J. Parallel Distrib. Comput. 74 (2014) 2818–2830

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

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

Hybrid circuit-switched network for on-chip communication in large-scale chip-multiprocessors

Hongyin Luo [a], Shaojun Wei [b], Deming Chen [c], Donghui Guo [a,*]

[a] Department of Electronic Engineering, Xiamen University, Fujian 361005, China
[b] Institute of Microelectronics, Tsinghua University, Beijing, China
[c] Department of Electrical and Computer engineering, University of Illinois at Urbana–Champaign, IL, USA

HIGHLIGHTS

- Circuit switching and packet switching are both provided in the HCS network with a hybrid scheme.
- The channel preserver provides circuit switching by actively analyzing the packet head info.
- The router in the HCS network is a buffer-less design, and a packet is transferred in a pipelined manner.
- An AMBA wrapper is implemented to enable the on-chip network to behave like a bus for AMBA IPs.
- Burst and locked transmissions are provided based on the circuit-switched scheme.

ARTICLE INFO

Article history:
Received 9 October 2012
Received in revised form
11 February 2014
Accepted 7 May 2014
Available online 20 May 2014

Keywords:
Chip-multiprocessors
Network-on-chip
Hybrid packet-circuit switching
AMBA compatibility
Pipeline channel

ABSTRACT

Large-scale chip-multiprocessors (CMPs) need a scalable communication structure characterized by low cost, low power, and high performance to meet their on-chip communication requirements. This paper presents a hybrid circuit-switched (HCS) network for on-chip communication in the large-scale CMPs. The HCS network, which is Advanced Microcontroller Bus Architecture (AMBA) compatible, is composed of bufferless switches, pipeline channels, and network interfaces. Furthermore, packets are transferred in a hybrid transmission scheme. If a message has only one packet, the transmission scheme for this message is packet switching. Conversely, if a message contains multiple packets, the transmission scheme for this message is circuit switching. We evaluate HCS networks with different channel depths and then compare the HCS network with the Stanford elastic buffer (EB) network. Our results show that the HCS network with two-depth channel requires 83% less power and occupies 32% less area compared with the EB network. Furthermore, under maximum frequency and single traffic, the HCS network with two-depth channel provides 37% lower zero-load latency, 390% higher maximum throughput per unit power, and 19% higher maximum throughput per unit area compared with the EB network.

1. Introduction

Technology scaling continuously improves transistor performance and transistor integration capacity [7]. Considering the so-called "power wall" problem, improving the execution frequency of microprocessors (with a single CPU core) through technology scaling is no longer feasible. Instead, chip-multiprocessors (CMPs) have emerged as an alternative architecture to improve processor performance through parallelism [8]. Small-scale CMPs, such as Tegra 3 [30] or OpenSPARC T2 [32], contain a few CPU cores. To interconnect computing elements, memories, and peripherals, a small-scale CMP normally uses shared buses or point-to-point (P2P) links as the on-chip communication structure. However, given that shared buses scale poorly on performance and P2P links scale poorly on area and energy [19,16], these components are unsuitable for use in large-scale CMPs. Thus, large-scale CMPs, such as Intel 48-core SCC [18] or Tilera TILEPro64 [5], require a scalable communication structure to meet their on-chip communication requirements.

Network-on-chip (NOC) [6,9,34,12], which uses techniques similar to computer networks for providing high communication scalability and throughput, has emerged as a promising solution for the on-chip communication problem of large-scale CMPs. In an NOC, computing elements, memories, and peripherals are

* Corresponding author.
E-mail addresses: hyluo@xmu.edu.cn (H. Luo), dhguo@xmu.edu.cn (D. Guo).

http://dx.doi.org/10.1016/j.jpdc.2014.05.003
0743-7315/

interconnected through a network comprising numerous communication nodes. Moreover, data packets are transferred between these nodes. Many NOC designs, such as the MIT RAW [33], TRPS [14], Intel 80-core Terascale chip [17], Intel 48-core SCC [18], or Tilera TILEPro64 [5], have served as bases for prototype systems used to evaluate NOC performance. These NOC designs use routers with input buffers to construct the on-chip network. Meanwhile, virtual channels (VCs) [18] and multi-networks [5] were also proposed to boost network throughput. The experimental results of these NOC designs show that the on-chip network is a major factor in the performance, area, and power consumption of the overall system. Low cost, low power, and high performance are the three key design targets to achieve the practicality of NOC [31].

Considering that metal wires in integrated circuits usually consist of many layers, wires are available in abundance in the on-chip space. On the other hand, the cost of logic units is significant because these units are placed in the two-dimensional on-chip space [9]. Thus, an on-chip network with low cost should use wide communication channels and few logic units to provide low latency and high throughput. The elastic buffer (EB) network [22] is based on such a design technique.

EB, EB router, and EB flow control [22–24] have recently been proposed for on-chip network design. An EB is actually a flip-flop (FF) with latches that are controlled by the EB control module. The master and slave latches of the EB, the enable pins of which are qualified by the clock signal, independently provide two storage locations. An EB router uses EB, instead of first-in–first-out (FIFO), as the input and output buffer. Furthermore, the EB flow control scheme ensures that the packets are accurately transferred between EBs. Refs. [22] and [23] demonstrated that EB flow control is more efficient than the VC [15] and wormhole flow control [28]. Furthermore, [23] also showed that the single-stage EB router outperforms the wormhole router in terms of area, energy consumption, zero-load latency, and throughput efficiency.

However, EB design is not power efficient compared with the FF. Based on the design in [23], we observe that the power consumption of an EB is higher than that of a regular FF. This power inefficiency is attributed to the fact that: an EB is a latch-based design. On the other hand, the EDA tool, which constructs a circuit design based on the FF technique, cannot optimize the latch-based design well.

We recently proposed a hybrid circuit-switched (HCS) on-chip network for large-scale CMPs [20]. The HCS network, which is composed of switches without buffers, pipeline channels, and network interfaces, is especially designed for Advanced Microcontroller Bus Architecture (AMBA) based CMPs. The HCS network has the same AMBA interface as the advanced high-performance bus (AHB). Thus, the HCS network is AMBA-compatible. In this paper, we present the detailed architecture of the HCS network, the mapping of the AMBA protocol onto the on-chip network, and a new transmission scheme in the HCS network, that is, the full packet-switched scheme. We evaluate HCS networks with different channel depths, EB networks with different channel widths, and the full packet-switched scheme in terms of area, power consumption, zero-load latency, maximum throughput, and throughput efficiency. We also compare the HCS network with the Stanford EB network [23].

The key contribution of this paper lies in two aspects: the mapping of the AMBA protocol onto an NoC and the hybrid circuit-switched scheme. **First**, to solve the standard communication interface issue for large-scale CMPs with on-chip network, the AMBA protocol is mapped onto the HCS network and then used as the standard communication protocol. Hence, AMBA IPs with AMBA 2.0 protocol could be connected to the HCS network through the network interface. **Second**, by combining the advantages of circuit switching and packet switching, the hybrid circuit-switched scheme achieves the improvement in performance and the reductions in area and cost, as well as makes the burst and locked transmissions defined in AMBA available.

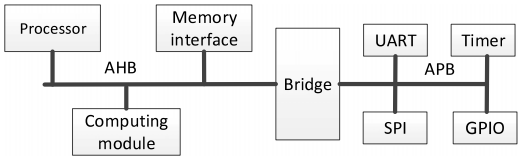

Fig. 1. Typical AMBA-based SOC system.

The remainder of this paper is organized as follows: Section 2 describes the mapping of the AMBA protocol onto the on-chip network. Section 3 discusses the architecture, switch, and pipeline channel of the HCS network. Section 4 describes the evaluation methodology. The results and discussions are given in Section 5. Section 6 presents related work. Finally, Section 7 concludes this paper.

2. Mapping of AMBA protocol onto on-chip network

2.1. Background

CMPs have two types of modules: the computation and communication modules. For enhancing design efficiency, these two types of modules are separated by a standard interface and are designed independently. In a small-scale CMP with a shared bus or P2P link, a standard communication protocol, such as AMBA protocol [3] or WISHBONE protocol [1], is used to maintain the consistency of communication interfaces within different computation modules. However, for large-scale CMPs with an on-chip network, no standard communication protocol exists. To use the well-designed AMBA IP, we map the AMBA protocol onto on-chip network.

The AMBA protocol defines a standard communication interface for bus [3] or P2P link [4]. In this paper, only the AMBA 2.0 described in [3] is considered. Under this circumstance, two kinds of bus are typically used: the AHB and the advanced peripheral bus (APB). The AHB is used for high-performance interconnection between multiple system modules including processors, memory interfaces, and high-performance computation modules. The APB is for low-power interconnection between peripherals.

In a CMP with shared bus, the communication interface can be categorized as either a master or a slave interface. A device with a master interface (master device) launches a message transfer, whereas a device with slave interface (slave device) responds to this transfer by providing data or response signals. A typical AMBA-based system-on-chip is illustrated in Fig. 1, where the AHB and APB are connected through a bridge. High-performance modules are generally interconnected through AHB. AHB contains a fully functional arbiter and decoder and can connect any system module with the master or slave interface. On the other hand, in APB, the AHB–APB bridge is the only device with a master interface. Thus, APB contains only a decoder and can only connect modules with the slave interface.

In a CMP with shared bus, data are normally transferred through the following three steps: (1) a master device requests the ownership of the bus channel; (2) the master device is granted permission to occupy this bus channel, after which the data and control signals are delivered to the corresponding slave device through the bus; (3) after the requested slave device responds to the master device with a message, the transmission is finished. This bus channel is then released and is free to be used by other devices.

Furthermore, in AMBA 2.0, more transmission features are provided by AHB, including burst, locked, and split transmission. In burst transmission, multiple data are transferred (read or write) in a streaming scheme. During this period, the bus is not released