

# Performance evaluation of directory protocols on an optical broadcast-based distributed shared memory multiprocessor

İpek Abasıkeleş, M. Fatih Akay\*

Computer Engineering Department, Çukurova University, 01330 Adana, Turkey

## ARTICLE INFO

### Article history:

Received 24 January 2009

Received in revised form 7 June 2009

Accepted 25 June 2009

Available online 6 August 2009

### Keywords:

Directory protocols

Distributed shared memory

Multiprocessors

Performance evaluation

Optical interconnects

Simulation

## ABSTRACT

Recent advances in the development of optical technologies suggest the possible emergence of broadcast-based optical interconnects within cache-coherent distributed shared memory (DSM) multiprocessor architectures. It is well known that the cache-coherence protocol is a critical issue in designing such architectures because it directly affects memory latencies. In this paper, we evaluate via simulation the performance of three directory-based cache-coherence protocols; strict request-response, intervention forwarding and reply forwarding on the Simultaneous Optical Multiprocessor Exchange Bus (SOME-Bus), which is a low-latency and high-bandwidth broadcast-based fiber-optic interconnection network supporting DSM. The simulated system contains 64 nodes, each of which has a processor, a cache controller, a directory controller and an output channel. Simulations have been conducted for each protocol to measure average processor utilization, average network latency and average number of packets transferred over the network for varying values of the important DSM parameters such as the ratio of the mean channel service time to mean thread run time ( $T/R$ ), probability of a cache block being in modified state ( $P(M)$ ), the fraction of write misses ( $P(W)$ ) and home node contention rate. The results reveal that for all cases, except for low values of  $P(M)$ , intervention forwarding gives the worst performance (lowest processor utilization and highest latency). The performance of strict request-response and reply forwarding is comparable for several values of the DSM parameters and contention rate. For a contention rate of 0%, the increase of  $P(M)$  makes reply forwarding perform better than strict request-response. The performance of all protocols decreases with the increase of  $P(W)$  and contention rate. However, the performance of strict request-response is the least affected among other protocols due to the negative impact of the increase of  $P(W)$  and contention rate. Therefore, for the full contention case (i.e. contention rate of 100%); for low values of  $P(M)$ , or for mid values of  $P(M)$  and high values of  $P(W)$ , strict request-response performs better than reply forwarding. These results are significant in the sense that they provide an insight to multiprocessor architecture designers for comparing the performance of different directory-based cache-coherence protocols on a broadcast-based interconnection network for different values of the DSM parameters and varying rates of contention.

© 2009 Elsevier Ltd. All rights reserved.

## 1. Introduction

As we slowly approach the physical limits of processor and memory speed, it is becoming more attractive to use multiprocessors to increase computing power. Two kinds of parallel processors have become popular: tightly coupled shared

\* Corresponding author.

E-mail address: [mfakay@cu.edu.tr](mailto:mfakay@cu.edu.tr) (M.F. Akay).

memory multiprocessors and distributed-memory multiprocessors [1]. DSM systems are a successful hybrid of these two parallel computer classes. They provide the shared memory abstraction in systems with physically distributed memories and consequently combine the advantages of both approaches [2]. The key strength of DSM systems is that communication occurs implicitly as a result of conventional memory access instructions (i.e. loads and stores), which makes them easier to program [3]. Also, these systems are claimed to be scalable because they can achieve an almost linear performance increase with growing system size and investment. Because of these advantages, DSM is one of the most attractive solutions for building large-scale, high performance multiprocessor systems [4].

One of the fundamental communication problems in DSM systems that significantly affects scalability, is the increase in remote memory access latency as the number of processors increase in the system. Caches play a vital role in reducing latency by maintaining coherent copies of shared data locally [5]. Large-scale shared memory multiprocessors use caches to reduce average memory latency and increase effective network bandwidth. Unfortunately, the existence of multiple caches in a shared memory environment can lead to data inconsistencies in the system. This is referred to as the cache-coherence problem. Any system that allows multiple copies of a data item to exist at the same time must solve this problem [6].

The two most common types of cache-coherence protocols are snooping and directory-based protocols. Directory protocols maintain a directory entry per memory block that records which processor(s) currently cache the block. On a miss, a processor sends a coherence message over an interconnect to a directory, which often forwards message(s) to processor(s) currently caching the block. These processors may forward data or acknowledgments to the requesting processor and/or directory [7]. The two major performance goals of a directory protocol are: (i) to reduce the number of network transactions generated per memory operation, which reduces the bandwidth demands placed on the network and the communication assists; and (ii) to reduce the number of actions, especially network transactions, that are on the critical path of the processor, thus reducing uncontended latency. Network latency is important because memory access time depends on that latency [8]. Directory-based coherence protocols offer a scalable performance path beyond snooping-based ones by allowing a large number of processors to share a single global address space over physically distributed-memory [9]. Most large shared memory multiprocessors use directory protocols [10].

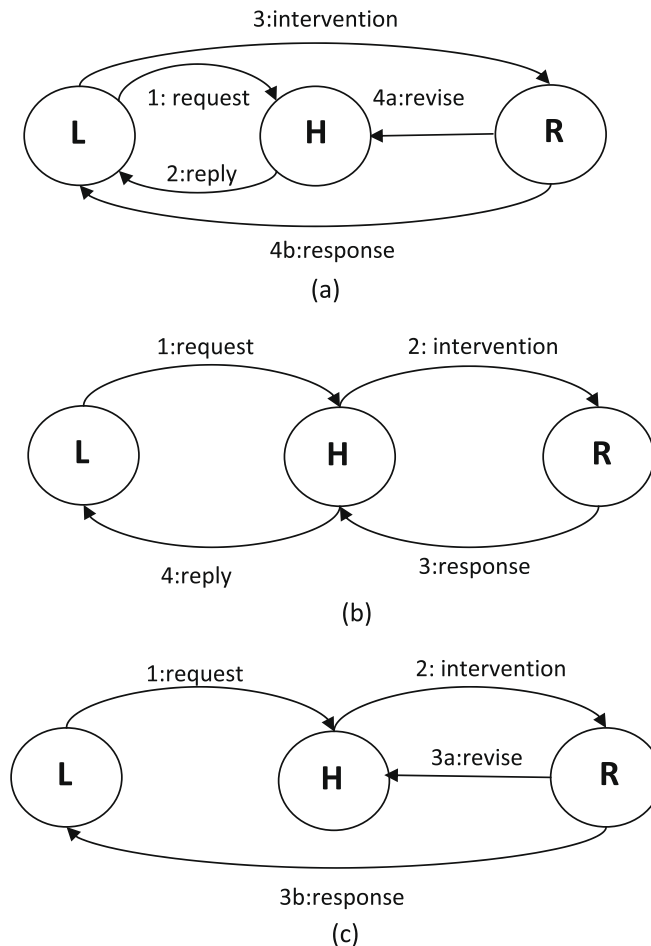


Fig. 1. Directory protocol diagrams for a modified cache state (a) strict request-response, (b) intervention forwarding and (c) reply forwarding.

Download English Version:

<https://daneshyari.com/en/article/455544>

Download Persian Version:

<https://daneshyari.com/article/455544>

[Daneshyari.com](https://daneshyari.com)