication and one-to-all communication are becoming more common. For example, in CMPs with cache coherent shared memory systems, the cache coherence protocols exhibit one-to-many communication characteristics to keep the ordering of different requests or to invalidate shared data on different cache nodes [6]. In [7], it has been observed that 5–10% of the network traffic is one-to-many in nature, ranging from scientific workloads to commercial workloads, in communication traces of different cache coherence protocols and operand network. Therefore, efficient support of one-to-many communications in CMPs, particularly hardware multicast support, will benefit a wide range of applications by boosting the network performance with reduced power consumption. Unfortunately, up to date, there is only very limited number of chip router designs that actually support multicasting [6–8].

In addition, the following issues make multicast supporting even more complicated. The first issue is topology irregularity. The large number of cores on a CMP unquestionably offers high parallelism in computation. To better utilize these vastly available computation resources, virtualization of the chip becomes a

* Corresponding author at: Department of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, Zhejiang 310027, PR China.
E-mail addresses: baikeina@yahoo.com.cn (X. Wang), Mei.Yang@unlv.edu (M. Yang), yingtao@egr.unlv.edu (Y. Jiang), liupeng@isee.zju.edu.cn (P. Liu).

0141-9331/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.micpro.2010.08.003

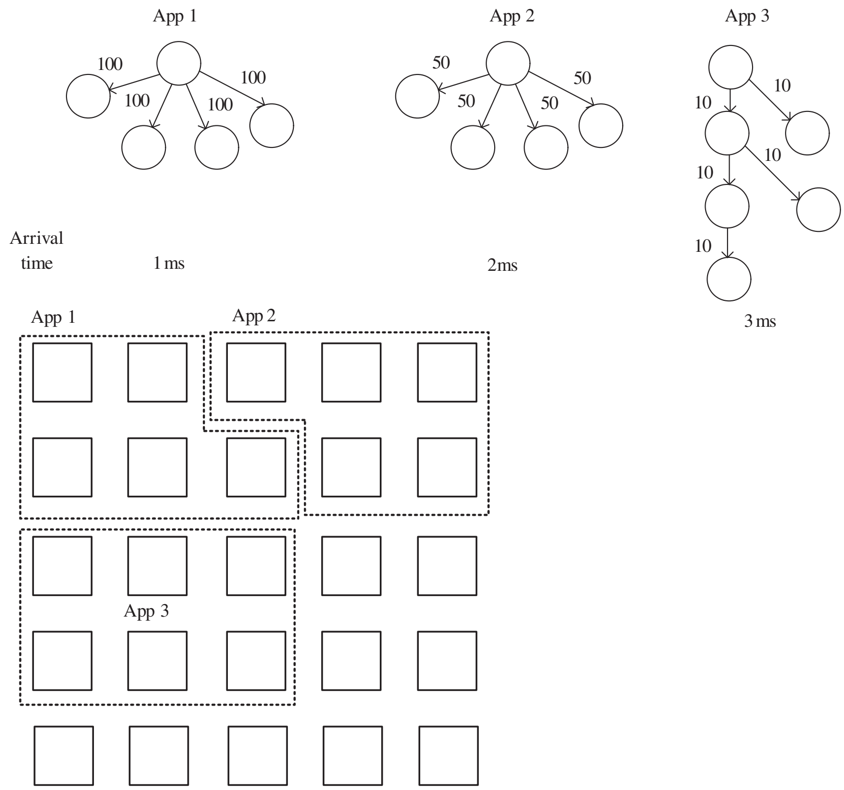

Fig. 1. Sub-network partition and task mapping of multiple applications on a 5 × 5 based mesh NoC.

necessity [9], where resources can be distributed among different virtual machines [3]. Applying virtualization [8] at the NoC level basically allows a single NoC-based CMP to be shared by multiple applications with each mapped to different sub-networks of the chip [10] either statically [11] or dynamically [12]. Fig. 1 shows an example with three applications arriving at 1 ms, 2 ms, and 3 ms. The three applications are allocated to three sub-networks which may not be regular shapes (e.g., 2D mesh, torus). On the other hand, virtualization requires traffic isolation [8]; that is, communication between nodes in a virtualized region is limited to the sub-network only. The irregular sub-network and traffic isolation requirements together negate regular 2D mesh oriented routing algorithms, like *XY* routing, odd–even routing, etc. [13].

The second issue is unpredictability of the application communication behavior. Different types of applications, such as desktop, server, embedded systems, will be executed on general purpose CMPs. It is impossible to pre-characterize the communication patterns among the cores inside a sub-network. As a result, customized NoC routing approaches (like the ones using routing tables [14]) may not be feasible.

Hence, it is important to design an efficient multicast mechanism which supports irregular topologies without the need of a routing table. In this paper, an irregular sub-network oriented multicast strategy is first proposed. Following this strategy, an irregular sub-network oriented multicast routing algorithm, namely, Alternative Recursive Partitioning Multicast (AL + RPM), is developed based on RPM [13], an efficient deterministic multicast routing algorithm proposed for regular mesh topology. To our best knowledge, our approach is the first multicast routing approach, as opposed to the broadcast-based one [8], that targets to irregular sub-networks.

In the rest of the paper, Section 2 reviews the existing work on multicast routing schemes in NoCs. Section 3 presents the preliminaries. Section 4 describes the irregular sub-network oriented multicast routing strategy and algorithm. Section 5 reports the performance evaluation of AL + RPM. Finally, Section 6 concludes the paper.

2. Related work

Multicast communication has been extensively studied in computer networks and interconnection networks [13]. However, due to the power and area constraints pertaining to NoCs, supporting multicast in NoCs has a different set of requirements. Particularly, an efficient multicasting approach for NoCs should result in low network latency and low power and area consumptions. A simple multicasting approach is to send a multicast packet as multiple unicast packets. However, such a scheme suffers from very large network latency and high power consumption [7]. Below reviews existing multicasting approaches proposed for mesh-based NoCs.

The multicast problem of regular mesh topology has been studied as in [13]. Generally, there are two types of multicast routing strategies, namely, path-based [15] and tree-based [6,7,16–18] multicast. Path-based multicast routing is to deliver the packet to each destination sequentially following one path [15]. Path-based multicast is attractive for its simplicity in hardware design. However, if the destination nodes are widely spread, path-based multicast may suffer higher latency compared to tree-based multicast. It is shown in [7] that path-based multicast may increase network latency by 48% compared to a unicast router.

Tree-based multicast routing [6,7,16–18] is to deliver the packet along a common path as far as possible and replicate packets (branch) for a unique set of destination nodes when necessary. Several tree-based multicast approaches have been proposed for NoCs with regular mesh topology. The virtual circuit tree-based multicast (VCTM) [7] avoids sending redundant packets as multiple unicast. However, it uses a lookup table based multicasting