Future Generation Computer Systems 56 (2016) 238–246

Contents lists available at ScienceDirect

Future Generation Computer Systems

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

Multiple-combinational-channel: A network architecture for workload balance and deadlock free☆

Liwei Chen [a], Yipeng Wang [a,*], Huandong Wang [b,c,d], Wenxiang Wang [b,c,d], Hua Jing [b,c,d], Guangfei Zhang [e]

[a] Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China
[b] State Key Laboratory of Computer Architecture, Chinese Academy of Sciences, Beijing, China
[c] Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China
[d] Loongson Technology Corporation Limited, Beijing, China
[e] Shannon Laboratory, Huawei Technologies Co., Beijing, China

HIGHLIGHTS

- We combine low usage channels and keep high usage channels wide enough.
- Our method balances the workload of traffic in NoC.
- Our method can avoid protocol-level deadlocks.
- Our method achieves high throughput in practice.
- Our method reduces area and power without the loss of performance.

ARTICLE INFO

Article history:
Received 14 January 2015
Received in revised form
10 July 2015
Accepted 25 August 2015
Available online 7 September 2015

Keywords:
NoC
CC-NUMA
Channel
Load balance
Deadlock free

ABSTRACT

As the number of cores in chip multiprocessors (CMPs) increases rapidly, network-on-chips (NoCs) have become the major role in ensuring performance and power scalability. In this paper, we propose multiple-combinational-channel (MCC), a load balancing and deadlock free interconnect network, for cache-coherent non-uniform memory accessing (CC-NUMA). In order to make load more balancing and reduce power dissipation, we combine low usage channels and make high usage channels independent and wide enough, since messages transmitted over NoC have different widths and injection rates. Furthermore, based on the in-depth analysis of network traffic, we summarize four traffic patterns and establish several rules to avoid protocol-level deadlocks. We implement MCC on a 16-core CMPs, and evaluate the workload balance, area, power and performance using universal workloads. The experimental results show that MCC reduces nearly 21% power than multiple-physical-channel with similar throughput. Moreover, MCC improves 10% performance with similar area and power, compared to packet-switching architecture with virtual channels.

1. Introduction

Nowadays, many-core systems have been the mainstream solution to improve the performance and power efficiency for both general purpose chip multiprocessors (CMPs) and special purpose system-on-chips (SoCs). Hence, increasing frequency and computing ability of single processor simply is regarded as an inefficient way to improve system performance. On the contrary, data movement among interconnect network has become the bottleneck of system performance, and network complexity is increasing exponentially with the processor number. Therefore, the network-on-chip (NoC) plays a key role in defining the performance, area and power of CMPs [1].

Most researches focus on the NoCs with hundreds of cores, and the packet-switching NoC with virtual channels (VCs) [2–4] seems to be a very successful resolution for large-scale CMPs, as shown in

☆ The preliminary version of this paper entitled "MCC: A Load Balancing and Deadlock Free Interconnect Network for Cache Coherent Chip Multiprocessors" was published in the proceedings of the 15th IEEE International Conference on Computational Science and Engineering (CSE), Nicosia, Cyprus, 2012 [21].

* Corresponding author.
E-mail address: chenliwei@iie.ac.cn (L. Chen).

http://dx.doi.org/10.1016/j.future.2015.08.013
0167-739X/

Fig. 1(b). However, the maximum throughput of VCs is limited by the unique shared physical channel (PC). In addition, VCs with fixed widths of buffers and channels are not suitable for various lengths of packets. For example, transmitting a long packet will take more than one cycle, while transporting short packets will waste bandwidth. Furthermore, VCs need very complex control logic to avoid deadlocks and packet loss, and additional packing circuits to pack data into packets and unpack packets. These circuits both lead to increasing of area, power and latency. As a result, packet-switching with VCs is not fit for small or medium-scale SoCs, like embedded systems [5].

Multiple-physical-channel (MPC) [6–8] is another typical approach to increase network throughput, which keeps traffic classes separate and uses heterogeneous channels to transport different messages, as depicted in Fig. 1(c). However, since different messages have different injection rates, these channels are load unbalancing. Some channels are very busy while others may be idle. Additionally, with too many wires and buffers, the area of MPC is very large, and the power is extremely high. For instance, the RAW prototype [7] had allocated nearly 50% of the area to an on-chip network, and the network accounted for nearly 40% of the overall chip power.

Previous works simply shared the only one physical channel among messages as in VCs, or employed multiple channels for packets with different lengths as in MPC. However, the intrinsic characteristics, that messages have different lengths and injection rates in general purpose CMPs, are not fully considered, especially for cache-coherent non-uniform memory accessing (CC-NUMA), which requires additional messages for maintaining cache coherence. As we observe, read operations are the major part of messages in NoC, whereas other operations are usually much less than the read operations. Consequently, the workloads in MPC and VCs are unbalanced. It is necessary to combine some low usage channels and allocate channels with different widths according to the message characteristics.

In this paper, we propose multiple-combinational-channel (MCC), a heterogeneous and combinational network architecture for cache-coherent CMPs [21]. As shown in Fig. 1(d), MCC combines some low utilization channels into one combinational channel. As a result, MCC is more balanced and energy efficient than packet-switching with VCs, and compared to MPC, MCC saves wires and buffers. Furthermore, MCC keeps high utilization channels independent and wide enough, hence the performance and throughput of MCC are not reduced. Moreover, because MCC needs no VC or data packing, the interface between cores and routers is simple. Therefore, MCC is not only efficient for large-scale CMPs, but also suitable for small and medium-scale SoCs.

As many messages travel through the interconnect network simultaneously, deadlocks may happen without careful design. In this paper, we fully utilize the characteristics of network traffic, and classify the network traffic into four standard patterns. With the four patterns, we discover the dependence among different messages, and summarize several rules to maintain the right order of messages. By following these rules, MCC can avoid protocol-level deadlocks.

The key contributions of this paper are summarized as follows:

(1) We propose a novel interconnect network architecture, MCC, with combinational and heterogeneous channels. MCC is not only a balanced design between performance and power consumption, but also a balanced architecture for network traffic.
(2) We classify traffic in CC-NUMA into four basic patterns, and establish several rules to avoid protocol-level deadlocks based on the analysis of the four patterns. Furthermore, these rules are integrated into MCC.
(3) We apply MCC to a 16-core CC-NUMA. And experimental results demonstrate that MCC is more balanced compared to MPC, and saves area and power with similar throughput. Compared to packet-switching with VCs, MCC achieves higher performance, higher energy efficient and more balanceable workload.

The rest of this paper is organized as follows. Section 2 introduces the related works about interconnect architectures in NoCs. In Section 3, the characteristics of traffic in CC-NUMA are analyzed. The MCC architecture is proposed in Section 4, and the experimental results of our architecture are presented in Section 5. Finally, conclusions and future work are outlined in Section 6.

2. Related works

Recently, many literatures have been proposed to optimize every aspect of the interconnect network. In this paper, we focus on the architecture of network, especially the channel architecture. Previous works about the network architecture falls into two categories: bus-based network and NoC.

Bus-based network is an efficient architecture in the early stage of CMPs. However, due to increasing integration degree, traditional bus architecture [9] for chip communication limits bandwidth utilization and scalability. Some recently research works claimed that bus-based interconnects with global shared wires had some benefits [10]. However, these bus-based networks were only suitable for small or medium-scale SoCs.

NoC had been proposed as an alternative solution to bus-based network to provide wider communication bandwidth and better scalability. A packet-switching NoC was proposed in [11] to replace design-specific global on-chip wiring. However, packet-switching NoCs, which are inefficient for long packets, still could not provide enough bandwidth, and have a risk of deadlock.

VCs have been introduced into NoCs to avoid deadlocks and improve bandwidth [12]. At present, the packet-switching NoC with VCs is a widely-used network architecture, and many works continue optimizing the NoC architecture on this basis. Some researches [3,4] attempted to present NoCs based on packet-switching with VCs that could support a thousand connected components with high energy efficiency and good performance. A low-latency VC router with a regular chip-wide network was proposed in [2]. Elastic buffers [13,14] and bufferless [15] architectures were proposed to support multiple VCs and minimize buffering requirements with some performance loss. In [16], adding a small side buffer was proved to be more efficient than the technology of bufferless. Kim et al. [17] presented a hybrid packet-flit flow control and a packet quota system, which enable a lightweight router with support for VCs while simplifying the switch arbitration. However, packet-switching NoCs with VCs are inefficient for small or medium-scale CMPs and various messages with different lengths in CC-NUMA.

MPC was proposed for bus-based architecture [6] at first, but it could also be used in NoC [18]. MPC has multiple physical channels, which is fit for heterogeneous data transmissions and provides much higher bandwidth than VCs. Four separate and independent networks (*i.e.* two statical, two dynamical) were proposed in the RAW processor [7]. Balfour and Dally [8] proposed that splitting network traffic into two heterogeneous or homogeneous networks could improve performance and power efficiency. However, due to the large area and high power, MPC has poor scalability and could not be implemented in power-limited SoCs, like embedded systems.

Although both VCs and MPC were used to avoid deadlocks in many previous works, these methods still have potential deadlock risks without careful design. We present two concrete examples in Section 4.