ELSEVIER

Computer Communications 24 (2001) 1578–1588

computer communications

www.elsevier.com/locate/comcom

Performance analysis of *BusNet* protocol for backplane bus-based interprocessor communication

Minyoung Sung*, Naehyuck Chang, Jinsung Cho, Heonshik Shin

School of Computer Science and Engineering, Seoul National University, Seoul 151-742, South Korea

Received 27 January 2000; revised 19 December 2000; accepted 20 December 2000

Abstract

Nowadays, backplane bus-based multiprocessor systems often utilize the standard network protocol such as TCP/IP for communication between processors on the backplane bus. In such systems, it is common for the backplane bus to emulate the standard MAC protocols such as CSMA/CD. This paper aims to analyze the delay performance of the MAC emulation-based backplane network by constructing queueing models based on detailed bus operations. For this purpose, we choose BusNet as a target protocol. BusNet is an ANSI standard network protocol and its specification contains basic operations commonly used in most backplane buses. We investigate the throughput-delay characteristics in terms of packet size, block transfer scale, and arbitration scheme. We also compare the packet delay in BusNet with the IEEE 802.3 CSMA/CD network which BusNet is expected to be compatible with. The simulation result shows how an optimal block transfer scale can be determined in respect of the performance trade-off between BusNet and other real-time traffics.

Keywords: Backplane bus network protocol; Performance analysis; Packet transfer delay; BusNet

1. Introduction

Backplane buses are widely used for interconnection between processors, memory subsystems, and I/O devices in multiprocessor systems. Multiprocessor systems interconnected over a backplane bus provide cost-effective solutions for a range of parallel and distributed computing applications. Processors in such systems typically communicate with each other by accessing directly the distributed and/or unified shared memory through the backplane bus protocol. Recently, with rapid advances in computing and communication technologies, the application software is commonly involved in complicated interprocessor communication. As a result, the modern trend is to use standard network protocols such as TCP/IP on the backplane bus. This gives advantages such as portability, robustness, and rapid prototyping at the expense of slight performance degradation. These *backplane network protocols* commonly emulate the standard medium access control (MAC) protocols such as CSMA/CD in order to interface the existing upper protocol stack, i.e. IP. This approach gives an additional advantage in costs because application software can utilize standard network protocols without equipping every processor board with a network interface card (NIC). Despite the growth of usage and the benefit of backplane network protocols, however, there has been little work regarding the performance of the backplane bus as a communication medium. Eventhough the bandwidth of the backplane bus is generally larger than that of the local area network, detailed performance analysis is essential in order to guarantee both timely and logically correct behavior. In this paper, we develop and validate an analytic model to study the MAC and link layer characteristics of the backplane network.

Numerous works have been carried out concerning the performance of the MAC protocols, but have mostly focused on the conventional local area networks such as the token passing and CSMA/CD networks. In particular, CSMA/CD has been extensively analyzed for the past two decades [1,9,11,14,15,18]. Compared with the local-area networks, however, the backplane bus exhibits quite a different behavior in delay performance due to the physical channel characteristics. Several works deal with the physical features of the backplane buses. Bain and Ahuja perform a simulation study on a static priority arbiter, a rotating daisy chain, and an ideal FCFS arbiter [2]. Assuming fixed bus access time and uniformly distributed interrequest time, they measure the mean and the coefficient of variation of the waiting time for the data transfer request.

* Corresponding author. Tel.: +82-2-880-7295; fax: +82-2-874-3104.
E-mail address: minyoung@cselab.snu.ac.kr (M. Sung).

0140-3664/01/$ - see front matter
PII: S0140-3664(01)00298-5

Vernon and Leutenegger analyze the bus arbiters using a timed Petri-net model [19]. Their performance metric is the mean processor efficiency, i.e. the bus bandwidth allocated to each processor. They also discuss the fairness of the arbiters. Woodbury and Shin construct a queueing network model to study a workload effect on performance for a unibus multiprocessor [20]. Assuming prioritized arbitration, they present the waiting time for each priority class traffic. Recently, several works have been presented which investigate the commercial backplane buses in detail. Johnson et al. model the arbitration phase of the Futurebus + as an M/G/1 queue and estimate the performance achievable by the virtual port memory multiprocessor system [7]. Kettler and Strosnider give a formal model for the MCA bus by analyzing the components comprising the data transfers [8]. A similar approach for the Controller Area Network (CAN) bus is taken by Tindell et al. [17]. However, most of these researches do not give a suitable model for the packet-based communication over the backplane bus. They lack in consideration for the packet-oriented data transfer or concentrate on estimating bus bandwidth allocated to each processor. For the communication performance, their analytic models are inadequate or need further investigations.

In this paper, we aim to analyze the packet transfer delay in the backplane bus network by taking into account the physical bus channel characteristics. Assuming each communicating node is associated with a queue of infinite capacity, we devise queueing systems where a packet request is served by a group of basic bus transfers called *transactions*. Unlike traditional packet service models where the entire packet transmission is non-preemptive by nature, packet transmission on the backplane bus can be preempted by other ready nodes since the bus is shared on the transaction basis. The accuracy of the analytic model is validated by a simulation study using a bus simulator which is constructed to reflect bus details. Using the analytic model, we discuss the throughput-delay performance of the backplane network in comparison with that of the conventional CSMA/CD network. We also present extensive experimental results obtained for various values of packet size, block transfer scale, etc. For generality, we choose *BusNet* [22–24] as a target protocol. BusNet is an MAC and link-layer protocol specification for the standardized communication over the *VMEbus*. Since its specification contains only primitive features commonly used in most backplane hardware, the analysis can be extended to different kinds of backplane buses without any difficulty.

The rest of this paper is organized as follows. In Section 2, we introduce the concept of backplane bus network protocol. We also briefly describe the arbitration schemes and data transfer mechanism in VMEbus. In Section 3, we analyze the packet transfer time in BusNet. Section 4 presents numerical results of the analytic model and discusses the communication performance. This paper ends with concluding remarks in Section 5.

2. Backplane bus network protocol

2.1. BusNet: a standard backplane network protocol

At present, most embedded operating systems support standard network protocol over the backplane bus as well as over the NIC. As shown in Fig. 1, a software module emulates the standard MAC protocol, typically the Ethernet, between the backplane hardware and the ISO/OSI network layer. It consists of data structures, protocol states, and transfer primitives for packet-oriented communication. While most backplane network protocols are proprietary and platform dependent, some of them have been approved as international standards.

BusNet is an ANSI standard protocol originally developed by Force Computers Inc. BusNet emulates the CSMA/CD protocol on the VMEbus. Unlike other backplane network protocols, BusNet assumes only basic hardware features; the specification does not rely on special features such as read-modify-write or mailboxes. Fig. 2 shows a system using BusNet. Each processor or *participant* is identified by a logical address ranging from 0 to 30. It shares a segment of its local memory, called *BusNet region*, for the access from other participants. Each BusNet region consists of a BusNet header, participant maps, and packet buffers. The participant map contains protocol descriptor (PD), which includes status flags and pointers for the packet exchange between two peers. The PD is composed of `receive_status`, `transmit_status`, `buffer_offset`, `buffer_size`, and `sequence_number`. The `receive_status` is used by the receiver to control the access to the packet buffer. When the `receive_status` is in `RDY` state, the transmitter may write a packet to the buffer in the receiver. The `transmit_status` indicates to the receiver that the packet buffer has been filled with a new packet. The `buffer_offset` and `buffer_size` contain the address and size of the packet buffer, respectively. The `sequence_number` is included for error checking.

Table 1 illustrates the packet transmission steps. The `map_i` describes the participant map located in the shared memory of participant `i`. It is noteworthy that the entire transmission process is based solely on write cycles for efficiency. Checking the status is a local operation since the `receive_status` and `transmit_status` flags

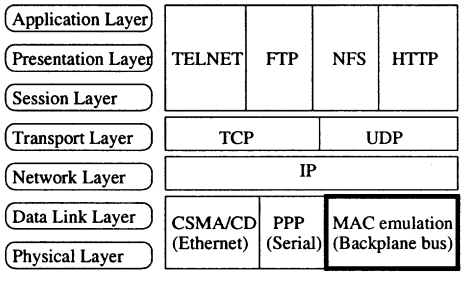

Fig. 1. The backplane bus network and the ISO/OSI model.