

Performance of Multiple-Bus Multiprocessor under Non-Uniform Memory Reference

M.A. Sayeed

Department of Electrical Engg.
King Fahd University of
Petroleum & Minerals
Dhahran 31261, Saudi Arabia.

M. Atiquzzaman

Department of Computer Science
and Computer Engineering
La Trobe University
Melbourne 3083, Australia.
atiq@latcsi.lat.oz.au

Abstract

Performance evaluation of multiple-bus multiprocessor systems is usually carried out under the assumption of uniform memory reference model. The objective of this paper is to study the performance of multiple bus multiprocessor system in the presence of hot spots. Analytical expressions for the average memory bandwidth and probability of acceptance of prioritized processors have been derived. Two new phenomenon, coined as bumping and knee effect, have been observed in the acceptance probabilities of the processors. The results are validated by simulation results.

1 Introduction

A variety of interconnection networks have been proposed and analyzed [1] in the literature. Because of modularity and fault tolerance of multiple bus systems, such systems have been widely investigated [2]. Bus and memory contentions limit the performance of multiple bus systems. In analyzing the performance, most authors assume a *uniform memory reference (URM)* model where all the memories are equally likely to be accessed by the processors [3]. Memory references in a multiprocessor system are not necessarily uniform. Favorite and localized memory references have been analyzed by several researchers.

A particular type of non-uniform memory reference arises in multiprocessor systems due to the use of variables for locking, global and barrier synchronization, pointers to shared queues, etc. The phenomenon is called hot spot contention and it degrades the performance of multiprocessor systems. Performance of crossbar multiprocessors, unbuffered MINs, and buffered MINs in the presence of hot spots have been developed in [4, 5, 6, 7]. The objective of this paper is to determine the performance of a multiple bus interconnection network under hot spot conditions. Analytical expressions for *memory bandwidth* and the *probability of acceptance* of processor requests are derived for processors having *different priorities* in the case of conflicts. It is shown that the bandwidth is affected by a change in the hot spot probability for high request rates. Moreover, for high request rates and a small number of buses, the degradation at low hot spot probabilities is mainly due to bus conflicts,




Figure 1: Multiple-bus multiprocessor systems.

and memory conflicts play an insignificant role in the degradation. We have observed two new phenomenon in the acceptance probabilities of the requests from prioritized processors. We call these the *bumping effect* and the *knee effect*. With an increase in the hot spot probability, the bumping effect results in an increase and decrease in the acceptance probabilities of the low and high priority processors respectively. The knee effect explains the fact that a system having prioritized processors may appear to be a crossbar system for some processors, while it may appear as a multiple-bus system for the rest of the processors. Results obtained from analytical model have been verified by simulation results, and have been found to be in close agreement.

The modeling assumptions are given in Section 2. Analytical models for determining the average memory bandwidth and the probability of acceptance of processor requests are presented in Section 3, followed by results in Section 4.

2 Assumptions

We model the multiple-bus system under the following assumptions.

- The system consists of N identical processing elements (PE), M identical memory modules (MM), and $B \leq \min(N, M)$ buses (see Figure 1). The processing elements and memory modules will be represented by $\{PE_0, PE_1, \dots, PE_{N-1}\}$ and $\{MM_0, MM_1, \dots, MM_{M-1}\}$ respectively.
- The system operates *synchronously*. Each processor generates a request at the start of a memory

cycle with probability r .

- Processor requests are characterized by temporal and spatial independence. Rejected requests are discarded.
- Processors have *priority* in the case of bus and memory conflicts.
- Memory module M_h is a hot memory for all PEs. The probability of PE i , $0 \leq i < N-1$, requesting MM_h is p_h , and of requesting MM_j , $j \neq h$, is given by $p_0 = (1 - p_h)/(M-1)$, where $p_h > 1/M$.

3 Performance Evaluation

Average memory bandwidth and probability of acceptance will be used as the measures of performance.

3.1 Average memory bandwidth

Average memory bandwidth (AMBW) of a multiprocessor system is defined as the average number of memory modules accessed simultaneously during a memory cycle. The average memory bandwidth of an N processor, M memory, B bus system with a processor request rate of r will be denoted by $AMBW(N, M, B, r)$. AMBW of a multiple-bus system for $r = 1$ is given by

$$AMBW(N, M, B, 1) = \sum_{k=1}^B kPr(k) + \sum_{k=B+1}^{\min(N, M)} BPr(k) \quad (1)$$

where $Pr(k)$ is the probability that exactly k memory modules are accessed in a cycle. For a URM, $Pr(k)$ is given [3] by $Pr(k) = \frac{k!S(N, k)}{M^N} \binom{M}{k}$, where $S(N, k)$ is the Stirling number of the second type.

3.1.1 Request Probability=1

First, let's consider $r = 1$. Let $E_k^{p, q}$ be the event that k distinct memory modules are requested during a cycle, where p and q are the number of references for MM_h and MM_j , $j \neq h$, respectively. Hence, $p + q = N$ for $r = 1$. Depending on the distribution of requests to hot and non-hot memories, we can divide $E_k^{p, q}$ into two classes.

Class 0: All the requests are to non-hot memory modules. This event will be denoted by $E_k^{0, N}$.

Class 1: i requests are to the hot memory, and $N - i$ requests to the non-hot memories. This event will be denoted by $E_k^{i, N-i}$, $i \neq 0$.

Probability that k distinct memory modules are requested during a cycle is, therefore, given by the sum of the probabilities of the two classes. It can be shown [4] that

$$Pr(k) = k!S(N, k) \binom{M-1}{k} p_0^N + (k-1)! \left(\binom{M}{k} - \binom{M-1}{k} \right) \sum_{i=1}^{N-k+1} S(N-i, k-1) \binom{N}{i} p_h^i p_0^{N-i} \quad (2)$$

Substituting Equation (2) in (1) gives the expression for the average memory bandwidth for $r = 1$.

3.1.2 Request Probability < 1

For $r < 1$, the probability of having exactly n PEs requesting memory modules at the beginning of a memory cycle is $\binom{N}{n} r^n (1-r)^{N-n}$. The conditional probability of having k distinct memory modules being requested given that n memory requests have been generated in a cycle is given by

$$Pr(E_k^{0, n} | n \text{ requests}) = k!S(n, k) \binom{M-1}{k} p_0^n \quad (3)$$

Therefore, Class 0 probability for $r < 1$ is given by

$$Pr(E_k^{0, *}) = \sum_{n=1}^N k!S(n, k) \binom{M-1}{k} p_0^n \binom{N}{n} r^n (1-r)^{N-n} \quad (4)$$

Similarly, the Class 1 probability of requesting k distinct MMs such that MM_h is one of them is given by replacing N by n in Equation (6) in [4] and multiplying it by the probability of n PEs requesting memories and $(N-n)$ PEs not requesting.

$$Pr(E_k^{i, *}) = \sum_{n=1}^N \left(\binom{M}{k} - \binom{M-1}{k} \right) \sum_{i=1}^{n-k+1} (k-1)! S(n-i, k-1) \binom{n}{i} p_h^i p_0^{n-i} \binom{N}{n} r^n (1-r)^{N-n}, \quad j \neq 0 \quad (5)$$

Average memory bandwidth is, therefore, given by

$$\begin{aligned} AMBW(N, M, B, r) = & \sum_{k=1}^B k \left\{ \sum_{n=1}^N k!S(n, k) \binom{M-1}{k} p_0^n \binom{N}{n} r^n (1-r)^{N-n} + \right. \\ & \sum_{n=1}^N \left(\binom{M}{k} - \binom{M-1}{k} \right) \sum_{i=1}^{n-k+1} (k-1)! S(n-i, k-1) \\ & \left. \binom{n}{i} p_h^i p_0^{n-i} \binom{N}{n} r^n (1-r)^{N-n} \right\} + \sum_{k=B+1}^{\min(N, M)} B \left\{ \sum_{n=1}^N k! \right. \\ & S(n, k) \binom{M-1}{k} p_0^n \binom{N}{n} r^n (1-r)^{N-n} + \sum_{n=1}^N \left(\binom{M}{k} - \right. \\ & \left. \binom{M-1}{k} \right) \sum_{i=1}^{n-k+1} (k-1)! S(n-i, k-1) \binom{n}{i} p_h^i p_0^{n-i} \\ & \left. \binom{N}{n} r^n (1-r)^{N-n} \right\} \\ = & \sum_{n=1}^N \binom{N}{n} r^n (1-r)^{N-n} AMBW(n, M, B, 1) \quad (6) \end{aligned}$$

If the average memory bandwidth of a system with $r = 1$ is known, Equation (6) can be used to calculate the average memory bandwidth of a system with $r < 1$.

3.2 Probability of acceptance

Probability of acceptance, denoted by P_a , is defined as the probability that a PE's request is accepted. If all the PEs have equal probability of being accepted in the case of bus or memory conflicts, P_a is given by $P_a = \frac{AMBW(N,M,B,1)}{Nr}$.

Processors are often prioritized [3]. The probability of acceptance for the n -th PE will be denoted by $P_a(n)$. We assume that PE _{j} has a higher priority than PE _{$j+1$} , for $j = 0, \dots, N-2$. PE _{n} can therefore, be blocked only by processors PE _{i} , $0 \leq i < n-1$. Let's assume that PE _{n} and i other PEs of higher priority than PE _{n} generate requests during a memory cycle. The analysis will be divided into two cases.

Bus Sufficient System (BSS): There is no possibility of PE _{n} being blocked due to bus conflicts, i.e., the system is either a crossbar, or the system is not a crossbar but $i < B$. It may be blocked due to memory conflicts with a higher priority PE.

Bus Deficient System (BDS): PE _{n} may be blocked due to bus and/or memory conflicts, i.e., the system is not a crossbar and $i \geq B$.

3.2.1 Analysis of Bus Sufficient System (BSS)

Since there is no bus conflict, $P_a(n)$ is equal to the probability that no other PE(s) of higher priority request the memory module requested by PE _{n} . We can further subdivide this case into two subcases.

Case BSS-H: PE _{n} requests the hot MM and the request is accepted.

Case BSS-NH: PE _{n} requests a non-hot MM and the request is accepted.

$P_a(n)$ for the BSS case is therefore, the weighted sum of the probabilities of cases BSS-H and BSS-NH. Therefore,

$$P_a(n)|_{\text{BSS}} = p_h P_a(n)|_{\text{BSS-H}} + (1 - p_h) P_a(n)|_{\text{BSS-NH}} \quad (7)$$

Analysis of Case BSS-H

The probability that PE _{n} requests MM _{h} and gets it is equivalent to the probability that none of the n higher priority PEs request MM _{h} . The probability that a particular set of i higher priority PEs do not request MM _{h} and the rest $(n-i)$ higher priority PEs do not generate a memory request is $r^i(1-p_h)^i(1-r)^{n-i}$. The set of i PEs can be chosen out of n PEs in $\binom{n}{i}$ ways, and i ranges from 0 to n . Therefore, the probability of acceptance for case BSS-H is given by

$$P_a(n)|_{\text{BSS-H}} = \sum_{i=0}^n \binom{n}{i} r^i (1-r)^{n-i} (1-p_h)^i \quad (8)$$

Analysis of Case BSS-NH

As in the case of BSS-H, let a set of i , $0 \leq i \leq n$, higher priority PEs request memories, and the other $(n-i)$ higher priority PEs do not request any memory during a cycle. Of the set of i requests, let a set of k , $0 \leq k < i$, be directed towards a set of j , $1 \leq j \leq k$, distinct non-hot MMs and the rest $(i-k)$ towards the MM _{h} . The probability of $(n-i)$ PEs not requesting, k PEs

requesting non-hot MMs, and $(i-k)$ PEs requesting the hot MM is $(1-r)^{n-i} (rp_h)^{i-k} (rp_0)^k$. The i PEs can be selected out of the n PEs in $\binom{n}{i}$ ways, and the $(i-k)$ requests can be selected out of i requests in $\binom{i}{i-k}$ ways. The j distinct non-hot MMs can be selected out of the $(M-2)$ non-hot MMs (i.e., excluding the non-hot MM requested by PE _{n}) in $\binom{M-2}{j}$ ways. The k requests can be distributed among the j MMs, such that none of them is empty, in $j!S(k, j)$ ways. Therefore, the probability that PE _{n} requests a non-hot MM and is not blocked by any PE of higher priority is given by

$$P_a(n)|_{\text{BSS-NH}} = \sum_{i=0}^n \binom{n}{i} r^i (1-r)^{n-i} \sum_{k=0}^i \binom{i}{i-k} p_h^{i-k} p_0^k \sum_{j=1}^k j!S(k, j) \binom{M-2}{j} \quad (9)$$

Substituting Equations (8) and (9) in Equation (7) and rearranging terms gives

$$P_a(n)|_{\text{BSS}} = \sum_{i=0}^n \binom{n}{i} r^i (1-r)^{n-i} \left\{ p_h (1-p_h)^i + (1-p_h) + \left(\sum_{k=0}^i \binom{i}{i-k} p_h^{i-k} p_0^k \sum_{j=1}^k j!S(k, j) \binom{M-2}{j} \right) \right\} \quad (10)$$

Equation (10) gives the probability of acceptance of prioritized PEs in a crossbar system, or in a multiple bus system having low processor request rates such that there is no bus conflict, or for PE _{n} where $n \leq B+1$.

3.2.2 Analysis of Bus Deficient System (BDS)

In this case the system is not a crossbar and $i \geq B$. PE _{n} is not blocked if it can be assigned a bus to access the requested MM and none of the i higher priority PEs (requesting memory) request the particular MM requested by PE _{n} . As before, we can sub-divide the analysis into two sub-cases.

Case BDS-H: PE _{n} requests the hot MM and gets it.

Case BDS-NH: PE _{n} requests a non-hot MM and gets it.

Analysis of Case BDS-H:

In a bus deficient system, if PE _{n} generates a request for the hot MM, it will be accepted if the i higher priority requesting PEs request no more than $(B-1)$ distinct non-hot MMs out of the $(M-1)$ non-hot MMs. Following a reasoning similar to that in Section 3.2.1, it can be shown [8] that the probability of PE _{n} requesting MM _{h} and getting accepted is given by

$$P_a(n)|_{\text{BDS-H}} = \sum_{i=0}^n \binom{n}{i} r^i (1-r)^{n-i} p_0^i \sum_{j=1}^{\min(i, B-1)} j!S(i, j) \binom{M-1}{j} \quad (11)$$

Analysis of Case BDS-NH:

In a bus deficient system, if PE_n generates a request for a non-hot MM, the request will be accepted if the i higher priority requesting PEs do not request the particular non-hot MM requested by PE_n and the i requests are directed to no more than $(B - 1)$ distinct MMs. Following a reasoning similar to that in Section 3.2.1, it can be shown [8] that the probability of PE_n requesting a non-hot MM and getting accepted is given by

$$P_a(n)|_{\text{BDS-NH}} = \sum_{i=0}^n \sum_{k=0}^i \binom{n}{i} r^i (1-r)^{n-i} \binom{i}{i-k} p_h^{i-k} p_0^k \sum_{j=1}^{\min(k, B-2+\lfloor k/i \rfloor)} j! S(k, j) \binom{M-2}{j}, k \leq i \quad (12)$$

Substituting Equations (11) and (12) in (7) and rearranging terms, $P_a(n)$ for a bus deficient system is given by

$$P_a(n)|_{\text{BDS}} = \sum_{i=0}^n \binom{n}{i} r^i (1-r)^{n-i} \left\{ p_h p_0^i \sum_{j=1}^{\min(i, B-1)} j! S(i, j) \binom{M-1}{j} + (1-p_h) \sum_{k=0}^i \binom{i}{i-k} p_h^{i-k} p_0^k \sum_{j=1}^{\min(k, B-2+\lfloor k/i \rfloor)} j! S(k, j) \binom{M-2}{j} \right\}, k \leq i \quad (13)$$

4 Results

Figure 2 shows the bandwidth vs. hot spot probability for various processor request rates. The bandwidth is found to decrease with increasing p_h , the reason being the increased contention for MM_h . The degradation is significant for high processor request rates.




Figure 2: Bandwidth vs. probability of hot spot for different request rates.

Figure 3 shows the variation of bandwidth with hotspot probabilities. When $B = 10$ or 8 , the system behaves like a crossbar and the bandwidth is maximum. However, the degradation in bandwidth is significant for $B = 6$ or lower. The bandwidth decreases

with increasing p_h . The degradation is significant when the number of buses is large, because in such cases the degradation is mainly due to memory contention. In systems with fewer number of buses, the contention is mainly for buses, and increased p_h does not have significant effect until p_h becomes very high when memory contention becomes the dominating factor contributing to the degradation.

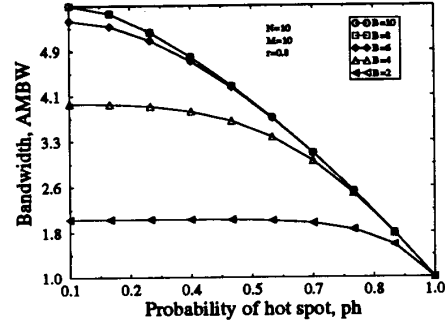


Figure 3: Bandwidth vs. probability of hot spot for different number of buses.

Figure 4 is a plot of probability of acceptance of different processors vs. the hot spot probability. Since PE_0 has the highest priority, it is never blocked. In general, the probability of acceptance of PE_n decreases with increasing n . But, an interesting behavior can be observed for large n , i.e. $n > 7$. The probability of acceptance increases first and then decreases. This behavior can be explained as follows. As the hotspot




Figure 4: Probability of acceptance for different processors vs. probability of hot spot.

probability increases initially, the memory contention among PEs increases and lower priority PEs stand a better chance of getting a bus. When the hotspot probability is considerably higher, even the low priority PEs may request MM_h with a higher frequency and thus the probability of acceptance decreases. This sort of *bumping effect* for lower priority PEs was found for request rates between 1 and 0.6. With decreasing r , the bump shifts towards the left, and is not observed for low values of r .

Figure 5 shows the probability of acceptance of different processors vs request rate. As expected, the probability of acceptance decreases with increasing r . It is observed that at low values of r , the low priority PEs are affected more with an increase in r . For high

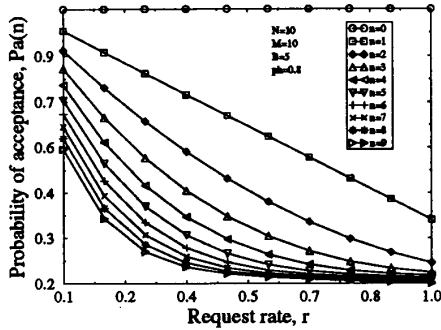


Figure 5: Probability of acceptance for different processors vs. processor request rates.

values of r , the degradation is insignificant for lower priority PEs since they have already reached the bottom line. The lower the priority of the PE, the faster its probability of acceptance reaches saturation as is evident by comparing the curves for PE_1 and PE_9 .

Figure 6 shows the probability of acceptance for different processors vs. the number of buses. We notice




Figure 6: Probability of acceptance for different processors vs. number of buses.

an interesting phenomenon, what we call as the *knee effect*. Again PE_0 is the highest priority PE and is, therefore, always accepted, but PE_1 is accepted approximately 80% of the time. Requests from PE_n , $n < B+1$, encounters only memory contention, and hence the system appears to PE_n as a crossbar system. Therefore, as the number of buses crosses $n-1$, $P_a(n)$ experiences a sharp rise because of the apparent transition from the bus deficient case to the bus sufficient case for PE_n . For PE_n , $n > B+1$, the probability of getting a bus is less due to the number of higher priority PEs being greater than the total number of buses. Therefore, for PE_n there is a rapid change in $P_a(n)$ for values between $B = n+1$ and $B = n+2$. The knee shifts towards the right for low priority processors. Extensive simulation results were found to be in close agreement to the analytical results.

5 Conclusions

We have developed analytical expressions for average memory bandwidth and probability of acceptance

of prioritized processors in a multiple bus system in the presence of hot spot conditions. Effects of different parameters like hot spot probability, processor request rate, and the number of buses on bandwidth and probability of acceptance have been presented.

It has been shown that the bandwidth decreases with increasing percentage of requests to hot memory module. Moreover, the degradation is significantly higher for high processor request rates. We have shown that for a large number of buses, the degradation is mainly due to memory conflicts. On the contrary, for fewer number of buses, bus contention is the dominating factor contributing to the degradation.

A *bumping effect* has been noticed in the probability of acceptance for the prioritized processors. It has been shown that for a fixed number of buses, the probability of acceptance for high priority processors decreases while that of the low priority processors increases with an increase in the percentage of hot memory requests, until a point is reached after which the probability of acceptance for all the processors decrease. With the number of buses remaining constant and for low values of request rates, the low priority processors are affected more with an increase in the request rate.

As the number of buses is increased, keeping the request rate and the hot spot probability constant, the acceptance probability of the processors reach saturation at different values. This give rise to a *knee effect* resulting in the high priority processors reaching saturation earlier than the low priority ones.

References

- [1] T.Y. Feng, "A survey of interconnection networks," *Computer*, vol. 14, pp. 12-27, December 1981.
- [2] T.N. Mudge, J.P. Hayes, and D.C. Winsor, "Multiple-bus architectures," *Computer*, vol. 20, pp. 42-48, June 1987.
- [3] Y.C. Liu and C.J. Jou, "Effective memory bandwidth and processor blocking probability in multiple-bus systems," *IEEE Transactions on Computers*, vol. C-36, no. 6, pp. 761-764, June 1987.
- [4] M. Atiquzzaman and M.M. Banat, "Effect of hot-spots on the performance of crossbar multiprocessor systems," *Parallel Computing*, vol. 19, no. 4, pp. 455-461, April 1993.
- [5] M. Atiquzzaman and M.S. Akhtar, "Effect of non-uniform traffic on the performance of multistage interconnection networks," *9th International Conference on Systems Engineering*, Las Vegas, pp. 31-35, July 14-16, 1993.
- [6] M. Atiquzzaman and M.S. Akhtar, "Effect of hot spots on the performance of multistage interconnection networks," *FRONTIERS 92: The Fourth Symposium on the Frontiers of Massively Parallel Computation*, Virginia, pp. 504-505, October 19-21, 1992.
- [7] M. Atiquzzaman and M.S. Akhtar, "Performance of buffered multistage interconnection networks in non uniform traffic environment," *7th International Parallel Processing Symposium*, California, pp. 762-767, April 13-16, 1993.
- [8] M.A. Sayeed and M. Atiquzzaman, "Performance of multiple-bus multiprocessor under non-uniform memory reference," Tech. Rep. 9/92, La Trobe University, Department of Computer Science, Melbourne,, June 1992.