DESIGN AND ANALYSIS OF A GRACEFULLY-DEGRADING INTERLEAVED MEMORY SYSTEM

by

Kifung C. Cheung, Gurindar S. Sohi, Kewal K. Saluja
University of Wisconsin
Madison, Wisconsin

and

Dhiraj K. Pradhan
University of Massachusetts
Amherst, Massachusetts

Computer Sciences Technical Report #751

February 1988

DESIGN AND ANALYSIS OF A GRACEFULLY-DEGRADING INTERLEAVED MEMORY SYSTEM

Kifung C. Cheung, Gurindar S. Sohi, Kewal K. Saluja
University of Wisconsin
Madison, Wisconsin

and

Dhiraj K. Pradhan
University of Massachusetts
Amherst, Massachusetts

Abstract

A hardware mechanism has been proposed to reconfigure an interleaved memory system. The reconfiguration scheme is such that, at any instant all fault-free memory banks in the memory system can be utilized in an interleaved manner. The design of the hardware that enables the reconfiguration is discussed. The reconfiguration scheme proposed in this paper is analyzed for a number of distinct benchmark programs. It is shown that memory system performance degrades gracefully as the number of faulty banks increase if the memory system uses the proposed reconfiguration scheme.

1. INTRODUCTION

In a computer system that consists of a processing unit (CPU) connected to a memory system, the rate at which the CPU can process information is limited by the rate at which the information can be transmitted between the CPU and the memory. This is the well-known von Neumann bottleneck [1]. Consequently, a decrease in the bandwidth of a memory system will directly affect the performance of the overall computer system.

There are two main approaches to attain a memory system with a high bandwidth. The first involves the use of a high-speed buffer or cache memory (an excellent survey can be found in [2]) and the second involves the use of several memory *banks* or *modules* connected in an interleaved fashion [3, 4]. Though the use of cache memories has become widespread, their utility is limited by their size. While cache memories are very effective for instructions and scalar data items [2, 5], they have not proven to be effective for numeric processing machines that utilize large data structures (such as arrays). For such systems, in order to achieve a high-bandwidth memory system, one is forced to use interleaved banks of memory. Of course, the best effect is achieved by using a cache memory for instructions and scalar data and an interleaved memory for non-cacheable data.

In an interleaved memory system that consists of N independent memory banks (or modules), by associating address latches and data latches with each bank, N different memory accesses can be overlapped. In this C-access method [3] the memory system can accept a stream of memory requests from the processor and service each request, one at a time, thereby increasing the available bandwidth of the memory system to N times the bandwidth of a single bank. A processing system that utilizes a cache memory for instructions and a C-access interleaved memory system for data is shown in Figure 1[1].

The bandwidth of interleaved memories has been studied extensively using analytical and simulation techniques [6-9]. Apart from the referencing behavior of programs, the main factor that influences the bandwidth of interleaved memory banks is the manner in which the addresses are distributed amongst the banks, i.e., the memory organization

[1] Throughout this paper C-access configuration of the memory banks is assumed for the interleaved memory.

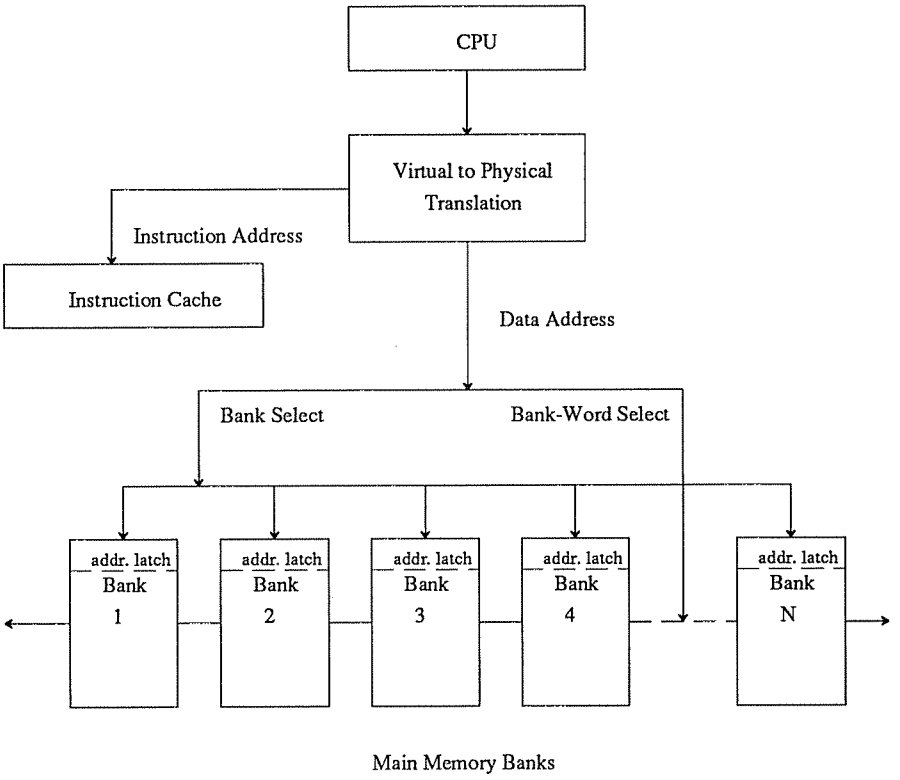

Figure 1: A Processor With an Interleaved Memory System.

[9]. Generally the number of banks, N, that are used to build an interleaved memory is a power of 2, i.e., $N = 2^q$ where q is an integer. In such a system, q bits of the address suffice to select a bank and the remaining bits are used to select a word within a bank. If the q bits are the high-order bits of the address space the scheme is a *high-order interleaving* scheme whereas a *low-order interleaving* scheme results if the low-order q bits are used to select the bank.

We should mention that an interleaving scheme is not restricted to using only a power of 2 number of banks. Interleaving schemes that utilize a prime number of memory banks have been investigated [10] and implemented [11]. However, the utility of such a scheme for high-performance machines is limited because of the complex logic that is needed to determine the appropriate bank/word from a given address.

In a high-order interleaved memory system, consecutive memory addresses in the linear address space lie in the same bank. Therefore, if the memory is referenced sequentially, consecutive memory references access the same bank and no increase in bandwidth is obtained. In a low-order interleaved memory system consecutive addresses lie in different banks. Thus, if the memory is accessed sequentially, consecutive references will access different banks thereby increasing the bandwidth of the memory. Since the memory referencing pattern for most programs is generally sequential (because of sequential instructions and array structures with a constant stride of unity), a low-order interleaved memory system generally has a higher bandwidth than a high-order interleaved memory system.

A low-order interleaving scheme has a major drawback - it is not modular, i.e., a failure in a single bank affects the entire address space [12]. If no precautions are taken to handle such a situation, the bandwidth of the memory and consequently the performance of the processor could be degraded to an intolerable extent.

In this paper, we study the organization of interleaved memories such that faults in the memory system degrade the performance in a graceful manner. We restrict our study to an interleaved memory system that starts out with 2^q memory banks and uses a low-order interleaving scheme. The ideas presented in this paper can easily be extended to other interleaved memory schemes.

Section 2 describes the motivation and design objectives of the memory system. In section 3, a new reconfiguration scheme is presented. Section 4 presents the design of the hardware needed to implement the reconfiguration scheme proposed in section 3. The reconfiguration scheme is evaluated using *trace-driven* simulation in section 5. A final section presents concluding remarks.

2. FAULTS IN INTERLEAVED MEMORIES

Consider a memory that consists of one or more *groups* of interleaved memories. A group consists of 2^r banks (where r is an integer) that are fully interleaved using low-order interleaving. Thus, the banks within a group can be selected using an r-bit bank selection address field. Different groups can have a different number of banks in them.

For example, group G_1 may consist of 4 banks while group G_2 may have only 1 bank. If the total number of banks in the memory system is 2^k where k is an integer, then there is only one group.

A normal interleaved memory consists of a single group of 2^q banks. If each bank contains 2^p words, the total addressable main memory of the system is 2^n (where $n=p+q$) words. Using a standard low-order interleaving scheme, q bits, i.e., $A_{q-1}A_{q-1}\cdots A_0$ of the n-bit address $A_{n-1}A_{n-2}\cdots A_0$ (where A_{n-1} is the most significant bit), are used to select the bank and the remaining p bits, i.e., bits $A_{n-1}A_{n-2}\cdots A_q$ are used to select a word within a bank.

In this paper, we are interested in investigating the situation when one or more banks of memory fail. Therefore, the fault model that we shall use in this paper is that a fault(s) results in the loss of a complete bank(s) of memory. We assume that a mechanism that detects the presence of a faulty bank exists. Such a fault-detection scheme is not the subject matter of this paper. Our main thrust is to evaluate the loss in performance when a fault is reported and how might the memory system be organized so that the resulting degradation in performance is graceful.

Consider what happens when a bank is deleted from a memory system that contains a single group of banks. This is exactly the situation when a bank fails in a standard interleaved memory system. The number of banks in memory is reduced to 2^q-1 and the total addressable physical memory is reduced to $(2^q-1)2^p$ words. However, since 2^q-1 is not a power of 2 the banks no longer form a single group and the system loses its capacity to interleave memory requests. Without interleaving, the bandwidth of the memory system reduces to the bandwidth of a single bank. Such a loss in memory bandwidth can be catastrophic to the performance of a high-speed CPU.

When a fault occurs, program execution must be halted and the address translation mechanism informed about the faulty bank. Correct information would be recovered from a backup store and program execution restarted (or restored form a later point if a recovery scheme is used). Unfortunately, if the memory system is not able to recover a significant portion of its fault-free performance, overall processor performance will suffer. What could we possibly do to salvage some of this lost memory performance?

Two approaches come to mind.

2.1. Use of Redundancy

The first approach involves the use of spare memory banks. After a faulty bank is detected, a spare bank can take its place. An organization of a memory system with one spare bank is shown in Figure 2. For each memory bank i ($0 \leq i \leq N$), the corresponding bank select signal BS_i passes through a switch W_i. Each switch W_i is controlled by a *Bank Status Indicator* (BSI_i) signal which corresponds to the status (normal/faulty) of the bank i. Under normal operation of the memory system the switches connecting the bank select signals to the memory banks are closed at all times. If a memory bank i is found to be faulty then the spare bank, bank N in the Figure 2, is used to substitute for the faulty bank. This can be achieved by opening the switch W_i and closing the switch X_i by the *BSI* control signals.

For the reconfiguration of the memory banks in a system with more than one spare banks, multiplexors can be used in place of the switches. However, as more banks become faulty, the system will eventually run out of spare banks if the spare banks

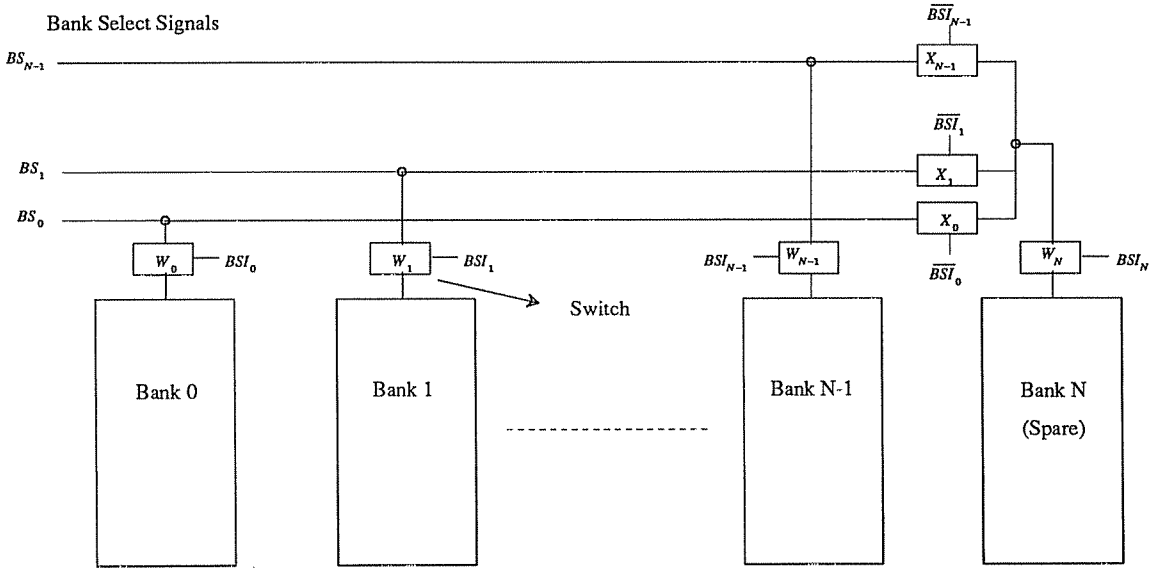

Figure 2: Organization of a Memory System with a Spare Bank.

cannot be replaced. Once all the spare banks have been exhausted, another fault-tolerance scheme must come into play.

2.2. Reconfiguration of Non-faulty Memory Banks

An alternative approach is to reconfigure the remaining non-faulty banks in order to salvage some of the lost performance. Such an approach could also be used if a system has spares but runs out of them eventually. The remaining banks have to be reconfigured so that interleaving is possible. Before proceeding further, let us see how the presence of a faulty bank affects memory system performance. A faulty bank degrades memory performance in two ways: (i) the number of fault-free banks is reduced thereby reducing the available bandwidth and (ii) the amount of available memory is reduced thereby increasing the chances of a page fault.

How might we organize the fault-free banks so that the performance is not degraded to an intolerable extent? A simple solution that could be used to salvage some of the lost bandwidth is to reduce the number of addressable banks to the nearest power of two, i.e., 2^{q-1}, thereby achieving a maximum bandwidth of 2^{q-1} words per memory cycle. While the hardware that allows this translation and the resulting address translation and bank selection mechanism is quite straight forward, $2^{q-1}-1$ banks of fault-free physical memory are not addressable and therefore are unutilized[2]. The resulting memory configuration with 2^{q-1} banks is likely to result in a higher page fault rate than a memory system that uses X $(2^{q}-1 \leq X < 2^{q-1})$ banks. Furthermore, a memory system with X $(2^{q}-1 \leq X < 2^{q-1})$ banks could potentially result in a higher bandwidth than a memory system with 2^{q-1} banks (as we shall see in section 5). Therefore, we must use a reconfiguration scheme that uses as many fault-free banks as possible to salvage the memory bandwidth and, at the same time, minimize the degradation due to the smaller memory size. Also, the hardware needed to implement the reconfiguration scheme should be simple enough so that it does not degrade fault-free memory performance significantly.

[2]Note that even such a scheme requires additional hardware to implement the reconfiguration.

3. THE RECONFIGURATION SCHEME

The proposed scheme reconfigures the remaining banks using a combination of high-order interleaving and low-order interleaving. We distinguish between 3 address spaces: (i) a *virtual* address space that is seen by the program, (ii) a *logical* address space and (iii) a *physical* address space. The difference between the logical and the physical address spaces will become clear in the following discussion.

Addresses in the logical and physical spaces are specified as a bank number and an address within the bank. In the absence of faults, there is a single group of banks and the logical and physical address spaces are the same. When a fault occurs, the state of the physical memory is incorrect. The faulty bank is switched out and the memory reconfigured. We assume that program execution can be restarted (or restored if a recovery scheme is used) from the backup memory. Addresses generated by the user program are still complete virtual addresses; the program does not know about the loss of a memory bank. The logical address space is reduced in a systematic manner. The virtual memory management process is informed about the loss of banks and the new logical configuration of the memory; it views the loss of a memory bank as the loss of a few page frames (equal to one bank) of memory. The virtual to logical translation process makes sure that no information is placed in the unavailable logical space and, for interleaved access to the elements of a page, it places a page entirely within the banks of a single group. As the program executes, pages are brought in from the backup store into the remapped main memory. The logical addresses are translated into physical addresses by the reconfiguration hardware (described in section 4) depending upon the actual banks that have failed.

A logical address specifies a *logical* bank number and a word within the logical bank. For a single faulty bank, there are 2^q-1 non-faulty logical banks and 1 faulty logical bank. The number of faulty logical banks is the same as the number of faulty physical banks. The faulty logical bank is numbered 2^q-1 and the non-faulty logical banks are numbered from 0 through 2^q-2. Non-faulty logical banks are partitioned into sets. Thus, if 2^q-1 logical banks were available, they would be partitioned into q sets. These q sets form 2 subsets; subset $S_0(2^{q-1})$ consisting of a single group of 2^{q-1} logical banks

and subset $S_1(2^{q-1}-1)$ defined recursively as consisting of two subsets $S_0(2^{q-2})$ and $S_1(2^{q-2}-1)$. Therefore, $S_0(2^{q-1})$ has one group $G_0(2^{q-1})$ that has 2^{q-1} logical banks and $S_1(2^{q-1}-1)$ is made up of group $G_{10}(2^{q-2})$ which has 2^{q-2} logical banks and the subset $S_1(2^{q-2}-1)$ which has $2^{q-2}-1$ banks. This recursive partition stops when S_1 has only one logical bank. Another way of looking at this partioning of logical banks into groups is as follows. Write down the number of non-faulty logical banks as a binary number $b_{n-1}b_{n-2}\cdots b_i \cdots b_0$. There is a group with 2^i banks if bit i of the binary number is 1. As we shall see in the following paragraph, the concept of sets is useful in understanding the logical address decoding process. An example of the partitioning of 7 logical memory banks into groups is given in Figure 3.

The 2^k banks within a group $G_i(2^k)$ are organized for low-order interleaving; high-order address bits are used to determine the group. If there is only one group, e.g., in the fault-free case, no group selection needs to be done. The low-order q bits of the address select the logical bank and the high-order p bits of the address select the word within the logical bank. With one fault, the number of groups becomes q with the number of

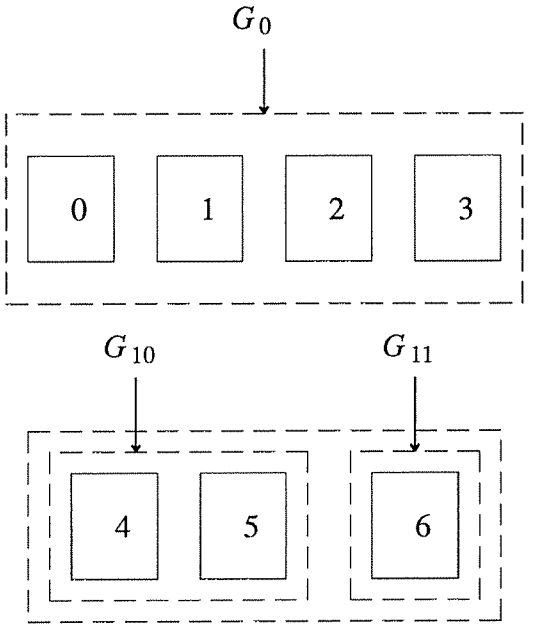

Figure 3: Partitioning Logical Banks into Groups.

logical banks 2^q-1 and q bits suffice to uniquely identify 2^q-1 non-faulty logical banks and one faulty logical bank. An address is decoded as follows: the most significant bit of the address, A_{n-1}, is used to select either subset $S_0(2^{q-1})$, i.e., group $G_0(2^{q-1})$ or subset $S_1(2^{q-1}-1)$. If group $G_0(2^{q-1})$ is selected, then bits A_{q-2} ... A_0 are used to select one of 2^{q-1} logical banks within the group and bits A_{n-2} ... A_{q-1} are used to address the word within the logical bank. If $S_1(2^{q-1}-1)$ is selected, then bit A_{n-2} of the address is used to select either $G_{10}(2^{q-2})$ (with bits A_{q-3} ... A_0 used to select a logical bank within this group) or $S_1(2^{q-2}-1)$ and so on. Note that this group identification scheme resembles the decoding scheme used to decode Huffman-encoded information. Once the group number has been determined from the address, the appropriate p bits are used to select the word within the logical bank.

The logical banks must now be mapped onto the *physical* banks of the memory system. For example, in a system with 8 banks if physical bank 3 is faulty, logical bank 7 (the unavailable logical bank) must be mapped onto physical bank 3 and logical banks 0 through 6 must be mapped onto the remaining physical banks. The logic that decodes the address and generates the appropriate bank select and word select signals is now more complex than a simple decoder. We call this logic the *Address Transliterator (AT)*. Each memory address now passes through the AT before it is forwarded to the memory system (Figure 4). The design of the AT is discussed in detail in the next section. The inputs to the AT are n-bit logical memory address and a 2^q-bit vector, the *Bank Status Indicator (BSI)*, that indicates the status of each bank. The BSI vector consists of a single bit flag for each bank. The flag is set to 0 if the bank is fault-free (available) and 1 if the bank is faulty (unavailable). The BSI vector is updated as faults occur and are detected. The output from the AT is the appropriate physical bank address and the address of the word within the selected bank.

On the occurrence of another fault, the program is stopped and the memory is remapped again. Thus, in the presence of a second fault the smallest logical group that contains only one bank is eliminated. The AT hardware is informed (through the BSI vector) and it responds by modifying its group numbering accordingly. The program is then restarted and continues to execute, albeit with degraded memory performance.

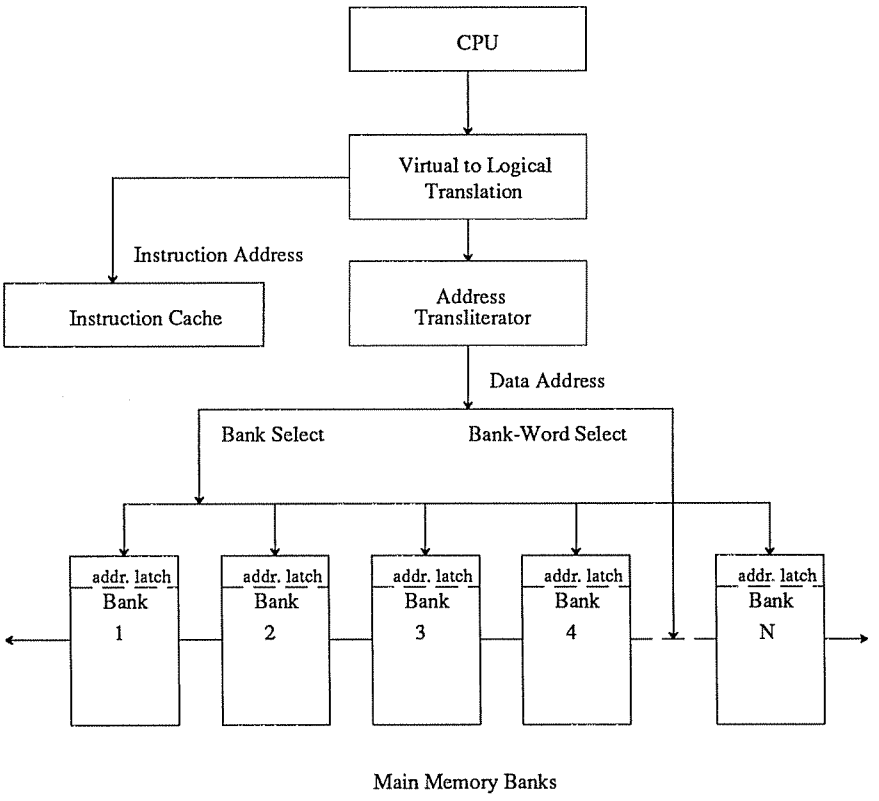

Figure 4: Interleaved Memory System with an Address Transliterator.

4. THE ADDRESS TRANSLITERATOR

The AT hardware consists of three parts as shown in Figure 5: i) a *Bank Fault Tally (BFT)*, ii) a *Bank Select Unit (BSU)* and iii) a *Word Select Unit (WSU)*. The inputs to the AT consist of the n-bit logical address and the BSI vector. The BFT determines the number of faulty banks using the BSI vector. The BSU is responsible for generating the physical bank select signals for a given logical address and the WSU is responsible for supplying the address within the selected bank. Before proceeding with the details of the AT, we would like to mention that the partitioning of the AT into distinct components is done solely to facilitate the understanding of its operation. For implementations of the AT, the above demarcation is not necessary. However, the overall operation of the logic will remain unchanged. Depending upon the technology used to realize the AT, such a demarcation may or may not offer optimum performance. In fact, it may be possible to

implement portions of the AT within the address decoders of the memory banks.

4.1. The Bank Fault Tally

The Bank Fault Tally (BFT) determines the number of faulty banks in the memory system from the information given in the 2^q bit BSI vector. The BFT provides a 2^q+1 bit output in the form of decoded fault count indicators $F = F_{2^q},...,F_0$, where $F_i = 1$ if and only if there are exactly i faulty banks in the memory system. Thus, for a fault-free memory system, F_0 is 1 and all the remaining fault count indicators are 0. Note that the number of faulty banks in the memory system is given by the number of 1's in the BSI vector. Also note that the BFT is a combinational circuit.

It is possible to have the outputs of the BFT in an encoded form. This will reduce the number of interconnecting lines between BFT and other subunits of the AT, but then the coded information may need to be decoded by other subunits. If, on the other hand, the outputs are produced in decoded form, as described above, then the BFT can be realized as a tally circuit [13]. Throughout this paper, we use the decoded form of the outputs for ease of understanding.

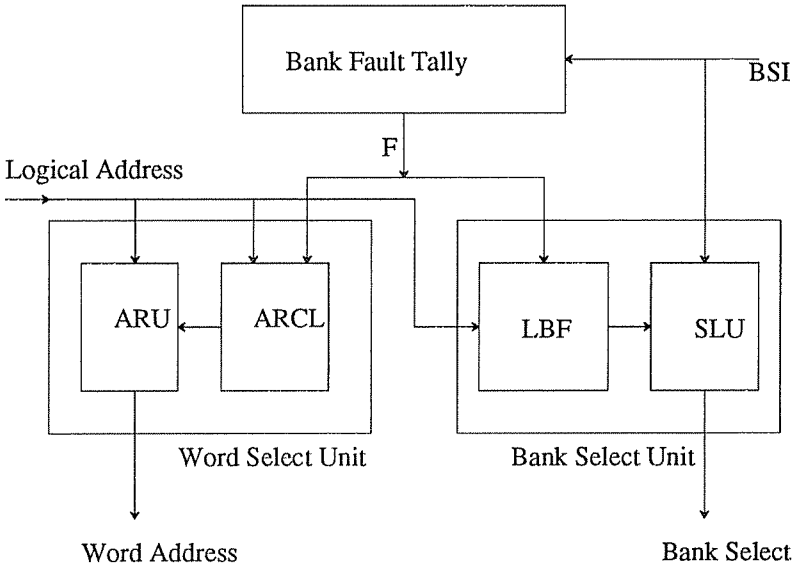

Figure 5: The Address Transliterator.

4.2. Bank Select Unit

The bank select unit consists of two subunits, the *Logical Bank Finder (LBF)* and the *Switching Logic Unit (SLU).*

4.2.1. Logical Bank Finder (LBF)

The LBF determines the logical bank number of the bank being accessed for a given logical address. Recall that in the presence of a single faulty bank, the faulty bank is represented by a string of 1's in the high-order q bits of the logical address. This result can be generalized for multiple faults. Table 1 shows the largest valid logical address assuming that the memory system consists of 8 banks and there are exactly k, $0 \leq k < 8$, faulty banks.

Table 1: Largest Valid Logical Address with k Faulty Banks

k	A_{n-1}	A_{n-2}	A_{n-3}	A_{n-4}	$\cdots$	A_0
7	0	0	0	1	$\cdots$	1
6	0	0	1	1	$\cdots$	1
5	0	1	0	1	$\cdots$	1
4	0	1	1	1	$\cdots$	1
3	1	0	0	1	$\cdots$	1
2	1	0	1	1	$\cdots$	1
1	1	1	0	1	$\cdots$	1
0	1	1	1	1	$\cdots$	1

In the general case, given a logical address, the LBF determines the logical bank number of the bank being addressed based on the following inputs:

i) $2q$ bits (i.e., high-order q bits and low-order q bits) of the address and

ii) 2^q+1 bits from the BFT indicating the fault count.

The q bits of the logical address used to select the logical bank depend upon the number of faults. The LBF has 2^q outputs, denoted as $B_{2^q-1}, \ldots, B_0$, that represent the decoded form of the logical bank numbers.

The operation of the LBF is best illustrated with an example. Table 2 gives the 3-bit logical bank number $L_2L_1L_0$ of the addressed bank as a function of 6 bits,

($A_{n-1}A_{n-2}A_{n-3}A_2A_1A_0$), and the number of faults, k, for the case of 8 (q=3) banks. In the fault-free case, the low-order bits $A_2A_1A_0$ are used to select the logical bank. In the presence of a single fault, the 7 fault-free logical banks are divided into three groups of 4, 2 and 1 banks respectively. If bit A_{n-1} is 0, the group of 4 logical banks is selected and bits A_1A_0 are use to select the logical bank within the group (see bold row in Table 2). If bit A_{n-1} is 1, then the set of 3 logical banks is selected and bit A_{n-2} is used to distinguish between the 2 groups of banks within the set. Other entries in the table can be interpreted in a similar fashion.

The LBF is a combinational circuit that implements a set of independent Boolean equations of the inputs A_i and F_i. For reasons of brevity, we do not present the exact Boolean equations for the LBF in this paper. The interested reader is referred to [14].

4.2.2. Switching Logic Unit (SLU)

The SLU maps a logical bank number obtained from the LBF onto a physical bank number, i.e., it provides the select signals for the physical memory banks. A faulty bank is never selected. If there are no faulty banks in the memory system, the physical bank

Table 2: Logical Bank Number Selection in Case of k Faulty Banks.

k	A_{n-1}	A_{n-2}	A_{n-3}	L_2	L_1	L_0
7	X	X	X	0	0	0
6	0	0	X	0	0	A_0
5	0	0	X	0	0	A_0
	0	1	0	0	1	0
4	0	X	X	0	A_1	A_0
3	0	X	X	0	A_1	A_0
	1	0	0	1	0	0
2	0	X	X	0	A_1	A_0
	1	X	X	1	0	A_0
1	**0**	**X**	**X**	**0**	A_1	A_0
	1	0	X	1	0	A_0
	1	1	0	1	1	0
0	X	X	X	A_2	A_1	A_0

number of every bank is same as the logical bank number as shown in Figure 6a. However, in the presence of faulty banks in the memory system, the logical bank numbers are remapped by the SLU onto physical bank numbers as follows. If there is one faulty bank in the memory system, say physical bank number m, then a logical bank number $m+i$, $i \geq 0$, is remapped to the physical bank number $m+i+1$. Thus, if the faulty bank is bank number 1, physical bank number 2 will be selected for logical bank number 1 as shown in Figure 6b. For multiple faults, the remapping mechanism is extended in a natural manner and always selects the next available fault-free physical bank. The inputs to the SLU consist of 2^q-bit BSI vector, containing the location of faulty banks in the memory system, and the 2^q outputs of the LBF. The outputs of the SLU are 2^q physical bank select signals, $(BS_0, \ldots, BS_{2^q-1})$, where BS_i is used to select the physical bank number i. This logic can, therefore, be realized using switches or multiplexors [14].

4.3. Word Select Unit

In a fault-free memory system, the high order p bits of the address determine the address of the word within a bank. However, if faulty banks exist in the memory system, the address of the word within a bank depends on the number of faulty banks in the system as well as the logical bank number of the selected bank. For example, for memory system with 8 banks ($q = 3$) and with one faulty bank, if a memory reference selects logical bank 5 then the word address within logical bank 5 is given by bits $A_{n-3}, A_{n-4}, \ldots, A_1$. In general, if for a given address some high order j bits are used to determine the logical bank number then the word address within the bank is given by $(A_{n-j}, \ldots, A_{q-j+1})$. The function of the WSU, therefore, is to extract the appropriate p bits that represent the address of the word within the selected bank from the n-bit logical address. The WSU accomplishes this by using 2 subunits, the *Address Reformulation Control Logic (ARCL)* and the *Address Reformulation Unit (ARU)*.

4.3.1. The Address Reformulation Control Logic (ARCL)

This subunit determines the p bits that represent the word address. The p bits are then extracted in the ARU by using shift-left by i bit or SL_i operations on the input address and retaining the high-order p bits. The ARCL determines the value of i, i.e.,

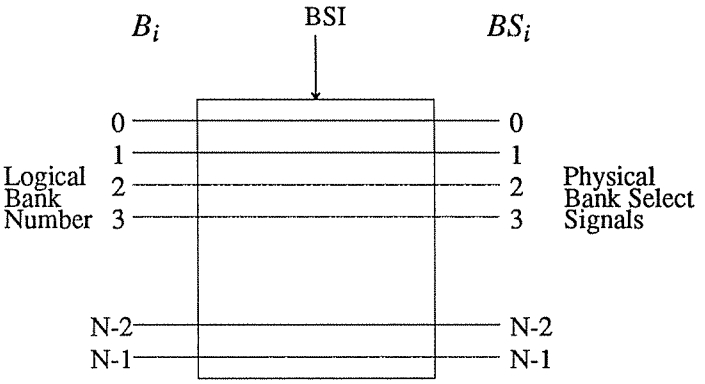

(a): Logical to Physical Bank Mapping with no faults.

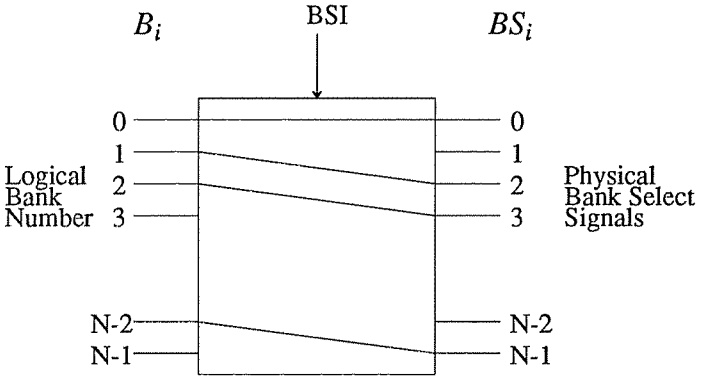

(b): Logical to Physical Bank Mapping with Physical Bank 1 faulty.

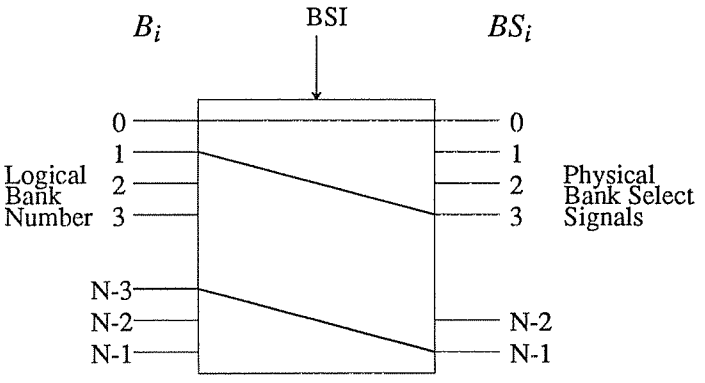

(c): Logical to Physical Bank Mapping with Physical Banks 1 and 2 faulty.

Figure 6: Operation of the Switching Logic Unit.

the SL_i signals from the information provided in the high-order q bits of the logical address and the number of faults obtained from the BFT. The SL_i signals are then provided to the ARU.

Rather than present the detailed Boolean equations, we again illustrate the operation of the ARCL with the help of an example. Table 3 presents the SL_i signals for q=3 and k faults. In the fault-free case, the high-order p bits themselves represent the word address, i.e., $SL_0 = 1$. In the case of 2 faulty banks, we have two groups of 4 and 2 banks respectively. If the group of 4 banks is addressed, i.e., $A_{n-1} = 0$ (see the bold row in Table 3), the logical address needs to be shifted left by 1 ($SL_1 = 1$) and the high-order p bits retained for the word address. If the group of 2 banks is selected ($A_{n-1} = 1$), then the logical address needs to be shifted left by 2 ($SL_2 = 1$) and the high-order p bits extracted.

Table 3: Control Signals of the ARCL in Case of k Faulty Banks.

k	A_{n-1}	A_{n-2}	A_{n-3}	SL_3	SL_2	SL_1	SL_0
7	X	X	X	1	0	0	0
6	0	0	X	0	1	0	0
5	0	0	X	0	1	0	0
	0	1	0	1	0	0	0
4	0	X	X	0	0	1	0
3	0	X	X	0	0	1	0
	1	0	0	1	0	0	0
2	**0**	**X**	**X**	**0**	**0**	**1**	**0**
	1	X	X	0	1	0	0
1	0	X	X	0	0	1	0
	1	0	X	0	1	0	0
	1	1	0	1	0	0	0
0	X	X	X	0	0	0	1

4.3.2. The Address Reformulation Unit (ARU)

The ARU accepts the n-bit address and the shift control signals SL_i as inputs and provides p-bit word address. The ARU is, therefore, a simple shifter and can be realized using multiplexors.

4.4. An Example

Through the following example we illustrate the operation of the complete AT. Let us consider a memory system consisting of 8 memory banks (q=3) with a 16 bit address (n=16). If physical banks 1 and 2 are faulty, the BSI vector will be 00000110, where the most significant bit of the BSI vector indicates the status of the physical bank 7. If the input address is 1001111111111100, from Table 2 we see that logical bank 4 of group G_{10} is selected. Therefore, 3 bits of the address ($A_{n-1}A_{n-2}A_0$) are used to identify the logical bank and the remaining bits of the address, ($A_{n-3}A_{n-4}...A_2A_1$), are used as the word address within the bank. The remapping of the logical bank numbers, i.e., outputs of the LBF, to the physical bank select signals by the SLU is shown in Figure 6c. As the address given above generates the bank address for the logical bank 4, the physical bank 6 is selected. The outputs of the subunits of the AT are summarized below in Table 4.

Table 4: Outputs of the Subunits of the AT; Input Address = 1001111111111100 and k = 2.

Unit	Inputs	Outputs
Bank Fault Tally (BFT)	BSI=00000110	$F_2 = 1$, $F_i = 0$ for $0 \le i \le 7$ and $i \neq 2$.
Logical Bank Finder (LBF)	F_i 's and $A_{n-1}A_{n-2}A_{n-3}$=100 $A_2A_1A_0$ =100	$B_4 = 1$, $B_i = 0$ for $0 \le i \le 7$ and $i \neq 4$.
Switching Logic Unit (SLU)	B_is and BSI vector	$BS_6 = 1$, $BS_i = 0$ for $0 \le i \le 7$ and $i \neq 6$.
Address Reformulation Control Logic (ARCL)	F_i 's and $A_{n-1}A_{n-2}A_{n-3}$=100	$SL_2 = 1$, SL 's = 0 for $0 \le i \le 3$ and $i \neq 2$
Address Reformulation Unit (ARU)	16 bit address and SL_i's	0111111111110

4.5. Logic Delays in the AT

For a conventional low-order interleaved memory system, shown in Figure 1, the address bits are transferred on two paths between the processor and the memory. The two paths are the bank select path and the word select path. The logic present in either path is a simple decoder. By using an AT, additional logic is inserted in both these paths.

Since the delays introduced by the extra logic will be of a different nature for different technologies and different implementations of the AT, we shall not attempt to quantify the delays in general. Rather, we give the reader a feel for the additional delays introduced by the AT and present the results for a conservative CMOS design.

The BFT does not contribute to the delay through the AT since its output does not change between faults. The longest path within the BSU is from the address inputs to the outputs of the LBF plus the delay through the SLU. The complexity of the LBF is such that it can be realized by a two level logic circuit (gates or a PLA); the SLU can be realized as a simple switch (or multiplexor). Each of these units will, therefore, contribute a small, fixed delay. Similarly, in the WSU, the critical path is from the address inputs to the outputs of the ARCL plus the delay through the ARU. The ARCL logic is simple enough that it can be realized as a two level logic circuit and, as commented in section 4.3.2, the ARU can be realized using multiplexors. The operations of the BSU and the WSU are carried out in parallel and, therefore, the critical path for the AT is the longer of the critical paths of the two units. Thus the delay through the AT is equivalent to the delay through a few levels of logic.

In order to get a better feel of the delays in the AT, we implemented the AT logic in CMOS VLSI using Magic [15]. The design was very conservative. Each unit of the AT was designed separately as described in this paper. The AT was designed for a memory system consisting of 16 memory banks and a 32-bit logical address. The details of complete design can be found in [14]. A timing simulation using the Crystal simulator[16] indicated a delay of 44ns through the AT. We believe that the delay can be reduced with more sophisticated design; relative delays within the WSU and BSU and the methods to reduce these delays are also discussed in [14]. Even for our conservative design of the AT, for a given processor and memory technology, the delay through the AT can be kept within a single CPU clock cycle [14]. Therefore, we believe that the AT logic will not degrade fault-free memory latency to a significant extent.

5. PERFORMANCE EVALUATION

We evaluated the performance of the proposed memory reconfiguration scheme using a *trace-driven* simulation analysis. A trace of instruction and data references was obtained for several benchmark programs. Data from the trace files was then fed into a program that simulated the memory system. Since the simulator is driven by actual memory reference traces which have no timing information and the simulator assumes that consecutive memory references in the trace file occur in consecutive CPU cycles.

The simulator written by us takes into account the memory structure and the virtual memory management process. The operational structure of the simulator is shown in Figure 7. The simulation model consists of a pipelined processor capable of issuing a memory request at each CPU cycle. The memory references are divided into instruction references and data references. For each memory reference a word is transferred between the memory system and the processor. We assume an instruction cache with a cycle time identical to the CPU cycle time to service instruction references. Thus, only data references go to the interleaved memory. The use of the bus alternates between the instruction

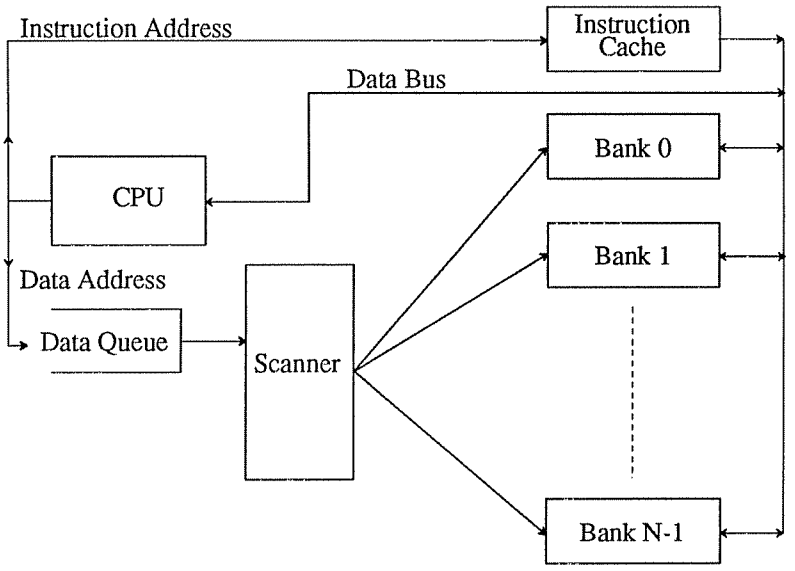

Figure 7: Structure of the Simulation Model.

cycles and the data cycles. During an instruction cycle, instructions are fetched from the cache and the data requests are buffered in a queue. We assume that all instruction requests are satisfied by the instruction cache [2, 5]. During a data cycle, data requests in the data queue are allowed to access memory if no bank conflict occurs. If an access to a busy bank is detected, subsequent requests are suspended till the next available memory cycle.

A program is allocated a fixed number of data pages (maximum of 32) for its use. A least recently used (LRU) replacement policy is employed to replace a page when a page fault occurs and no free page frame is available. The pages are loaded into memory on demand. The page size is 2K bytes. Initially, we assume that only one data page is present in memory. In case of bank failures, fewer data pages are allocated to the program. The reduction in the number of pages is proportional to the number of faulty banks. The pages are distributed amongst the groups of the reconfigured memory in proportion to the number of banks in the group. A page lies completely within a group of banks. Pages are first loaded into the group with the largest number of banks and are then loaded into the groups with fewer banks. For example, if there are 2 groups consisting of 8 and 4 banks respectively, a process will place 67% of its data pages in the group of 8 banks first and the remaining pages in the group of 4 banks next. The time to process a page fault is 2000 memory cycles [17].

Recall that faults degrade memory system performance in two ways: (i) the available memory bandwidth is reduced and (ii) a reduction in the available physical memory increases the probability of a page fault. We believe that a performance metric must take into account both factors of memory performance. Therefore, we combine the effects of reduced bandwidth and increase in page faults into a single metric, T, similar to the metric used by Smith [18]. The performance metric T is defined as:

T = total data trace length + time to process a page fault $\times$ number of data page faults

where the data trace length is the number of data references divided by the data bandwidth. T is an indicator of the memory access time for the trace when both bandwidth and page faults are taken into account.

We realize, however, that in some situations the bandwidth may be the more important metric and the number of page faults may be of secondary importance. In other cases, the page faults may be of major concern. Therefore, along with the metric T, we also present the bandwidth and the number of page faults for each one of our experiments.

5.1. Experiments and Results

To evaluate the performance of the reconfigurable memory system, we used the following benchmark programs: (i) *nroff*, a text formatting program, (ii) *compact*, a program for file compaction using adaptive Huffman encoding, (iii) *boyer*, a theorem proving program [19], (iv) *tak*, an execution of the Takeuchi function [19], (v) *spice*, a circuit simulation program, (vi) *mpla*, a PLA generation program for the Magic layout editor [15], (vii) *cripta*, an encryption program written in Lisp, and (viii) *csh*, a command interpreter for the UNIX operating system.

The number of instruction and data references traced for the above programs are given in Table 5. Total number of data pages in each trace are also given in this table.

Table 5: Statistics for the Benchmark Programs

Trace	Trace Records		Data Pages Touched
	Instruction	Data	
nroff	284966	175488	52
compact	234110	205468	22
boyer	217147	229871	216
tak	236628	250384	151
spice	258563	250996	55
mpla	255708	173266	94
cripta	147663	150386	221
csh	220367	211592	72

The traces for each of these programs were fed to the trace driven simulator and the bandwidth, the page faults and the value of the performance metric computed. To account for the presence of faulty banks, each trace is simulated for reduced number of addressable banks and the number of pages allocated to each program reduced in

proportion to the amount of memory lost. The T metric is normalized with respect to the fault-free case. The results of our experiments are presented in Tables 6-13. The bandwidth in these tables is defined as the number of busy banks per memory cycle.

Table 6: Result for *nroff*

N	Data				T
	Bandwidth	Page Faults	Page Fault Rate (%)	Pages	
16	3.417	53	0.030	32	1.00
15	3.078	53	0.030	30	1.04
14	3.079	53	0.030	28	1.04
13	3.083	63	0.036	26	1.16
12	3.098	122	0.070	24	1.91
11	3.021	131	0.075	22	2.03
10	3.009	165	0.094	20	2.47
9	3.025	204	0.116	18	2.96
8	3.011	246	0.140	16	3.50

Table 7: Result for *compact*

N	Data				T
	Bandwidth	Page Faults	Page Fault Rate (%)	Pages	
16	2.895	22	0.011	16	1.00
15	2.450	23	0.011	15	1.13
14	2.345	23	0.011	14	1.16
13	2.421	44	0.021	13	1.50
12	2.396	112	0.055	12	2.69
11	2.583	326	0.157	11	6.36
10	2.627	827	0.402	10	15.07
9	2.568	1908	0.929	9	33.88
8	2.544	3760	1.830	8	66.10

Table 8: Result for *boyer*

N	Data				T
	Bandwidth	Page Faults	Page Fault Rate (%)	Pages	
16	3.283	1465	0.637	32	1.00
15	2.809	1696	0.738	30	1.16
14	3.110	1962	0.854	28	1.33
13	3.043	2328	1.013	26	1.58
12	3.143	2754	1.198	24	1.86
11	2.957	3209	1.396	22	2.17
10	2.998	4454	1.938	20	2.99
9	2.223	6696	2.913	18	4.50
8	2.998	8895	3.870	16	5.96

Table 9: Result for *tak*

N	Data				T
	Bandwidth	Page Faults	Page Fault Rate (%)	Pages	
16	4.920	319	0.127	32	1.00
15	4.600	335	0.134	30	1.05
14	2.900	365	0.146	28	1.18
13	3.322	393	0.157	26	1.25
12	4.343	449	0.179	24	1.39
11	2.114	513	0.205	22	1.66
10	4.083	585	0.233	20	1.79
9	4.193	678	0.271	18	2.06
8	3.986	1286	0.514	16	3.82

Table 10: Result for *spice*

N	Data				T
	Bandwidth	Page Faults	Page Fault Rate (%)	Pages	
16	3.577	80	0.032	32	1.00
15	3.017	116	0.046	30	1.37
14	3.045	160	0.064	28	1.75
13	2.981	189	0.075	26	2.01
12	3.109	231	0.092	24	2.36
11	2.813	247	0.098	22	2.53
10	2.948	274	0.109	20	2.75
9	2.894	351	0.140	18	3.42
8	2.894	407	0.162	16	3.91

Table 11: Result for *mpla*

N	Data				T
	Bandwidth	Page Faults	Page Fault Rate (%)	Pages	
16	2.119	129	0.074	16	1.00
15	2.094	132	0.076	15	1.02
14	2.095	136	0.078	14	1.04
13	2.090	374	0.216	13	2.45
12	2.072	2200	1.270	12	13.20
11	2.029	2543	1.468	11	15.22
10	2.027	2670	1.541	10	15.97
9	2.059	2712	1.565	9	16.21
8	2.046	2738	1.580	8	16.35

Table 12: Result for *cripta*

N	Data				T
	Bandwidth	Page Faults	Page Fault Rate (%)	Pages	
16	2.569	564	0.375	32	1.00
15	2.466	629	0.418	30	1.11
14	2.477	808	0.537	28	1.41
13	2.399	1044	0.694	26	1.81
12	2.417	1139	0.757	24	1.97
11	2.190	1261	0.839	22	2.18
10	2.366	1490	0.991	20	2.57
9	2.201	1684	1.120	18	2.90
8	2.321	1869	1.243	16	3.20

Table 13: Result for *csh*

N	Data				T
	Bandwidth	Page Faults	Page Fault Rate (%)	Pages	
16	3.097	443	0.209	16	1.00
15	2.483	507	0.240	15	1.15
14	2.645	562	0.266	14	1.26
13	2.426	627	0.296	13	1.41
12	2.759	740	0.350	12	1.63
11	2.071	909	0.430	11	2.01
10	2.505	1143	0.540	10	2.48
9	2.434	1528	0.722	9	3.29
8	2.697	5029	2.377	8	10.62

From the tables, one observation is quite obvious - a decrease in the available physical memory increases the number of page faults. We also observe that, in most cases, the bandwidth when α banks ($15 \leq \alpha \leq 9$) are used is better than the bandwidth that could be achieved with 8 banks. If bandwidth alone is the major performance-determining factor, then the reconfiguration scheme could be used to reconfigure part or all of the remaining fault-free banks (for example, we could choose to reconfigure only 12 banks in the case of 1, 2, 3 or 4 faults). If the memory capacity is a limiting factor, then page faults play an important role in the overall memory access time. Our reconfiguration scheme allows the use of all the fault-free memory thereby minimizing the number of page faults. The resulting degradation in memory system performance (as measured by T) is quite graceful.

Based on the experimental results we can conclude that the proposed reconfiguration scheme allows for the graceful degradation of interleaved memory systems. In situations where the memory capacity is unimportant, the reconfiguration scheme is able to reconfigure the fault-free banks so that the resulting memory configuration has a better bandwidth than a memory configuration with the next lower power-of-2 number of banks. In situations where the memory system capacity is a limiting factor, the reconfiguration scheme is able to reconfigure the memory to minimize the number of page faults and, at the same time, recover part of the lost memory bandwidth.

6. CONCLUSIONS

In this paper, we have presented the design of an interleaved memory system whose performance degrades gracefully in the presence of faulty banks. We discussed the details of such a design and evaluated its performance using a trace-driven simulation. Our simulation results show that the performance of an interleaved memory system that employs the design proposed in this paper does indeed degrade gracefully in the presence of faults. Furthermore, the address translation mechanism needed for graceful degradation does not increase the memory latency significantly.

7. ACKNOWLEDGMENTS

The authors are thankful to Professor C. R. Kime his helpful comments and suggestions. This work was supported in part by National Science Foundation Grants DCR-8509194, CCR-8706722 and Air Force Grant AFOSR 84-0052.

References

[1] J. Backus, "Can Programming Be Liberated from the von Neumann Style? A Functional Style and Its Algebra of Programs," *Communications of the ACM*, vol. 21, pp. 613-641, August 1978.

[2] A. J. Smith, "Cache Memories," *ACM Computing Surveys*, vol. 14, pp. 473-530, Sept. 1982.

[3] K. Hwang and F. A. Briggs, *Computer Architecture and Parallel Processing*. New York: McGraw-Hill, 1984.

[4] P. M. Kogge, *The Architecture of Pipelined Computers*. New York: McGraw-Hill, 1981.

[5] J. E. Smith and J. R. Goodman, "A Study of Instruction Cache Organizations and Replacement Policies," *Proc. 10th Annual Symposium on Computer Architecture*, pp. 117-123, June 1983.

[6] G. Burnett and E. G. Coffman, "A Study of Interleaved Memory Systems," *Proc. AFIPS 1970 Spring Joint Computer Conference*, pp. 467-474, 1970.

[7] F. W. Terman, "A Study of Interleaved Memory Systems by Trace Driven Simulation," *Proc. Symposium on the Simulation of Computer Systems*, pp. 3-9, 1976.

[8] D. P. Bhandarkar, "Analysis of Memory Interference in Multiprocessors," *IEEE Trans. on Computers*, vol. C-24, pp. 897-908, September 1975.

[9] B. R. Rau, "Program Behavior and the Performance of Interleaved Memories," *IEEE Trans. on Computers*, vol. C-28, pp. 191-199, March 1979.

[10] D. H. Lawrie and C. R. Vora, "The Prime Memory System for Array Access," *IEEE Trans. on Computers*, vol. C-31, pp. 435-442, May 1982.

[11] D. J. Kuck and R. A. Stokes, "The Burroughs Scientific Processor (BSP)," *IEEE Trans. on Computers*, vol. C-31, pp. 363-376, May 1982.

[12] D. K. Pradhan, *Fault Tolerant Computing: Theory and Techniques*. Englewood Cliffs, New Jersey: Prentice Hall, 1986.

[13] C. Mead and L. Conway, *Introduction to VLSI Systems*. Reading, MA: Addison-Wesley, 1980.

[14] K. C. Cheung, "Organization and Analysis of Interleaved Memory Systems," M. S. Thesis, Dept. of Electrical Electrical and Computer Engineering, University of Wisconsin-Madison, Madison, WI, 1987.

[15] J. K. Ousterhout, G. T. Hamachi, R. N. Mayo, W. S. Scott, and G. S. Taylor, "The Magic VLSI Layout System," *IEEE Design & Test of Computers*, February, 1985.

[16] J. K. Ousterhout, "Crystal: A Timing Analyzer for nMOS VLSI Circuits," in *Third Caltech Conference on Very Large Scale Integration*. Computer Science Press, pp. pp. 57-70, 1983.

[17] M. Malkawi and J. H. Patel, "Performance Measurement of Paging Behavior in Multiprogrammed Systems," *Proc. 13th Annual Symposium on Computer Architecture*, pp. 111-118, June 1986.

[18] A. J. Smith, "A Modified Working Set Paging Algorithm," *IEEE Trans. on Computers*, vol. C-25, pp. 907-914, September 1976.

[19] R. P. Gabriel, *Performance and Evaluation of Lisp Systems*. Cambridge, MA: The MIT Press, 1985.