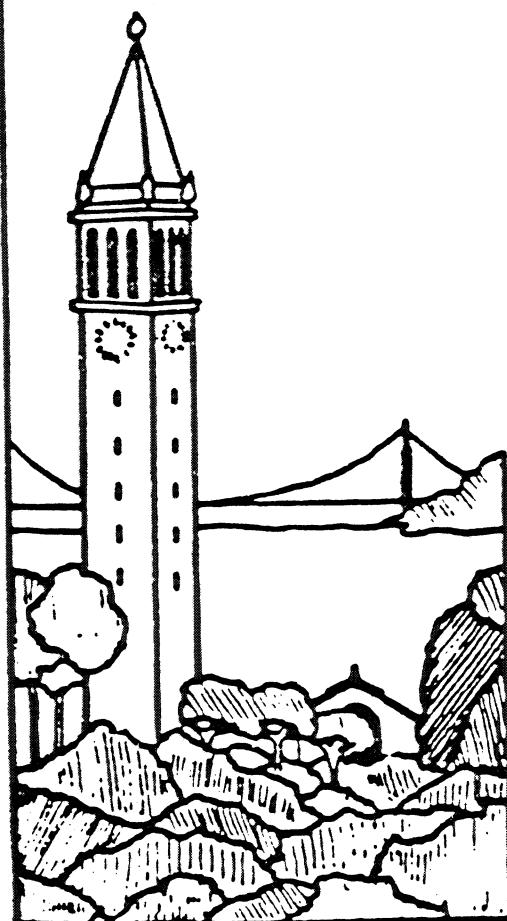


Performance Evaluation of Cache Prefetch Implementation

*John Tse
Alan Jay Smith*



Report No. UCB//CSD-95-877

June 1995

**Computer Science Division (EECS)
University of California
Berkeley, California 94720**

Abstract

Prefetching into CPU caches has long been known to be effective in reducing the cache miss ratio, but implementations of prefetching have been unsuccessful in improving CPU performance. The reasons for this are that prefetches interfere with normal cache operation by making cache address and data ports busy, the memory bus busy and the memory banks busy, and by not necessarily being complete by the time that the prefetched data is actually referenced. In this paper, we present the results of a very detailed cycle by cycle trace driven simulation of a uniprocessor memory system, in which we vary several relevant parameters in order to determine when and if prefetching is useful. We find that in order for prefetching to actually improve performance, the address array needs to be double ported, and the data array needs to either be double ported or fully buffered. It is also very helpful for the bus to be reasonably wide, bus transactions to be split and main memory to be interleaved. Under the best circumstances, i.e. with a significant investment in extra hardware, prefetching can significantly improve performance.

1. Introduction

It is well known that delays in accessing CPU memory are one of the major factors in limiting CPU performance [Smit82]. Cache memories are the principal technique used to improve CPU memory system performance, but cache misses continue to degrade performance. Cache studies have considered various factors, such as line (block) size, associativity, cache size, prefetching algorithms, etc. In this paper, we study the effectiveness of prefetching, where necessary taking into account its interaction with other cache design parameters.

Prefetching has long been known to significantly decrease the miss ratio of the CPU cache. This was shown in [Smit78], with additional results, for varying algorithms and workloads, presented in [Smit82,85]. These studies, and most of the others that appear in the literature (see below), however, have neglected timing effects; i.e. even though the miss ratio has decreased, does the machine get faster? In this paper, we present the results of a very detailed cycle by cycle trace driven simulation of a memory system, in which we vary all relevant parameters in order to determine when and if prefetching is useful. Our study differs from others that have used similar timing simulations in the detail of our model, in the fact that we are concentrating on uniprocessor design, and in the range of architectural parameters studied.

1.1. Previous Research, and Prefetch Algorithms

CPU cache prefetching involves fetching a block from main memory into the CPU cache when it has not been referenced, in the hope that it will be referenced soon. Prefetching algorithms have been concerned with two issues, which block to prefetch, and when to prefetch. The simplest choice of block to prefetch is the next sequential block after the block most recently referenced. Conceptually, this makes sense because instructions are fetched sequentially (except for branches), and data is often referenced sequentially (when in arrays) or at least locally (when the compiler allocates related variables in contiguous locations).

One simple prefetch algorithm is called *always prefetch* [Smit82]. With this algorithm, every time there is a reference to block i , the cache is examined for block $i+1$ (i.e. the next sequential block, in terms of ascending memory addresses); if block $i+1$ is absent from the cache, it is prefetched. A variation, which requires fewer prefetches and prefetch lookups (i.e. look in the cache to see if the block is there) is called *prefetch on misses*, which prefetches the next sequential cache block if and only if the access to the current

*The authors' research is or has been supported in part by the National Science Foundation under grants MIP-9116578 and CCR-9117028, by NASA under Grant NCC 2-550, by the State of California under the MICRO program, and by Intel Corporation, Apple Computer Corporation, Sun Microsystems, Digital Equipment Corporation, Philips Laboratories/Signetics, International Business Machines Corporation, and Mitsubishi Electric Research Laboratories.

cache block is a miss. A more complicated scheme known as *tagged prefetch* [Smit82, Gind77] keep the number of prefetch lookups low while issuing more prefetches than prefetch on misses. In this case, each cache block has a single bit, called the tag, which is set to zero whenever the block does not reside in the cache. When a block is referenced by the processor, its tag will be set to one. A block brought into the cache by a prefetch, however, retains its tag of zero. Whenever a tag changes from zero to one, a prefetch is initiated for the next sequential cache block. This is similar to always prefetching, but it avoids repeated cache lookups, and also does not prefetch a line which was prefetched and then replaced without having been referenced.

More sophisticated schemes like *threaded prefetching* [Kim93] and *bi-directional prefetching* [Varm92] have also been proposed. In threaded prefetching, cache block i has associated with it a list of pointers known as threads. Each thread points to a cache block which is most likely to be referenced after block i has just been accessed. Suppose that the processor is accessing cache block i in cycle T . If block j is referenced in cycle $T+1$, a new thread which contains the address of block j will be attached to block i in cycle $T+1$. The new thread can be stored together with block i in the instruction or data cache, or it can be stored in a separate cache. As soon as block i is reaccessed by the processor, all threads associated with block i will trigger the prefetching of block j and other cache blocks the threads point to. Bi-directional prefetching attempts to capture the forward and backward accessing behavior found in data caches. If the current and the previous sequential cache blocks are found in the cache, the next sequential cache block will be prefetched; otherwise, if the current and the next sequential cache blocks are found, the previous sequential cache block will be prefetched.

When a prefetching strategy detects a potential miss, it will issue a prefetch for a fixed number of cache blocks. Since the prefetching efficiency varies during the execution of a program, it may be advantageous to prefetch a variable number of cache blocks. Such an adaptive sequential prefetching scheme for shared memory multiprocessors is proposed in [Dahl93]. Counters are kept with each cache block which record the number of times prefetches on this block have been serviced by the main memory. By inspecting these counters, the idea is to be able to measure the effectiveness of prefetching dynamically and prefetch an appropriate number of cache blocks accordingly.

Cache prefetching strategies can also be used in special cache organizations like vector caches. Data prefetching strategies for vector cache memories are examined in [Fu91]. A stride-prefetch strategy is proposed which takes advantage of the vector stride information specified in a vector instruction that causes a cache miss. If element i in the vector data causes the miss, prefetches for blocks $i, i+stride, i+2*stride, \dots, i+p*stride$ will be issued, where p is the number of sequential blocks to be prefetched. Not unexpectedly, results show that stride-prefetch strategy performs better than the always prefetching strategy and no prefetching.

Two level cache memory systems have been shown to be an effective means to enhance system performance [Baer88] [Shor88]. Methods to evaluate the performance of cache prefetching in the second level caches are studied in [Smit93]. Results show that miss ratios alone are insufficient to evaluate the performance of cache prefetching in higher level caches. By using a detailed timing-based trace-driven simulation model, it is shown that hit-rate-only analysis may be extremely optimistic in predicting the benefits of cache prefetching. This paper uses a similar approach, but explores a larger design space in more detail.

One of the disadvantages of cache prefetching is the unavoidable increase in memory traffic because of prefetches which are never referenced. The limitations of cache prefetching on a bus-based multiprocessor system are investigated in [Tull93], where it is shown that when bus bandwidth is the bottleneck, prefetching will not improve performance even when it cuts the demand miss ratio.

Other recent work on prefetching worthy of mention is that in [Chi94, Gorn94, Pou94]. In [Gorn94], a prefetching scheme which combines hardware and software features is presented and analyzed. [Pou94] concentrates on prefetching in shared memory multiprocessors running scientific (vector) workloads. [Chi94] likewise considers prefetching in a vector environment with constant stride.

Another disadvantage of cache prefetching is that useless prefetches may pollute cache contents ("memory pollution") by displacing useful cache blocks from the cache and thus cause new cache misses which would not have happened had there been no prefetching; this effect is analyzed for disk caches in [Smit79b]. Memory pollution is most likely when cache sizes are small and block (line) sizes are large.

1.2. Research Issues

Cache prefetching has been implemented in at least two machines, the Amdahl 470V/6 [Smit78] and the Intergraph Clipper [Holl89]. In both cases, performance failed to improve, although for the latter machine, it was at least in part due to a clear implementation flaw. More generally, however, it is not clear that a decrease in the miss ratio attributable to prefetching will actually lead to an improvement in CPU performance. Prefetching can decrease performance even when the miss ratio decreases by: (a) making the cache address tag arrays busy for prefetch lookups; (b) making the cache data arrays busy for prefetch loads and replacements; (c) making the memory bus busy for prefetch address transfers and data fetches; and (d) making the memory system busy for prefetch fetches and replacements.

The issue we consider in this paper is the effect of prefetching on performance, which we evaluate using a detailed cycle by cycle simulator of the CPU memory system. We consider the effect of (a) various prefetch algorithms, (b) while varying various cache parameters such as cache size, block size, and cache associativity, and (c) also while adding hardware resources such double ported cache tags, cache data arrays, a wider memory bus, an interleaved memory system, and buffering for loads into the cache.

The rest of this paper is organized as follows: section 2 gives a detailed description of our architectural model while section 3 describes the methodology we used; section 4 gives a performance overview of cache prefetching with our baseline system; section 5 investigates the impacts of several essential system resources on the performance of cache prefetching strategies; section 6 compares their performance on the best system with that on the baseline system; finally, section 7 concludes the paper. Table 1.1 summarizes terminology used in this paper.

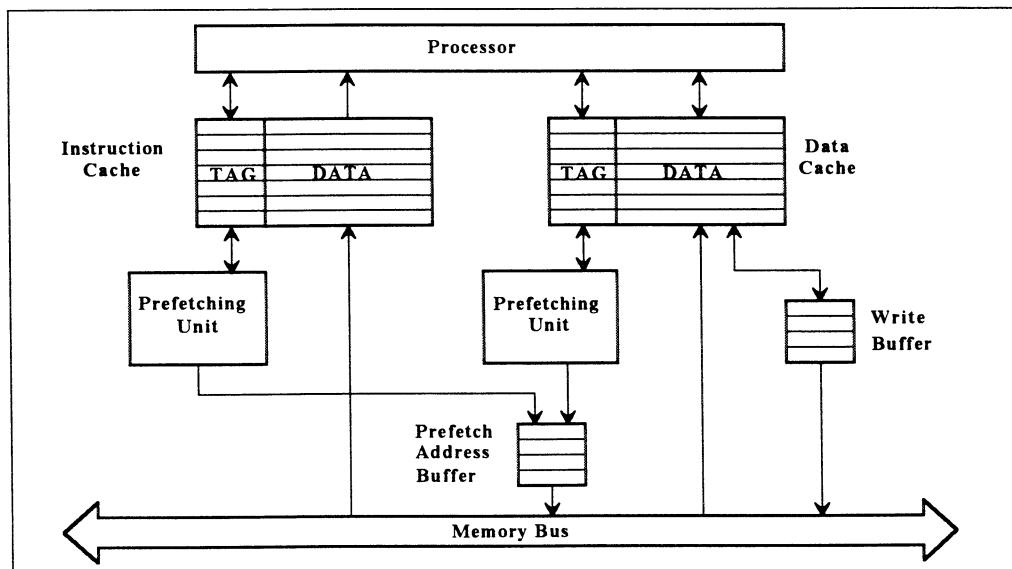


Figure 2.1. Overview of the architectural model.

Term	Definition
Cache Miss	An event in which the address referenced by the processor is not found in the cache and the address, if issued by the prefetching unit, has not been sent to the main memory yet.
Partial Cache Miss	An event in which the address referenced by the processor is not found in the cache but the address has been issued by the prefetching unit and sent to the main memory already.
True Miss Ratio	The ratio of the total number of cache misses to the total number of references the cache made by the processor.
Partial Miss Ratio	The ratio of the total number of partial cache misses to the total number of references to the cache made by the processor.
Total Miss Ratio	Sum of true miss ratio and partial miss ratio.
Issued Prefetch	A prefetch which is sent to the prefetch address buffer by the prefetching unit.
Lifetime of a Prefetch	It starts when a prefetch address is sent to main memory and ends when it is replaced from the cache or referenced by the processor.
Useful Prefetch	An issued prefetch whose address is referenced by the processor during its lifetime.
Useless Prefetch	An issued prefetch whose address is never referenced by the processor during its lifetime.
Aborted Prefetch	An issued prefetch which is discarded in the prefetch address buffer because: (1) the prefetch address buffer runs out of room; or (2) the prefetching address has already been referenced by the processor after the prefetch was issued but the prefetching address has not been sent to the main memory yet.
Prefetch Ratio	The ratio of the total number of issued prefetches to the total number of cache references.
Success Ratio	The ratio of the total number of useful prefetches to the total number of issued prefetches.
Global Success Ratio	The fraction of cache misses which are avoided or partially avoided (i.e. a partial cache miss). See section 4.2 for more details.

Table 1.1. Terminology used in this paper.

2. Architectural Model

This section gives a detailed description of the architectural model which our cycle-by-cycle simulator is based on. Figure 2.1 shows an overview of the model. In each case, we indicate the default design and the various options.

2.1. Processor

Our CPU model is a RISC type processor similar to the MIPS R3000. There are five stages in the pipeline: instruction fetch, instruction decode, ALU computation, read or write to memory, and register file update. The processor fetches an instruction from the instruction cache every clock cycle. Whenever there is a cache miss, the processor halts and waits for the required cache block to be filled. Execution resumes as soon as the required word, rather than the entire cache block, has arrived.

2.2. Caches

We use a split instruction and data cache. Both the instruction and data caches can have either single or double ported tag arrays and data arrays. When an array is double ported, one port is used by the processor and the other by the prefetcher, so in that case, prefetches and processor accesses (to tags, instructions or data) can proceed in parallel without interference. In one case, we also simulate a buffered data array and compare that with a double ported data array.

The default for all caches is 8-way set associative and virtually addressed with a total size of 64KB and a block size of 64 bytes, unless otherwise specified. The data cache uses a write back and write allocate policy. A read hit takes one clock cycle while a write hit takes two (lookup and then modify). The set associativity, cache size and block size are varied in some of our studies.

2.3. Write Buffer

The data cache sends dirty cache blocks back to the main memory via the write buffer. The write buffer is a four entry first-in-first-out (FIFO) queue. It waits until the memory bus is idle before it puts the oldest entry onto the bus. If the write buffer is full, it cannot accept new entries and an attempt to replace a dirty block causes the processor to stall until the write buffer has a free slot. Whenever there is a data cache miss, the write buffer is searched for the missing cache block. If the block is found, it is copied from the write buffer back to the data cache within the same clock cycle.

2.4. Prefetching Units

The two prefetch units, one for each cache, are responsible for issuing new prefetch requests to the main memory. During each clock cycle, each prefetch unit receives information like cache misses, cache hits, instruction types and branch target addresses from the processor and the caches. Based on this information, it decides whether to issue a new prefetch request or not. If it does, the prefetch address is looked up in the corresponding cache. The request is issued in the next clock cycle if the data is not found in the cache.

Issued requests from both prefetch units are not sent directly to the memory bus, though, but to a prefetch address buffer, each of which is organized as a FIFO queue with 16 entries. The oldest entry is sent to the memory bus only when the bus is free. If the buffer is full when a newly issued request arrives, the oldest entry is discarded from the buffer to make room for the new one. Whenever there is a cache miss, the address of the missing cache block is compared against every entry of the buffer; any entry which matches the address represents a failed prefetch (because it is issued too late) and is aborted.

2.5. Memory Bus

The memory bus selectively supports both split and non-split bus transactions [Levy78]. A bus transaction includes two parts: sending the address and receiving or sending the data. In a split bus transaction, the memory bus is idle between sending the address and receiving the data and other transactions are free to use it. In a non-split bus transaction, however, a transaction holds memory bus until it receives its data. By default, the memory bus operates with non-split transactions and is four bytes wide, unless specified otherwise.

Since multiple bus users may compete for the memory bus at the same time, a bus arbitrator is needed to resolve these conflicts. In our model, a fixed priority scheme is used to arbitrate the bus. The priority in descending order is as follows: (1) write back by the write buffer when it is full; (2) a cache miss; (3) returning data from the main memory; (4) write back by the write buffer when it is not full; (5) a new prefetch request.

2.6. Main Memory

The main memory consists of one or more banks, each four bytes wide, and is interleaved [Knut75] on the low order address bits. Each bank has an input queue to buffer requests while the bank is busy and an output queue to buffer returning data when the bus is busy. Because requests to distinct memory banks can be served out of order, a collating buffer holds the returning data until they can be returned to the caches in order. The default memory delay is set to be 16 processor clock cycles while the default number of memory banks is four. In order to simplify the memory hierarchy, we assume that the main memory always contains all referenced pages so that there can be no page faults.

3. Methodology

3.1. Trace Driven Simulation

Trace driven simulation was used for the studies here for the usual well known reasons; see [Smit94] for a discussion of the advantages of the technique. We collected address traces on a DECstation 5000, running version 4.2A of the DEC Ultrix operating system. The benchmark programs were compiled using version 2.3.3 of the gcc compiler for C programs and version 2.1 of the f77 compiler for FORTRAN programs, both with the highest level of optimization flag turned on. The MIPS pixie tool [DEC91] then inserted extra monitoring codes into the compiled binaries so that when we ran the pixified executables, address traces would be collected in destined files.

We ran our simulations on a DEC 3000, running version 1.2A of the OSF/1 operating system. We developed a trace driven simulator which read instruction and data addresses from the address trace, simulated the target memory hierarchical model described earlier on a cycle by cycle basis, and collected and reported the simulation results.

3.2. The Workload

In order to have a substantial, and hopefully representative workload sample, we selected 25 commonly used real programs from five different workload categories: computer-aided design tools (CAD), compiler-related tools (COMP), floating point intensive applications (FP), text processing programs (TEXT), and UNIX utilities (UNIX). We also chose existing inputs of reasonable size and complexity, whenever possible, to run these programs on. Tables 3.1 describes these programs and shows the breakdowns of memory references in the final traces.

Each group trace consisted of five different program traces in the same category and they were interleaved according to a round robin based, preemptive scheduler. Each program trace represented a single process and each process had a distinct process identifier which was added as a prefix to all of its virtual addresses. Since each process had its own distinct address space, there was no need to flush caches after every context switch. Whenever a process executed a system call, it was blocked and put into a sleeping queue for a random number of cycles which followed a uniform distribution in the range of 25,000 to 50,000 cycles. Control was switched to the next process in the ready queue. When the sleeping process woke up, it was put on the tail of the ready queue. In addition, no process was allowed to run more than 400,000 cycles continuously.

In order to minimize start-up effects, the first 70 million memory references of each group trace were not included in the final trace used; instead, they were used to warm-start the instruction and data caches. In order to save space and simulation time, we adopted the scheme described in [Przy90]; we only kept track of references to unique addresses and the last time they were referenced. The one million most recently used unique addresses were added to the front of the final traces, starting with the least recently used address and ending with the most recently used one. We also distinguished reads from writes to the same addresses so that the correct cache blocks would be dirty at the warm-start boundary.

Each final group trace contained roughly 40 million memory references other than those used to warm-start the caches. These traces were originally stored in an ASCII address trace format known as the dinero [Hill84] format. We used a trace compacting scheme similar to *mache* [Samp89] which basically converts each ASCII address into a one to four byte binary integer. We extended *mache* by keeping a counter with each instruction address. The counter informed the simulator of the number of sequential instructions to be executed before the next data read, data write or instruction branch. By keeping track of the program counter, we could now delete those sequential instructions in between from the trace. The size of the final traces compacted by our scheme was about two to five times smaller than that processed by *mache*, and our simulator ran as much as 20% faster. In our opinion, for the purposes of this study the trace lengths were more than sufficient. The purpose of this study was not to obtain absolute miss ratios for large caches, but to explore the behavior of prefetching algorithms. Such behavior should not be sensitive to the length of the trace, but should react primarily to the reference patterns over short lengths of the trace.

4. The Baseline System

4.1. Evaluation Metric

System Parameters	Default Value	System Parameters	Default Value
Cache tag array	double ported	Bus transaction	non-split
Cache data array	single ported	Bus width	4 bytes
Cache size	64K bytes	Memory latency	16 CPU cycles
Cache block size	64 bytes	Number of memory banks	4
Cache associativity	8	Memory bank width	4 bytes
Cache type	split	Prefetch lookahead distance	1 cache block

Table 4.1.1. Default system configuration for the baseline system.

Table 4.1.1 describes the default system settings in the baseline system. These values were chosen according to current technology and common design practice. Unless other values are explicitly stated, they will be used throughout the rest of the paper.

The main metric we used in our evaluation is known as "Cycles Per Instruction contributed by Memory accesses", or MCPI in short [Chen95]. Since we assume that instruction pipelining is perfect and that the processor is capable of executing one instruction per cycle, memory access penalty becomes the sole contributor to CPI. MCPI is given by the following equation:

$$MCPI = \frac{\text{total memory access penalty} - \text{processor stalls due to cache write hits}}{\text{total number of instructions executed}}$$

Our definition of MCPI differs from [Chen93] in that we exclude processor stalls due to cache write hits from our calculation, since write hit stalls are unavoidable and insensitive to cache prefetching.

We prefer MCPI to other evaluation metrics like cache miss ratios and effective memory access time for two reasons:

1. MCPI covers every aspect in performance which can be improved or degraded by a cache prefetching strategy. Not only does it reflect the reduction in cache miss ratios, but it also includes the effects of heavier data bus traffic.
2. MCPI excludes the portions in performance which cannot be affected by a cache prefetching strategy, for example the efficiency of the instruction pipelining.

Because MCPI is a function not only of the cache management algorithms, including prefetching, but also of the cache and system design parameters, its absolute value cannot be used as an evaluation metric. Accordingly, we use the relative MCPI, which compares the performance of a cache prefetching strategy with that when no prefetching is used in the same system. A cache prefetching strategy improves performance only when its relative MCPI is smaller than one. If it is greater than one, the strategy actually degrades performance. Relative MCPI is given as follows:

$$\text{Relative MCPI} = \frac{MCPI \text{ when a cache prefetching strategy is used}}{MCPI \text{ when no prefetching strategy is used in the same system}}$$

4.2. Performance of Cache Prefetching in the Baseline System

In this section, we evaluate the performance of cache prefetching strategies in the baseline system. Table 4.2.1 shows the relative and absolute MCPIs when different prefetching strategies are used. The

results there look discouraging: most strategies perform worse than if there is no prefetching at all. The average relative MCPI for all strategies is 1.2042. In other words, prefetching increases MCPI by more than 20% on average in the baseline system.

But all prefetching strategies do improve performance for the address trace fp. Fp consists of programs like linpack, matrix300 and tomcatv which work on data structures representing matrices and meshes. Since these data structures are large in size and highly sequential in nature, a prefetching strategy is able predict future address references accurately. This explains why aggressive yet simple strategies like always and tag can reduce MCPI by as much as 12%.

The address trace unix, on the other hand, gives the worst results when cache prefetching strategies are used. On average, MCPI increases by more than 48%. The unix trace consists of small programs that work on many different data sets. The memory accessing pattern is often non-sequential, and even backwards, especially in the data cache. Most forward predicting strategies therefore cannot predict future references correctly. Bi-dir, which is designed to predict both forward and backward accessing patterns, therefore performs significantly better than other strategies.

From the drastic difference in performance between fp and unix, one can deduce that a cache prefetching strategy is very sensitive to the type of programs the processor is running. In general, it favors a program which references memory sequentially but works poorly for one which accesses data non-sequentially or backwards.

Cache Prefetching Strategies	Relative MCPI					Trace avg.	Absolute MCPI					Trace avg.		
	Trace Name						Trace Name							
	cad	comp	fp	text	unix		cad	comp	fp	text	unix			
none	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.0863	0.0983	0.1759	0.0306	0.1460	0.1074		
always	1.5914	1.2924	0.8821	1.1770	1.7009	1.3288	0.1373	0.1271	0.1552	0.0360	0.2484	0.1408		
miss	1.2686	1.1505	0.9523	1.1306	1.6159	1.2236	0.1094	0.1131	0.1675	0.0346	0.2360	0.1321		
tag	1.3774	1.1561	0.8793	1.1495	1.6695	1.2463	0.1188	0.1137	0.1547	0.0352	0.2438	0.1332		
bi-dir	1.4329	1.1979	0.8885	1.1712	1.0818	1.1545	0.1236	0.1178	0.1563	0.0358	0.1580	0.1183		
thread	1.0004	0.9968	0.9979	0.9933	1.3515	1.0680	0.0863	0.0980	0.1755	0.0304	0.1974	0.1175		
prefetch avg	1.3341	1.1587	0.9200	1.1243	1.4839	1.2042	0.1151	0.1139	0.1618	0.0344	0.2167	0.1284		

Table 4.2.1. Relative and absolute MCPI for the baseline system.

Table 4.2.2 shows the composition of processor stalls in percentage. Altogether there are five reasons why the processor has to stall:

(1) Cache Misses

The processor is stalled because the memory address it is referencing is not found in the cache and it must wait until the required memory blocks are sent from the main memory. Note that this category does not include stalls due to conflicts with active prefetches. It only includes the raw memory latency of demand misses.

Although all prefetching strategies are able to reduce stalls due to cache misses, they introduce three new categories of stalls, described below, which reduce their effectiveness:

(2) Prefetch (cache) Stalls

The processor is stalled because it cannot access the data port of a cache, which is being used by an active prefetch to load a block into the cache. In our architectural model, a regular memory reference by

the processor has higher priority to access the data port of a cache than a prefetch. However, if a prefetch has already acquired the data port, the processor has to wait until the prefetch finishes its access.

In the baseline system, this category represents a significant fraction of total processor stalls. Conflicts over cache data ports between prefetches and regular processor references seriously affect the performance of cache prefetching. We will study this issue in greater detail in section 5.2.

(3) *Prefetch (bus) Stalls*

A demand miss is delayed because it cannot access the data bus. As in the case of cache data ports, a demand miss has higher priority to access the data bus than a prefetch. But if a prefetch is already using the bus, a demand miss has to wait until the bus is free again.

Conflicts over the data bus represents the second biggest drawback of cache prefetching. Because in the baseline system, the data bus only supports non-split transactions, a prefetch will lock the bus until the prefetched block returns from the main memory. A split transaction data bus would have relieved this contention. We will study this issue more carefully in section 5.6.

(4) *Prefetch (mem) Stalls*

A demand miss is delayed because the memory bank it needs to access is being used by an active prefetch. Here, both demand misses and prefetches have the same priority in accessing the memory banks because the main memory does not distinguish between them.

Since a demand miss or a prefetch will not release the data bus until it finishes accessing the memory bank in the baseline system, there can only be one access to the main memory at any time. This explains why this category is always zero in table 4.2.2. When the effect of split transaction bus is considered in section 5.6, multiple accesses to the main memory become possible and this category of stalls will become more significant.

(5) *Write Buffer Stalls*

The processor is stalled because the write buffer is full. In table 4.2.2, this represents less than 1% of total stalls. This frequency of this kind of stall is insignificant because a four entry deep write buffer is highly effective in buffering dirty cache blocks back to the main memory [Smit79a].

CPU Stall Reasons	Cache Prefetching Strategies						Prefetch Average
	none	always	tag	miss	bidir	thread	
cache misses	99.27	51.56	66.88	53.69	61.07	82.74	63.19
prefetch (cache)	0.00	33.97	19.68	32.28	30.38	8.23	24.91
prefetch (bus)	0.00	13.94	12.86	13.47	7.98	8.38	11.33
prefetch (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.72	0.52	0.58	0.56	0.57	0.65	0.58

Table 4.2.2. Stall breakdowns in percentage for the baseline system.

Table 4.2.3 gives the miss ratios found in the instruction and data caches. We note that the miss ratios there, for the non-prefetching case, are comparable to the design target miss ratios (DTMRs) proposed in [Smit85,87], and to the other miss ratios reported from measured real systems and standard workloads [Gee93a]. We distinguish three kinds of cache miss ratios when a cache prefetching strategy is used; note that these terms are also defined in table 1.1.

The *True Miss Ratio* is the ratio of the total number of true cache misses to the total number of processor references to the cache. A true cache miss is defined as an event in which the address referenced

by the processor is not found in the cache and the address, if issued by the prefetching unit, has not been sent to the main memory yet.

The *Partial Miss Ratio* is the ratio of the total number of partial cache misses to the total number of processor references to the cache. A partial cache miss is defined as an event in which the address referenced by the processor is not found in the cache but the address has already been issued by the prefetching unit and sent to the main memory.

The *Total Miss Ratio* is the sum of the true miss ratio and the partial miss ratio.

The true miss ratio is a better indicator of the accuracy of a cache prefetching strategy than total miss ratio. It is not fair to count a partial cache miss as a genuine demand miss because the prefetching strategy is already halfway to bringing the required block to the cache. If the memory latency were smaller, the prefetch might have already finished and the partial cache miss might not have been a miss at all.

As shown in table 4.2.3, cache prefetching strategies do not reduce total cache miss ratios or true cache miss ratios as much as one would have expected, or has been reported earlier in [Smit82,85]. The reason for this is that the line size used in the base system (64-bytes) is much larger than that used in the earlier studies (16-bytes) and prefetching declines rapidly in effectiveness with larger line sizes.

Cache Type	Cache Miss Ratios	Cache Prefetching Strategies						Prefetch Average
		none	always	miss	tag	bi-dir	thread	
Instr. Cache	True Miss Ratio	0.0011	0.0008	0.0009	0.0008	0.0008	0.0007	0.0008
	Partial Miss Ratio	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002
	Total Miss Ratio	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010
Data Cache	True Miss Ratio	0.0109	0.0065	0.0085	0.0065	0.0065	0.0101	0.0076
	Partial Miss Ratio	0.0000	0.0002	0.0002	0.0001	0.0001	0.0002	0.0002
	Total Miss Ratio	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078

Table 4.2.3. Instruction and data cache miss ratios.

Table 4.2.4 shows the success ratios and global success ratios for the instruction and data caches. The *success ratio* is the ratio of the total number of useful prefetches issued to the total number of prefetches issued. A *useful prefetch* is one which fetches a line that is referenced before it is replaced. Success ratio alone is not sufficient to evaluate the accuracy of a prefetching strategy, because a prefetching strategy may issue only a few prefetches and attain a high success ratio, yet fail to catch most demand misses.

A better metric to evaluate the accuracy of a cache prefetching strategy is the *global success ratio* (GSR), which is the fraction of cache misses which are avoided or partially avoided (i.e. partial cache misses). A GSR of zero implies that a prefetching strategy does not save any cache misses while a GSR of one means that it catches all of them. GSR is defined by the equation:

$$GSR = \frac{\text{total number of correct prefetches}}{\text{total number of correct prefetches} + \text{total number of true cache misses}}$$

The fallacy of success ratio as a metric is illustrated by *thread*. Note that *thread* achieves a very high success ratio, yet its GSR tells us that it does not perform much better than other prefetching strategies. *Thread* is a very conservative strategy; it will not issue a prefetch unless the current cache block has been accessed at least twice. It attains a high success ratio but it fails to capture many demand misses.

Table 4.2.4 helps to explain the small decrease in miss ratios found in table 4.2.3. Because the global success ratios are low (26% for the instruction cache and 36% for the data cache), cache prefetching strategies are not very successful in reducing cache miss ratios in the baseline system. Note again that these results are quite different than those in [Smit82,85] because the baseline system has the longer line size of 64-bytes.

Cache Type	Prefetch Success Ratios	Cache Prefetching Strategies					Prefetch Average
		always	miss	tag	bi-dir	thread	
Instr. Cache	Success Ratio	0.6151	0.6638	0.6613	0.6190	0.9752	0.7069
	Global Success Ratio	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640
Data Cache	Success Ratio	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
	Global Success Ratio	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612

Table 4.2.4. Success ratios and global success ratios.

Figure 4.2.1 gives the compositions of prefetches for the instruction and data caches. We classify prefetches into three different categories:

(1) *Useful Prefetches*

They are prefetches whose addresses are referenced by the processor when they are still residing in the caches or being prefetched from the main memory.

(2) *Useless Prefetches*

They are prefetches whose addresses are not referenced by the processor before they are replaced in the caches. These constitute "memory pollution".

(3) *Aborted Prefetches*

They are prefetches which are discarded in the prefetch address buffer because: (a) the prefetch address buffer overflows; or (b) the prefetch address has already been referenced by the processor after the prefetch was issued but the prefetch address has not been sent to the main memory yet.

Figure 4.2.1 explains why the global success ratio is low for the instruction cache. More than 40% of all prefetches for the instruction cache are aborted. Most of these aborted prefetches are actually correct prefetches, but the prefetching unit fails to send them to the main memory in time because the data bus is busy and they are not issued early enough. Notice that incorrect prefetches are seldom discarded in the prefetch address buffer because: (1) from the simulation results, the prefetch address buffer seldom overflows; and (2) incorrect prefetches will not be referenced when they are still in the prefetch address buffer. As a result, incorrect prefetches are always sent to the main memory successfully while many correct prefetches are aborted.

Although very few prefetches are aborted in the data cache, many prefetches are useless, as is shown in figure 4.2.1. They degrade the performance of cache prefetching by increasing the databus traffic and replacing useful blocks from the data cache.

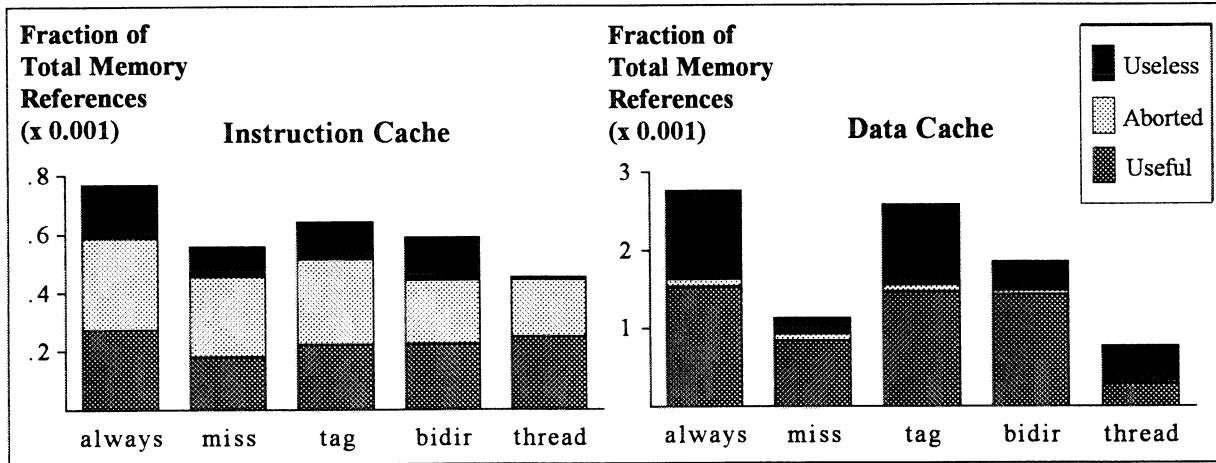


Figure 4.2.1. Average distribution of prefetches.

4.4.3. From Ideal To Baseline

In this section we will study some of the architectural constraints that limit the performance of a cache prefetching strategy. The major limitations which reduce the effectiveness of cache prefetching are, conflicts and delays in accessing the caches, the data bus and the main memory. By running simulations on a hypothetical system in which these limitations are made ideal, we will be able to evaluate the extents to which these constraints affect a cache prefetching strategy. Note that the ideal system that we describe below is just that - *ideal* - some aspects of it are not feasible for a reasonable implementation, or even possible. Our ideal system is made so by:

1. The cache is *ideal* if there is a special access port to the tag array for prefetch lookups; and there is a special access port to the data array for prefetch loads. This port is as wide as the data bus so that there is no need to buffer prefetched blocks coming off the data bus.
2. The data bus is *ideal* if there is a private bus connecting the prefetching unit and the main memory only; and the width of this bus is as large as the prefetch block size so that an entire prefetched block takes only one cycle to transfer from the main memory to the destination cache.
3. The main memory is *ideal* if memory banks are dual ported for regular and prefetch accesses; and prefetches take zero time to access these memory banks.

Prefetching Strategies	CBM							
<i>none</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
<i>always</i>	0.3952	1.1724	0.6636	0.4937	0.8169	1.1245	1.1655	1.2943
<i>miss</i>	0.7166	1.1214	0.8547	0.7716	0.9787	1.0719	1.1118	1.2049
<i>tag</i>	0.3952	1.1117	0.6509	0.4835	0.7890	1.0632	1.1039	1.2186
<i>bi-dir</i>	0.4568	1.0336	0.6509	0.5341	0.7555	1.0354	1.0549	1.1408
<i>thread</i>	0.7142	0.9608	0.8810	0.7783	0.9832	0.9424	0.9840	1.0596
<i>prefetch average</i>	0.5356	1.0800	0.7402	0.6122	0.8647	1.0475	1.0840	1.1836

Table 4.3.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used.

Table 4.3.1 shows the performance of the eight systems in which each component can be either ideal or realistic. We identify each system by a three letter word: the first letter can be C or c, meaning that

the caches are ideal or realistic respectively; the second letter can be B or b, meaning that the data bus is ideal or realistic respectively; and the third letter can be M or m, meaning that the main memory is ideal or realistic respectively.

In the most ideal system (CBM), cache prefetching reduces MCPI by 46% on average. Since there is no penalty in using the cache ports, the data bus and the main memory, system CBM favors aggressive strategies that issue lots of prefetches. Hence always prefetch and tag prefetch perform the best.

Cache prefetching, however, increases MCPI by 18% on average in the real system (cbm). Because resources are limited, aggressive strategies like always and tag are heavily penalized and perform the worst among all strategies.

We analyzed the effects of the three factors, cache, bus and main memory, by using the technique of a sign table [Jain91]. Table 4.3.2 shows the results of this analysis. The effects due to caches alone account for more than 82% of the variations in average relative MCPI. Fortunately, making the caches ideal is feasible but making the data bus and main memory ideal are not.

Effects or Interactions	Magnitude	Variations Explained (%)
Cache	-0.2053	82.14
Bus	-0.0746	10.86
Memory	-0.0335	2.19
(Cache) (Bus)	-0.0396	3.06
(Bus) (Memory)	0.0225	0.99
(Cache) (Memory)	-0.0168	0.55
(Cache) (Bus) (Memory)	-0.0105	0.22

Table 4.3.2. Results from the sign table analysis on average relative MCPI.

5. Effects of System Resources on Cache Prefetching

In this section we consider each of the design parameters separately, and analyze their effect on the effectiveness of prefetching.

5.1. Single vs. Double Ported Cache Tag Arrays

Table 5.1.1 shows the impact of cache tag ports on relative MCPI. We see that the impact of a second cache tag port on performance is enormous when cache prefetching strategies are used. The average relative MCPI when the cache tag arrays are single ported is more than twelve times bigger than that when the arrays are double ported. Conflicts over the cache ports account for an average of 59% of total stalls when the tag arrays are single ported [Appendix B.1].

When the cache tag arrays change from double to single ported, relative MCPI increases fifteen fold for *always* and *thread*, since both strategies look up the cache tag arrays immediately after each memory access made by the processor. The situation for *bi-dir* is even worse: its relative MCPI is almost twice that of *always* and *thread*, because it looks up the tag arrays twice for every memory access: once for the next cache block and once for the block immediately before. Strategies *miss* and *tag*, on the other hand, are only slightly affected by the type of cache tag ports. Because both strategies look up the tag arrays very rarely, their relative MCPIs increase by only 0.05 when the tag arrays change from double to single ported.

In general, if a prefetch strategy looks up the cache tag arrays frequently, extra access ports to the tag arrays are vital. If these ports are not available, contention for them will be so damaging to performance that the strategy becomes virtually useless.

Prefetching Strategies	Cache Tag Arrays	
	single ported	double ported
<i>none</i>	1.0000	1.0000
always	15.9536	1.2943
miss	1.2415	1.2049
tag	1.2676	1.2186
bi-dir	29.0202	1.1408
thread	15.6638	1.0596
prefetch average	12.6293	1.1836

Table 5.1.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same type of cache tag arrays.

5.2. Single vs. Double Ported Cache Data Arrays, and Buffering

Table 5.2.1 gives the relative MCPI when the cache data arrays are single and double ported; note that the tag arrays are double ported in this case. The issue of single or double ported cache data arrays is far less important than that of the cache tag arrays, because prefetch strategies access the data arrays far less frequently than they do to the tag arrays. When the data arrays are single ported, however, all prefetching strategies give a relative MCPI greater than one. But when the data arrays are double ported, all strategies except thread are able to reduce MCPI and give a relative MCPI smaller than one. On average, relative MCPI decreases significantly by about 0.32.

From [Appendix B.2], conflicts over cache ports total about 13% of all stalls on average when the data arrays are single ported. When they are double ported, these conflicts completely vanish and result in no processor stalls.

Prefetching Strategies	Cache Data Arrays		
	single ported	double ported	single ported (buffered)
<i>none</i>	1.0000	1.0000	1.0000
always	1.2943	0.8169	0.9135
miss	1.2049	0.9787	0.9846
tag	1.2186	0.7890	0.8781
bi-dir	1.1408	0.7555	0.8402
thread	1.0596	0.9832	0.9898
prefetch average	1.1836	0.8647	0.9212

Table 5.2.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same type of cache data arrays.

Double ported data arrays are no doubt a highly desirable feature for cache prefetching strategies. However, adding an extra access port to a cache data array is very costly. A more practical solution would be to provide some buffering for the data arrays. When using a buffering, when a prefetched block arrives

from the main memory, it is first stored in a buffer. When the data array becomes free, the contents of the buffer are loaded into the data array. If the processor needs to access the data array while the load is taking place, the prefetch load is halted to allow access to the processor. After the processor finishes its access, the load will resume. In this way, the processor does not need to wait to access the data arrays and the data bus will be freed once a prefetched block gets off the bus.

Table 5.2.1 shows that such a buffering scheme indeed effectively improves the performance of cache prefetching at a lower hardware cost. Relative MCPI decreases by 0.26 on average and is comparable to that when the cache data arrays are double ported.

5.3. Cache Size

Figure 5.3.1 depicts the relative MCPI for various cache sizes. Performance of a prefetch strategy generally improves as the cache size increases, although even for a 1MB cache, in the base system prefetching still does not improve performance.

When a cache is larger, a prefetched block is more likely to reside in the cache for a longer period of time before it is replaced. Consequently, it has a higher chance of being referenced by the processor while it is still in the cache. Moreover, a useless prefetch is less likely to replace a useful block, i.e. to pollute the cache, if the cache is larger. As shown in [Appendix B.3], the fractions of useless prefetches in both caches indeed drop dramatically with larger caches.

Another reason for the improvement with cache size is that there is less interference caused by prefetches in large caches. The bigger the caches, the fewer the cache misses; hence prefetches are less likely to interfere with normal cache operations.

In general, cache prefetching strategies should be used only when there is sufficient space in the caches; otherwise the effects of interference and cache pollution due to useless prefetches will overshadow any benefits brought about by cache prefetching.

5.4. Cache Block Size

Figure 5.4.1 illustrates the effect of cache block size on prefetching. Note that we assume that in all cases, the block size is the same in the instruction cache, the data cache, and is the transfer size for both demand and prefetch transfers. Figure 5.4.1 shows that when cache blocks are 16 or 32 bytes long, most prefetch strategies perform better than when there is no prefetching. But as the cache block size increases, relative MCPI rises above one and prefetching hurts performance. This confirms some results in [Smit78], which showed that the miss ratio improvement for prefetching increases with decreasing block size.

The increase in relative MCPI with increasing block size is mainly a result of more cache port conflicts. [Appendix B.4] shows that cache port conflicts account for only 2.19% of total stalls on average when cache blocks are 16 bytes long. But when cache blocks are 16 times larger, the percentage rises to 20.10%. Because the port width of the data arrays remains the same but the cache block size increases, a cache needs more cycles to load a prefetched block into its data array. Consequently, a single prefetch loading holds up the data array for more consecutive cycles and creates more chances for port conflicts.

Note also that the bigger the cache block, the less useful cache prefetching is. It is because fetching a larger cache block is very similar to fetching a smaller one and always prefetching the next one.

5.5. Cache Associativity

Figure 5.5.1 illustrates the impact of cache associativity on prefetching. When the caches change from direct mapped to two-way mapped, relative MCPI decreases by 0.07 on average, but remains almost constant as cache associativity increases further.

[Appendix B.5] shows that the prefetch distributions exhibit similar trends. There are more prefetches issued, both useless and useful ones, in a direct mapped cache than in a set associative cache. This is because in a direct mapped cache, a cache block can be placed in only one slot. A prefetched block, therefore, has a higher chance of replacing a useful block from the cache. The prefetch unit detects more potential misses and issues more prefetches. Bus traffic increases and more conflicts arise. In a multiple mapped cache, a prefetched block can be placed in more than one slot, and hence the problem of cache pollution is less serious than that in a direct mapped one.

From these results, we conclude that a cache prefetch strategy performs more effectively in set associative caches than in direct mapped caches.

5.6. Split vs. Non-Split Bus Transaction

Prefetching Strategies	Bus Transaction	
	non-split	split
<i>none</i>	1.0000	1.0000
always	1.2943	1.0937
miss	1.2049	1.0100
tag	1.2186	1.0335
bi-dir	1.1408	1.0547
thread	1.0596	0.9309
prefetch average	1.1836	1.0246

Table 5.6.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same type of bus transaction.

Table 5.6.1 shows the relative MCPI with the two types of bus transactions. Relative MCPI decreases by 0.16 on average when bus transactions change from non-split to split. For both caches, most aborted prefetches become useful prefetches when the data bus supports split transactions [Appendix B.6]. Recall that an aborted prefetch is one which is canceled in the prefetch address buffer before it can be sent to the main memory. Although most aborted prefetches correctly predict future memory references, they fail to be sent to the main memory because the data bus is too busy. When split transactions are available, a transaction does not have to hold up the bus during its entire access to the main memory. It is therefore easier for the prefetching unit to acquire the bus and send prefetch requests to the main memory. Many would-be aborted prefetches now become useful. True miss ratios decrease and the prefetch strategy is able to improve MCPI more.

The type of bus transaction also affects the impact of other system resources on cache prefetching. For example, the number of memory banks has very little effect on relative MCPI unless the data bus supports split transactions. We will consider this issue later in section 5.9.

5.7. Bus Width

Figure 5.7.1 shows the relative MCPI when the bus width varies. In order to take the full advantage of a wider bus, we assume that the cache data ports are as wide as the data bus so that when a prefetched block arrives, a cache can load the block into its data array in a single cycle and hence no buffering is needed.

Figure 5.7.1 shows that as the bus width increases, relative MCPI begins to fall below one for the base system. The major reason is because there are fewer cache port conflicts. [Appendix B.7] shows that

the percentage of total stalls contributed by cache port conflicts drops by 10% on average when the bus width increases from 4 bytes to 16 bytes. Since a cache block is 64 bytes long by default, a cache takes 16 cycles to finish a prefetch loading when the bus is 4 bytes wide. But when the bus is 16 bytes wide, it only takes 4 cycles to do so. A wider bus therefore relieves the contention for cache data ports.

However, the above argument holds only when the cache data ports are as wide as the data bus. If the ports are smaller in size, a prefetch load will take multiple cycles to finish. In this case, we need some buffering schemes for the data arrays or we need to lock the bus until the prefetch load finishes.

5.8. Memory Latency

Figure 5.8.1 illustrates the performance of prefetch strategies as the memory latency increases. Although relative MCPI decreases as the memory latency increases from 8 to 64 processor cycles, it starts to rise as the memory latency increases further. Two reasons account for these U-shaped curves:

(1) fewer cache port conflicts

When the memory latency increases, prefetch loadings happen less frequently in the same period of time. Conflicts over cache ports occur more rarely and hence the processor is stalled less. This explains why relative MCPI falls. [Appendix B.8] shows that the percentage of total stalls contributed by these conflicts drops by 13.4% on average when the memory latency increases from 8 to 128 processor cycles.

(2) more data bus conflicts

As the memory latency increases, a prefetch has to hold the bus longer and is more likely to conflict with a regular demand miss. Hence the processor is stalled more. This explains why relative MCPI starts to rise when the memory latency is larger than 64 cycles. [Appendix B.8] shows that the percentage of total stalls contributed by bus conflicts rises by 10.6% on average when the memory latency increases from 8 to 128 processor cycles.

If the performance gap between memory and processor speed continues to increase, cache prefetching will become increasingly useful until the memory latency passes the ideal latency for that architecture. For our base system, the ideal latency is around 64 processor cycles.

5.9. Number of Memory Banks

If bus transactions are non-split, having more memory banks cannot improve the performance of cache prefetching significantly, because a non-split bus transaction does not allow multiple prefetches to access the main memory concurrently. Therefore, in this section, we assume that bus transactions are split.

Figure 5.9.1 depicts the impact of the number of memory banks on cache prefetching. Relative MCPI decreases by 0.18 on average when the number of memory banks increases from one to four. When there are more memory banks, relative MCPI levels off.

One reason for this improvement is because there are fewer conflicts over the memory banks. [Appendix B.9] shows that these conflicts account for 16.0% of total stalls when there is only one memory bank. But when there are 16 of them, the percentage drops to 1.8%.

Another advantage to multiple banks is that multiple prefetches can occur in parallel when there are more memory banks. As long as the cache doesn't experience a demand miss, any memory reference can initiate a prefetch. However, if there is only one memory bank for a prefetch to access, then the prefetches must be processed sequentially by the memory.

5.10. Bus Traffic

Bus traffic is another aspect of a cache prefetching strategy. Previous measurements [Smit85] shows that prefetching increases bus traffic from 20% to 40%. This additional traffic may seriously affect

the performance of the system. In order to study the effect of bus traffic, we added an extra synthetic load on the bus by simulating a direct memory access (DMA) of variable speed made by some imaginary IO devices. The priority of these DMA accesses to acquire the data bus is higher than that of ordinary prefetches but lower than those of demand misses and write backs.

Figure 5.10.1 illustrates the effect of bus traffic on prefetching. As the amount of bus traffic increases, relative MCPI starts to converge to one. This is because there is less and less bus bandwidth to send prefetch requests to the main memory. [Appendix B.10] shows that when the amount of bus traffic increases, more and more prefetches become aborted until all issued prefetches are aborted. Since it becomes increasingly difficult to access the data bus, very few prefetch requests can be sent to the main memory successfully. When the bus utilization used by DMA reaches 80%, almost no prefetches can acquire the data bus and the processor runs as if there is no prefetching.

Figure 5.10.1 seems to suggest that heavier bus traffic actually improves the relative performance of cache prefetching since relative MCPI decreases. This is true in the baseline system because prefetching does not improve performance. Heavier bus traffic helps to reduce the amount of these undesirable prefetches and hence relative MCPI decreases. But for a system in which prefetching is beneficial, heavier bus traffic will reduce the amount of desirable prefetches and relative MCPI will rise. In any case, relative MCPI converges to one as the data bus becomes more congested.

We also note that shared memory multiprocessor systems, particularly those with a shared bus, are typically limited in their throughput by the bus bandwidth [Gee93b]. In such a case, prefetching will almost certainly lower overall throughput.

5.11. Prefetch Lookahead Distance

Ideally, we would like to issue a useful prefetch well in advance so that by the time the prefetch target is referenced by the processor, the block will already be in the cache. This can be accomplished by prefetching block $p+LA$ instead of block p , where p is the original block address requested by the prefetching strategy and LA is the lookahead distance.

Figure 5.11.1 shows the relative MCPI for different prefetch lookahead distances. For all strategies except *thread*, relative MCPI rises with increasing lookahead distance. This is because the chance that block $p+LA$ will be referenced by the processor decreases as LA increases since the effect of spatial locality is diminishing. As a result, more prefetches are not referenced by the processor before they are replaced from the caches. [Appendix B.11] shows that the number of useless prefetches increases with higher LA . More bus and cache port conflicts occur and performance degrades.

Kim et al. suggested a better way to increase the lookahead distance for *thread* algorithm in [Kim93]. Recall that *thread* records all blocks brought into the cache by the block currently accessed by the processor. When the current block is later reaccessed, the recorded blocks will be prefetched if they are not already in the cache. To increase the lookahead distance, we can keep track of all the blocks brought into the cache by the next block rather than the current block. We implemented this scheme and found out that relative MCPI decreases with increasing lookahead distance. Our results are consistent with those presented in [Kim93]. Because this scheme is able to maintain a high prefetch accuracy, increasing lookahead distance does not generate more useless prefetches. Since we are prefetching earlier, it is more likely that a prefetched block is already in the cache when the processor references it. Hence relative MCPI improves.

From our results, increasing lookahead distance simply by prefetching block $p+LA$ instead of block p is generally not a good idea.

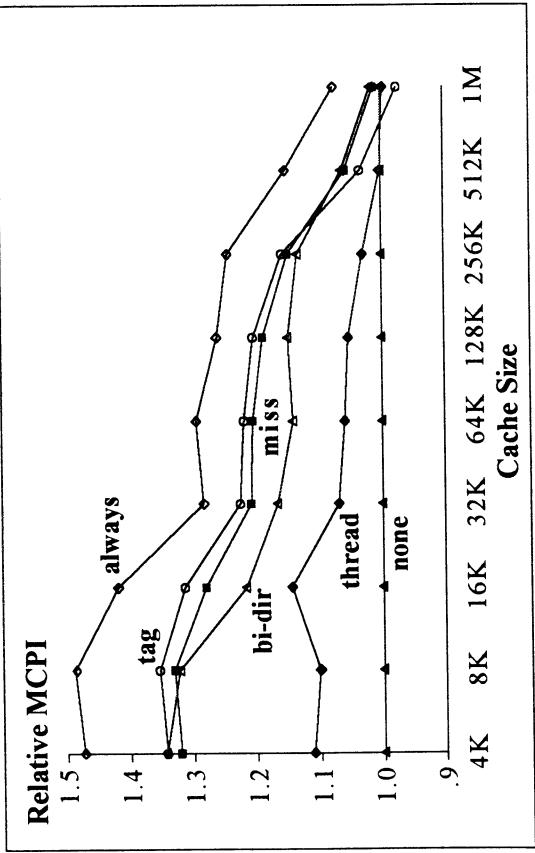


Figure 5.3.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same cache size.

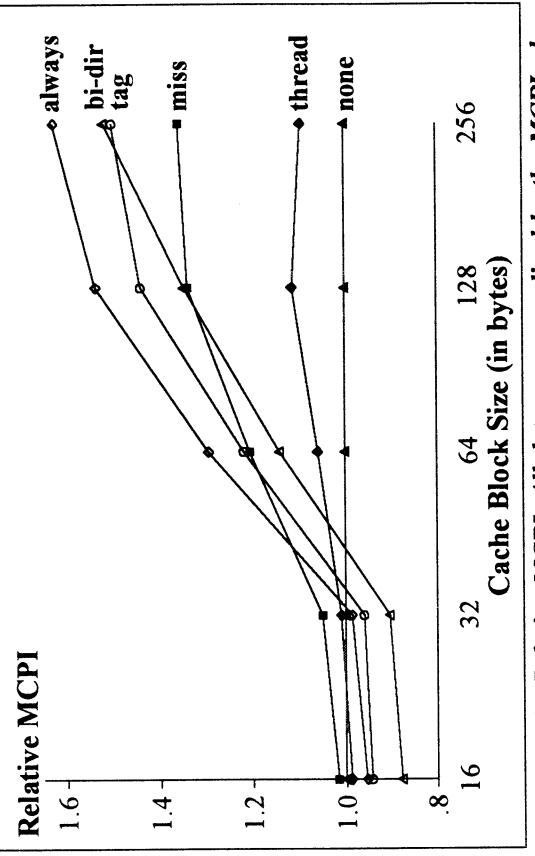


Figure 5.4.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same cache block size.

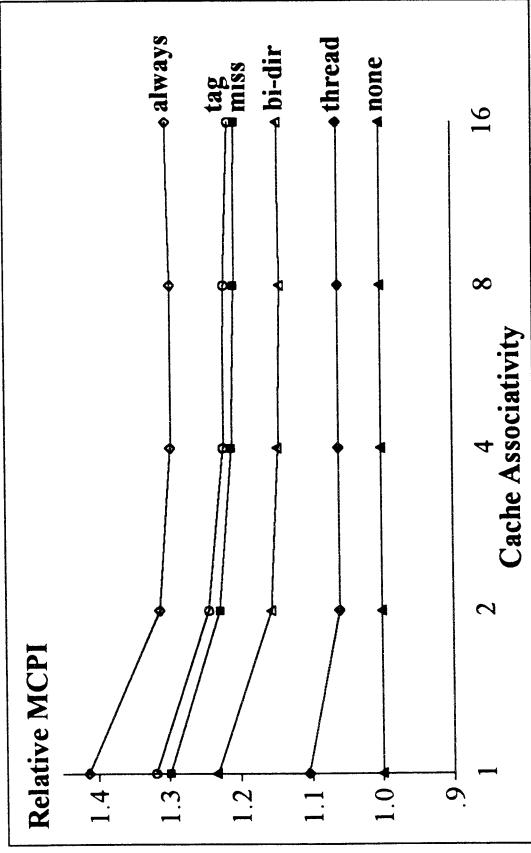


Figure 5.5.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same cache associativity.

Figure 5.7.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same bus width.

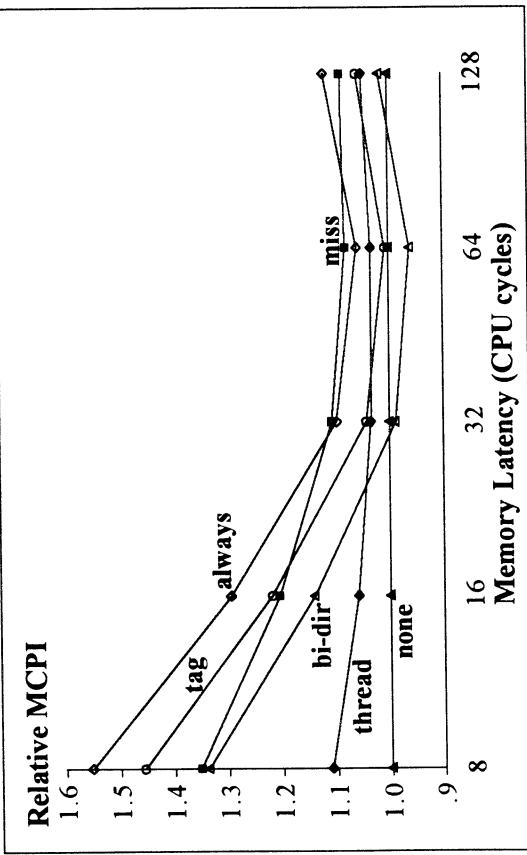


Figure 5.8.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same memory latency.

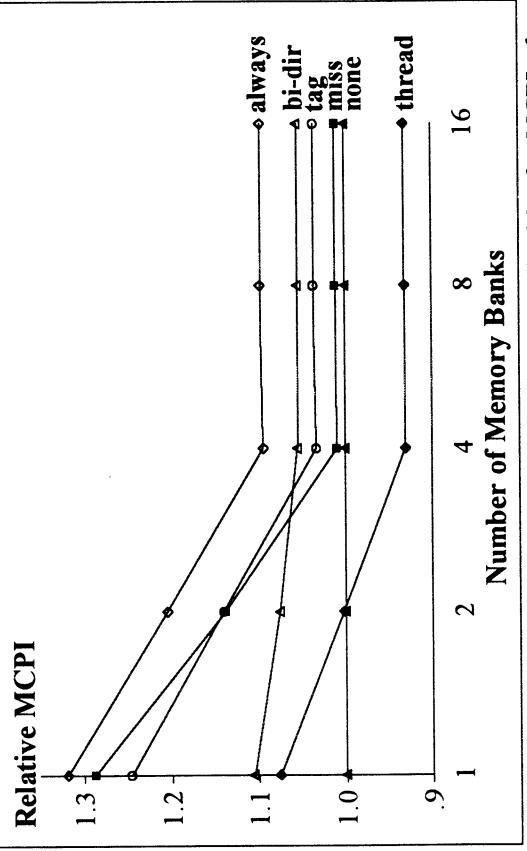


Figure 5.9.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same number of memory banks.

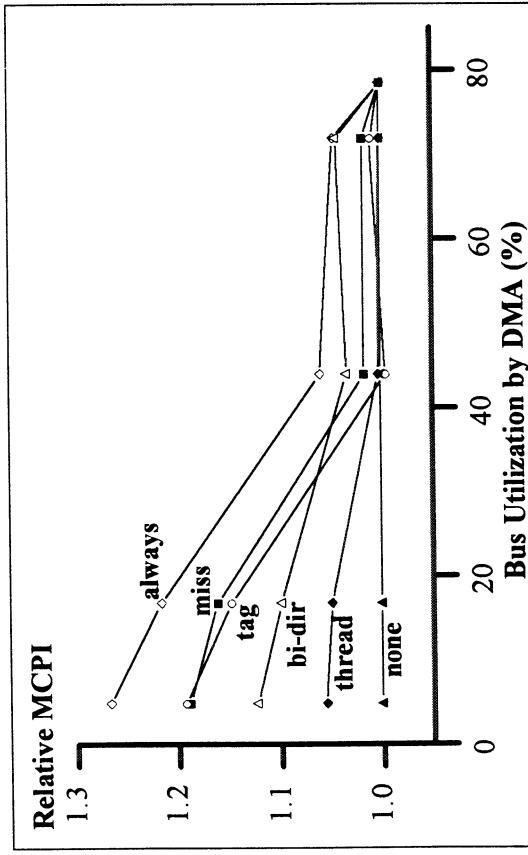


Figure 5.10.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same DMA bus utilization.

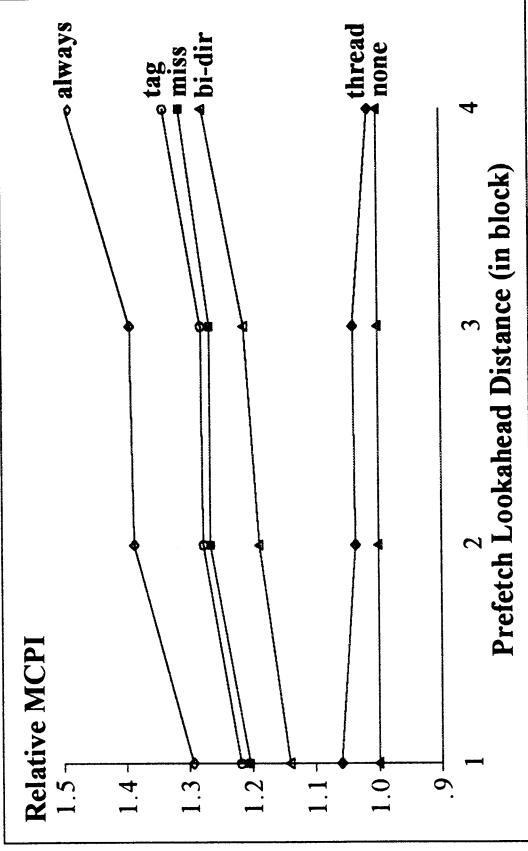


Figure 5.11.1. Relative MCPI. All data are normalized by the MCPI when no prefetching strategies are used with the same prefetch lookahead distance.

6. System Designs

6.1. The Best System for Cache Prefetching

System Parameters		Worst Value	Relative MCPI	Best Value	Relative MCPI	% Improve
Cache tag array		single	12.6293	double	1.1836	90.63
Cache data array		single	1.1836	double	0.8647	26.94
Cache size		8KB	1.3193	1MB	1.0508	20.35
Cache block size		256 bytes	1.4206	16 bytes	0.9562	32.69
Cache associativity		1	1.2733	16	1.1836	7.04
Bus transaction		non-split	1.1836	split	1.0246	13.43
Bus width		4 bytes	1.1836	16 bytes	0.8630	27.09
Memory latency		8 cycles	1.3613	64 cycles	1.0284	24.45
No. of memory banks		1	1.2067	16	1.0260	14.97
Lookahead distance		4 blocks	1.2868	1 block	1.1836	8.02

Table 6.1.1. The worst and the best values for each system parameter. The average relative MCPI shown here are the same as those reported in section 5 where only one particular parameter is changed in the baseline system. The last column gives the improvement in average relative MCPI when a particular parameter changes from its worst value to its best.

Table 6.1.1 gives a summary of the impacts of various system resources on the performance of cache prefetching. We see that when a system parameter changes from its worst value (an unfavorable setting for prefetching) to its best (a favorable one), the improvement in the performance of cache prefetching ranges from 7% to as high as 90%. In this section, we will focus our attention on a system in which each one of these parameters is chosen favorably (but feasibly) for prefetching. We call it the best system; the configuration for the best system is given by the best values shown in table 6.1.1 except for the cache size. Since the cache size is one of the most important factors in system performance, we choose the instruction and data cache size in the best system to be the same as that in the baseline system (64K bytes) in order to facilitate our comparisons.

Cache Prefetching Strategies	Relative MCPI					Absolute MCPI				
	Trace Name					Trace avg.	Trace Name			
	cad	comp	fp	text	unix		cad	comp	fp	text
none	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.4292	0.6273	2.3140	0.2090
always	0.5941	0.5133	0.3110	0.4295	0.8686	0.5433	0.2550	0.3220	0.7197	0.0897
miss	0.6980	0.6313	0.5202	0.6031	0.9186	0.6743	0.2996	0.3960	1.2038	0.1260
tag	0.5902	0.5121	0.3111	0.4428	0.8678	0.5448	0.2533	0.3212	0.7200	0.0925
bi-dir	0.8069	0.7308	0.4583	0.6879	0.8912	0.7150	0.3463	0.4584	1.0605	0.1437
thread	0.7836	0.7076	0.7882	0.7073	0.8941	0.7761	0.3363	0.4438	1.8238	0.1478
prefetch avg	0.6946	0.6190	0.4778	0.5741	0.8881	0.6507	0.2981	0.3883	1.1056	0.1200

Table 6.1.2. Relative and absolute MCPI for the best system.

Cache Prefetching Strategies	Relative CPI					Trace avg.	Absolute CPI					
	Trace Name						Trace Name				Trace avg.	
	cad	comp	fp	text	unix		cad	comp	fp	text		
<i>none</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.5253	1.7301	3.4329	1.3141	1.5319	1.9068
<i>always</i>	0.8861	0.8239	0.5361	0.9097	0.9629	0.8237	1.3515	1.4255	1.8399	1.1955	1.4749	1.4575
<i>miss</i>	0.9153	0.8666	0.6770	0.9371	0.9770	0.8746	1.3961	1.4993	2.3240	1.2315	1.4966	1.5895
<i>tag</i>	0.8851	0.8235	0.5361	0.9117	0.9626	0.8238	1.3499	1.4249	1.8403	1.1982	1.4747	1.4576
<i>bi-dir</i>	0.9460	0.9029	0.6354	0.9508	0.9693	0.8808	1.4428	1.5620	2.1810	1.2494	1.4848	1.5840
<i>thread</i>	0.9391	0.8941	0.8573	0.9536	0.9699	0.9228	1.4325	1.5470	2.9420	1.2531	1.4857	1.7321
prefetch avg	0.9143	0.8622	0.6484	0.9326	0.9684	0.8652	1.3945	1.4917	2.2254	1.2255	1.4833	1.5641

Table 6.1.3. Relative and absolute CPI for the best system.

Table 6.1.2 and 6.1.3 show the relative and absolute MCPI and CPI for different address traces in the best system. The results are much better than those found in the baseline system. All prefetching strategies perform better than when there is no prefetching. The relative MCPI for all strategies averages 0.65, meaning that prefetching reduces MCPI by 35% on average relative to the baseline system.

Prefetching performs the best on the address trace fp. MCPI decreases by more than 52% on average. *Always* and *thread* even reduce MCPI by as much as 69%. Prefetching continues to perform the worst on the address trace unix, for which the reduction in MCPI averages 11%.

The best system favors aggressive strategies like *always* and *tag* which issue lots of prefetches. The baseline system, on the other hand, favors conservative strategies like *thread* and *bi-dir* which only issue prefetches with high chances of being referenced. This difference arises from the fact that there are more resources and bandwidth available in the best system for prefetching than in the baseline system. Strategies which make better use of these resources by aggressively issuing prefetches win in the best system. But when resources are limited, these strategies run into too many conflicts. The type of prefetching strategy a system should use is therefore highly dependent on the availability of system resources.

CPU Stall Reasons	Cache Prefetching Strategies						Prefetch Average
	<i>none</i>	<i>always</i>	<i>miss</i>	<i>tag</i>	<i>bidir</i>	<i>thread</i>	
cache misses	81.25	67.78	74.75	67.98	74.65	76.47	72.33
prefetch (cache)	0.00	0.00	0.00	0.00	0.00	0.00	0.00
prefetch (bus)	0.00	3.95	0.95	3.86	2.43	0.92	2.42
prefetch (mem)	0.00	0.01	0.00	0.01	0.00	0.00	0.00
write buffer	1.01	0.19	0.50	0.19	0.40	0.91	0.44
write hits	17.74	28.09	23.80	27.97	22.51	21.70	24.81

Table 6.1.4. Stall breakdowns in percentage for the best system.

Table 6.1.4 shows the stall compositions. There are no conflicts over the cache data ports because the data arrays are dual ported. Contention over the data bus is minimal and it accounts for less than 2.5% of total stalls on average. Since there are 16 memory banks available, conflicts over memory banks occur very rarely. Because conflicts are rare, the best system does not heavily penalize useless prefetches. An

aggressive strategy takes advantage of this fact by issuing more prefetches and is therefore better able to improve MCPI.

Table 6.1.5 lists the miss ratios found in the instruction and data caches. Prefetching strategies lower both the total miss ratios and true miss ratios significantly in the best system. This is because prefetch accuracy is greatly improved. As shown in table 6.1.6, both the success ratios and global success ratios are high. The global success ratios in the instruction and data caches average 68% and 45% respectively. For comparison, the global success ratios in the baseline system averaged about 26% and 36% for the instruction and data caches respectively.

Figure 6.1.1 shows what happens to the prefetches. Because there is abundant bus bandwidth, most prefetch requests can be sent to the main memory successfully and almost no prefetches are aborted. On average, over 80% and 69% of all prefetches are useful in the instruction and data caches respectively.

Cache Type	Cache Miss Ratios	Cache Prefetching Strategies						Prefetch Average
		none	always	miss	tag	bi-dir	thread	
Instr. Cache	True Miss Ratio	0.0024	0.0004	0.0013	0.0004	0.0007	0.0007	0.0007
	Partial Miss Ratio	0.0000	0.0011	0.0000	0.0011	0.0015	0.0009	0.0009
	Total Miss Ratio	0.0024	0.0015	0.0013	0.0015	0.0022	0.0016	0.0016
Data Cache	True Miss Ratio	0.0252	0.0076	0.0160	0.0076	0.0081	0.0204	0.0119
	Partial Miss Ratio	0.0000	0.0060	0.0000	0.0060	0.0095	0.0006	0.0045
	Total Miss Ratio	0.0252	0.0136	0.0160	0.0136	0.0176	0.0210	0.0164

Table 6.1.5. Instruction and data cache miss ratios in the best system.

Cache Type	Prefetch Success Ratios	Cache Prefetching Strategies					Prefetch Average
		always	miss	tag	bi-dir	thread	
Instr. Cache	Success Ratio	0.8509	0.8867	0.8773	0.8333	0.7942	0.8485
	Global Success Ratio	0.8432	0.4514	0.8101	0.7050	0.5768	0.6773
Data Cache	Success Ratio	0.6705	0.6289	0.6769	0.7682	0.7501	0.6989
	Global Success Ratio	0.6039	0.3308	0.5921	0.5553	0.1792	0.4523

Table 6.1.6. Success ratios and global success ratios in the best system.

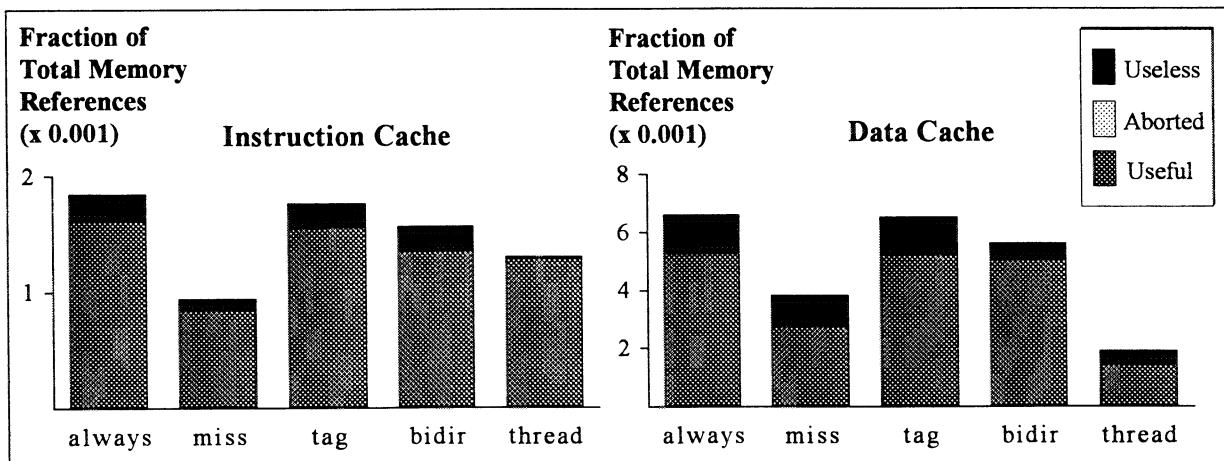


Figure 6.1.1. Average distribution of prefetches in the best system.

6.2. Practical System Designs

Table 6.1.7. shows some reasonable system configurations and the performance of cache prefetching in these systems. The fact that prefetching improves performance in most systems (except systems D and G) is encouraging: it means that prefetching enhances performance not only in the best system but also in systems that can be feasibly implemented. Prefetching strategies perform poorly in systems D and G because the data bus bandwidth is too small (no split transaction, only one memory bank and the data bus is only four bytes wide).

System Parameters	A	B	C	D	E	F	G	H	I	J	K	L
Cache tag array	double											
Cache data array	double	double	buffer	double	double	buffer	double	double	double	double	double	double
Cache size	64KB	64KB	64KB	64KB	64KB	64KB	256KB	256KB	256KB	256KB	1MB	1MB
Cache block size	16	16	16	64	64	64	64	64	64	64	64	64
Cache associativity	2	2	2	2	2	2	2	2	2	2	2	2
Bus transaction	non-sp	split	split	non-sp	split	split	non-sp	split	non-sp	split	non-sp	split
Bus width	4	8	16	4	8	16	4	8	16	16	16	16
Memory latency	16	16	16	16	16	16	16	16	16	16	16	16
No. of mem. banks	1	4	4	1	4	4	1	4	4	4	4	4
Lookahead distance	1	1	1	1	1	1	1	1	1	1	1	1
Relative MCPI	0.8978	0.6204	0.6192	1.0923	0.7305	0.7482	1.0669	0.6961	0.7833	0.7094	0.6805	0.6184
Absolute MCPI	0.1773	0.1144	0.1147	0.3156	0.0711	0.0704	0.1338	0.0303	0.0333	0.0303	0.0126	0.0112

Table 6.1.7. Performance of cache prefetching in some common systems.

7. Conclusion

In memory systems with appropriately designed hardware, prefetching can significantly reduce average memory latency. For a cache to prefetch effectively, it is necessary that the cache tag arrays be double ported, and that the data arrays either be double ported or buffered. The cache should also be at least two way set associative. Prefetching is most effective when the cache is large, the block size small, and the memory bus wide. Using a split transaction bus and interleaving main memory help considerably. Such well designed systems achieve their highest performance with aggressive prefetch strategies which prefetch blocks frequently.

Conversely, in less well endowed memory systems, prefetching either decreases or leaves unchanged the memory system performance. In this case, additional conflicts and bus traffic may overshadow any benefit brought about by prefetching. Such systems function best with conservative strategies which only issue prefetches with high chances of being referenced by the processor.

In this paper, we have only considered single level caches i.e. caches with a processor referencing them on one side, and main memory on the other side. We hope in future work to explore two additional issues. One is the use of prefetching in multiple level caches e.g. on-chip and on-board caches, and the other is the use of prefetching in a multiprocessor system with shared memory.

8. References

- [Baer88] J. L. Baer and W. H. Wang, "On the Inclusion Properties for Multi-Level Cache Hierarchies", Proceedings of the 15th Symposium on Computer Architecture, pages 73-88 (1988).
- [Chen95] Tien-Fu Chen and Jean-Loup Baer, "Effective Hardware-Based Data Prefetching for High-Performance Processors", IEEETC, 44, 5, pages 609-623 (May 1995).
- [Chi94] Chi-Hung Chi, "Compiler Optimization Technique for Data Cache Prefetching Using a Small CAM Array", Proc. 1994 ICCP, Penn State University, vol I, pages 263-266 (August 1994)
- [Dahl93] Fredrik Dahlgren, Michel Dubois, and Per Stenstrom, "Fixed and Adaptive Sequential Prefetching in Shared Memory Multiprocessors", Proceedings of the 1993 International Conference on Parallel Processing, pages I56-I63 (August 1993).
- [DEC91] "Pixie", DEC Ultrix manual page (1991).
- [Fu91] John W. C. Fu and Janak H. Patel, "Data Prefetching Strategies for Vector Cache Memories", Proceedings of the Fifth International Parallel Processing Symposium, pages 555-560 (May 1991).
- [Gee93a] Jeffrey Gee, Mark Hill, Dionisios Penvmatikatos, and Alan Jay Smith, "Cache Performance of the SPEC Benchmark Suite" IEEE MICRO, 13, 4, pages 17-27 (August 1993).
- [Gee93b] Jeffrey Gee and Alan Jay Smith, "Absolute and Comparative Performance of Cache Consistency Algorithms" Technical Report UCB/CSD-93-753 (June 1993)
- [Gind77] J.D. Gindele, "Buffer Block Prefetching Method", IBM Tech. Disc. Bull., 20, 2, pages 696-697 (July 1977).
- [Gorn94] Edward Gornish and Alexander Veidenbaum, "An Integrated Hardware/Software Data Prefetching Scheme for Shared-Memory Multiprocessors", Proc. 1994 ICCP, Penn State University, Vol. II, pages 281-284 (August 1994).
- [Hill84] Mark D. Hill, manual page on DinerolIII, University of California, Berkeley, (October 1985).
- [Holl89] Walter Hollingsworth, Howard Sachs and Alan Jay Smith, "The Fairchild CLIPPER: Instruction Set Architecture and Processor Implementation" Communications of the ACM, 32, 2, pages 200-219 (February 1989).
- [Jain91] Raj Jain, "The Art of Computer Systems Performance Analysis", John Wiley & Sons Inc., pages 283-292 (1991).
- [Kim93] Seong Baeg Kim, et. al., "Threaded Prefetching: An Adaptive Instruction Prefetch Mechanism", Microprocessing and Microprogramming 39:1, pages 1-15 (November 1993).
- [Knut75] D. E. Knuth and G. S. Rao, "Activity in an Interleaved Memory", IEEE Transactions on Computers TC-24:9, pages 943-944 (Sept. 1975).
- [Levy78] J. V. Levy, "Buses: The Skeleton of Computer Structures", Computer Engineering: A DEC View of Hardware Systems Design, Digital Press (1978).
- [Poul94] David Poulsen and Pen-Chung Yew, "Data Prefetching and Data Forwarding in Shared Memory Multiprocessors", Proc. 1994 ICCP, Penn. State University, Vol. II, pages 276-280 (August 1994).
- [Przy90] Steven A. Przybylski, Cache and Memory Hierarchy Design - A Performance-Directed Approach, Morgan Kaufmann Publishers, pages 181-186 (1990).
- [Samp89] Dain Samples, "Mache: No-Loss Trace Compaction", Performance Evaluation Review 17:1, pages 89-97 (May 1989).
- [Shor88] R. T. Short and H. M. Levy, "A Simulation Study of Two-Level Caches", Proceedings of the 15th Annual Symposium on Computer Architecture, pages 81-88 (1988).
- [Smit78] Alan Jay Smith, "Sequential Program Prefetching in Memory Hierarchies", IEEE Computer 11:12, pages 7-21 (December 1978).
- [Smit79a] Alan Jay Smith, "Characterizing the Storage Process and Its Effects on Main Memory Update", Journal of the ACM 26:1, pages 6-27 (Jan. 1979).
- [Smit79b] Alan Jay Smith, "Sequentiality and Prefetching in Data Base Systems", IBM Research Report RJ 1743, March 19, 1976, and ACM Transactions on Data Base Systems, 3:3, pages 223-247 (Sept 1979)
- [Smit82] Alan Jay Smith, "Cache Memories", Computing Surveys 14:3 (September 1982).
- [Smit85] Alan Jay Smith, "Cache Evaluation and the Impact of Workload Choice", Proceedings of the 12th International Symposium on Computer Architecture, pages 64-75 (June 1985).
- [Smit87] Alan Jay Smith, "Line (Block) Size Selection in CPU Cache Memories", IEEE Transactions on Computers, C-36, 9, September, 1987, pages 1063-1075. (Also see correction IEEETC, 38, 6, June, 1989, p. 927.)
- [Smit93] Robert B. Smith, James K. Archibald and Brent E. Nelson, "Evaluating Performance of Prefetching Second Level Caches", Performance Evaluation Review 20:4, pages 31-42 (May 1993).
- [Smit94] Alan Jay Smith, "Trace Driven Simulation in Research on Computer Architecture and Operating Systems", Proceedings New Directions in Simulation for Manufacturing and Communications (SIM94), Tokyo, Japan, August 1-2, 1994, ed. Morito, Sakasegawa, Yoneda, Fushimi, Nakano, pages 43-49 (August 1994).
- [Tull93] Dean M. Tullsen and Susan J. Eggers, "Limitations of Cache Prefetching on a Bus-Based Multiprocessor", Proceedings of the 20th Annual International Symposium on Computer Architecture, pages 278-288 (May 1993).
- [Varm92] Anujan Varma and Gunjan K. Sinha, "A Class of Prefetch Schemes for On-Chip Data Caches", Technical Report, Computer Science Department, University of California, Santa Cruz (1992).

APPENDIX

A.1.	The Baseline System	A1
A.2.	From Ideal to Baseline	A6
B.1.	Single vs. Double Ported Cache Tag Arrays	B1
B.2.	Single vs. Double Ported Cache Data Arrays	B4
B.3.	Cache Size	B7
B.4.	Cache Block Size	B13
B.5.	Cache Associativity	B17
B.6.	Split vs. Non-Split Bus Transaction	B21
B.7.	Bus Width	B24
B.8.	Memory Latency	B27
B.9.	Number of Memory Banks	B31
B.10.	Bus Traffic	B35
B.11.	Prefetch Lookahead Distance	B39
C.1.	The Best System for Cache Prefetching	C1

A.1. The Baseline System

Trace Name	True Miss Ratios (Instruction Cache)							True Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
cad	0.0001	0.0007	0.0008	0.0007	0.0007	0.0006	0.0007	0.0080	0.0072	0.0070	0.0066	0.0066	0.0076	0.0070
comp	0.0024	0.0019	0.0020	0.0019	0.0021	0.0015	0.0019	0.0066	0.0034	0.0046	0.0035	0.0037	0.0061	0.0043
fp	0.0010	0.0007	0.0008	0.0007	0.0007	0.0007	0.0007	0.0178	0.0014	0.0094	0.0015	0.0018	0.0156	0.0059
text	0.0009	0.0006	0.0007	0.0006	0.0007	0.0006	0.0006	0.0017	0.0008	0.0012	0.0009	0.0009	0.0015	0.0011
unix	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0206	0.0199	0.0203	0.0199	0.0197	0.0196	0.0199
trace average	0.0011	0.0008	0.0009	0.0008	0.0008	0.0007	0.0008	0.0109	0.0065	0.0085	0.0065	0.0065	0.0101	0.0076
Trace Name	Partial Miss Ratios (Instruction Cache)							Partial Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
cad	0.0000	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0001	0.0002	0.0001	0.0002	0.0002
comp	0.0000	0.0006	0.0003	0.0004	0.0004	0.0005	0.0004	0.0000	0.0005	0.0005	0.0005	0.0002	0.0000	0.0004
fp	0.0000	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0000	0.0001	0.0001	0.0001	0.0001	0.0000	0.0001
text	0.0000	0.0001	0.0001	0.0001	0.0001	0.0002	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
unix	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0001	0.0001	0.0001	0.0008	0.0002
trace average	0.0000	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002
Trace Name	Total Miss Ratios (Instruction Cache)							Total Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
cad	0.0001	0.0009	0.0009	0.0009	0.0009	0.0008	0.0009	0.0080	0.0074	0.0072	0.0068	0.0068	0.0077	0.0072
comp	0.0024	0.0025	0.0023	0.0023	0.0025	0.0020	0.0023	0.0066	0.0039	0.0051	0.0040	0.0039	0.0062	0.0046
fp	0.0010	0.0009	0.0009	0.0009	0.0009	0.0009	0.0009	0.0178	0.0015	0.0095	0.0015	0.0018	0.0157	0.0060
text	0.0009	0.0007	0.0008	0.0007	0.0008	0.0007	0.0008	0.0017	0.0008	0.0013	0.0009	0.0009	0.0015	0.0011
unix	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0206	0.0200	0.0205	0.0200	0.0198	0.0204	0.0201
trace average	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078

Table A.1.1. Instruction and data cache miss ratios in the baseline system.

Trace Name	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
cad	0.6569	0.6654	0.6741	0.6681	0.9823	0.7293	0.4470	0.4356	0.4478	0.5335	0.7012	0.5130
comp	0.5340	0.5581	0.5637	0.5248	0.9620	0.6285	0.6618	0.6189	0.6825	0.7189	0.6351	0.6634
fp	0.8315	0.8384	0.8382	0.8278	0.9942	0.8660	0.9759	0.9750	0.9772	0.9717	0.7858	0.9371
text	0.6396	0.6493	0.6545	0.6286	0.9375	0.7019	0.7566	0.7344	0.7763	0.7110	0.8837	0.7724
unix	0.4133	0.6076	0.5758	0.4457	1.0000	0.6085	0.1395	0.1300	0.1422	0.3676	0.2092	0.1977
trace average	0.6151	0.6638	0.6613	0.6190	0.9752	0.7068	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
Trace Name	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
cad	0.3591	0.2569	0.3270	0.3189	0.4047	0.3333	0.3555	0.2261	0.3156	0.3744	0.0630	0.2669
comp	0.3178	0.2040	0.2515	0.2524	0.3228	0.2697	0.4878	0.3011	0.4699	0.4291	0.0725	0.3521
fp	0.2335	0.2120	0.2295	0.2166	0.2890	0.2361	0.9209	0.4716	0.9177	0.9028	0.1236	0.6673
text	0.3243	0.2445	0.3199	0.2698	0.3142	0.2945	0.5711	0.3122	0.5315	0.5409	0.1195	0.4150
unix	0.2571	0.1667	0.1966	0.2516	0.0586	0.1861	0.1276	0.1022	0.1243	0.0629	0.1058	0.1046
trace average	0.2984	0.2168	0.2649	0.2619	0.2779	0.2639	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612

Table A.1.2. Success ratios and global success ratios in the baseline system.

Bus Utilization Breakdown (%)	none					always				
	cad	comp	fp	text	unix	cad	comp	fp	text	unix
Icache read	2.41	5.55	2.37	2.13	0.02	1.62	4.49	1.85	1.46	0.02
Dcache read	7.40	5.25	20.16	1.51	15.18	6.34	2.64	1.64	0.71	13.58
Dcache write	0.76	1.13	4.42	0.41	2.18	0.87	1.12	4.60	0.46	2.07
prefetch	0.00	0.00	0.00	0.00	0.00	9.46	8.44	20.25	2.57	14.25
total	10.57	11.93	26.95	4.05	17.38	18.29	16.69	28.33	5.19	29.91
Bus Utilization Breakdown (%)	miss					tag				
	cad	comp	fp	text	unix	cad	comp	fp	text	unix
Icache read	1.80	4.61	1.88	1.62	0.02	1.64	4.39	1.86	1.45	0.02
Dcache read	6.35	3.64	10.77	1.09	13.98	5.88	2.75	1.70	0.76	13.63
Dcache write	0.80	1.13	4.48	0.44	2.08	0.81	1.13	4.58	0.45	2.07
prefetch	5.36	5.10	10.48	1.64	12.25	7.45	6.73	20.10	2.36	13.61
total	14.32	14.48	27.61	4.80	28.33	15.79	15.00	28.23	5.01	29.34
Bus Utilization Breakdown (%)	bi-dir					thread				
	cad	comp	fp	text	unix	cad	comp	fp	text	unix
Icache read	1.71	4.77	1.89	1.58	0.02	1.36	3.59	1.67	1.39	0.02
Dcache read	5.94	2.94	2.01	0.77	14.42	6.98	4.87	17.73	1.34	13.91
Dcache write	0.85	1.13	4.61	0.46	2.17	0.76	1.14	4.43	0.42	2.13
prefetch	8.06	6.73	19.92	2.40	2.65	1.75	2.70	3.88	1.02	7.87
total	16.56	15.56	28.43	5.21	19.26	10.85	12.30	27.72	4.17	23.93

Table A.1.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none						always					
	cad	comp	fp	text	unix	avg	cad	comp	fp	text	unix	avg
cache misses	99.92	99.28	98.60	97.95	99.93	99.13	58.00	62.09	22.35	59.94	57.75	52.10
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	0.00	27.72	25.03	73.82	26.62	17.44	34.13
prefetch(bus)	0.00	0.00	0.00	0.00	0.00	0.00	13.83	12.59	2.27	12.13	24.80	13.12
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.08	0.72	1.40	2.05	0.07	0.86	0.03	0.29	1.55	1.30	0.01	0.64
CPU Stall Reasons	miss						tag					
	cad	comp	fp	text	unix	avg	cad	comp	fp	text	unix	avg
cache misses	71.22	73.81	60.68	72.25	61.77	67.95	62.00	66.56	22.72	62.08	58.86	54.53
prefetch(cache)	18.00	15.99	31.84	16.39	14.86	19.42	23.95	21.75	73.53	24.65	16.77	32.13
prefetch(bus)	10.73	9.78	6.03	9.82	23.35	11.94	13.58	11.32	2.20	11.78	24.35	12.65
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.06	0.42	1.45	1.58	0.01	0.70	0.04	0.36	1.55	1.48	0.01	0.69
CPU Stall Reasons	bi-dir						thread					
	cad	comp	fp	text	unix	avg	cad	comp	fp	text	unix	avg
cache misses	60.29	67.15	23.84	62.41	89.75	60.69	90.00	85.77	86.56	81.33	72.67	83.25
prefetch(cache)	25.41	21.94	72.71	25.53	5.73	30.26	4.84	7.43	9.66	8.82	9.73	8.10
prefetch(bus)	14.26	10.55	1.90	10.79	4.49	8.40	5.20	6.19	2.40	7.97	17.59	7.87
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.05	0.36	1.55	1.26	0.05	0.65	0.08	0.62	1.37	1.88	0.03	0.80

Table A.1.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory.

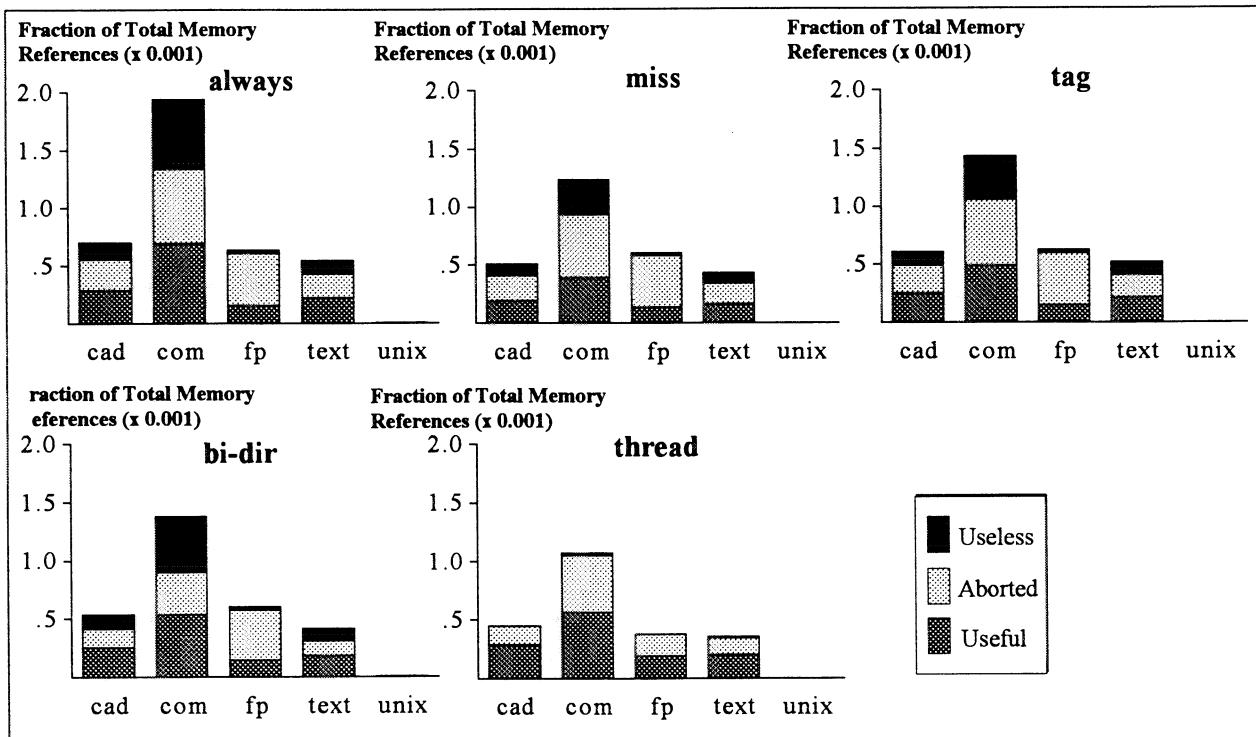


Figure A.1.1. Average distribution of prefetches for the instruction cache.

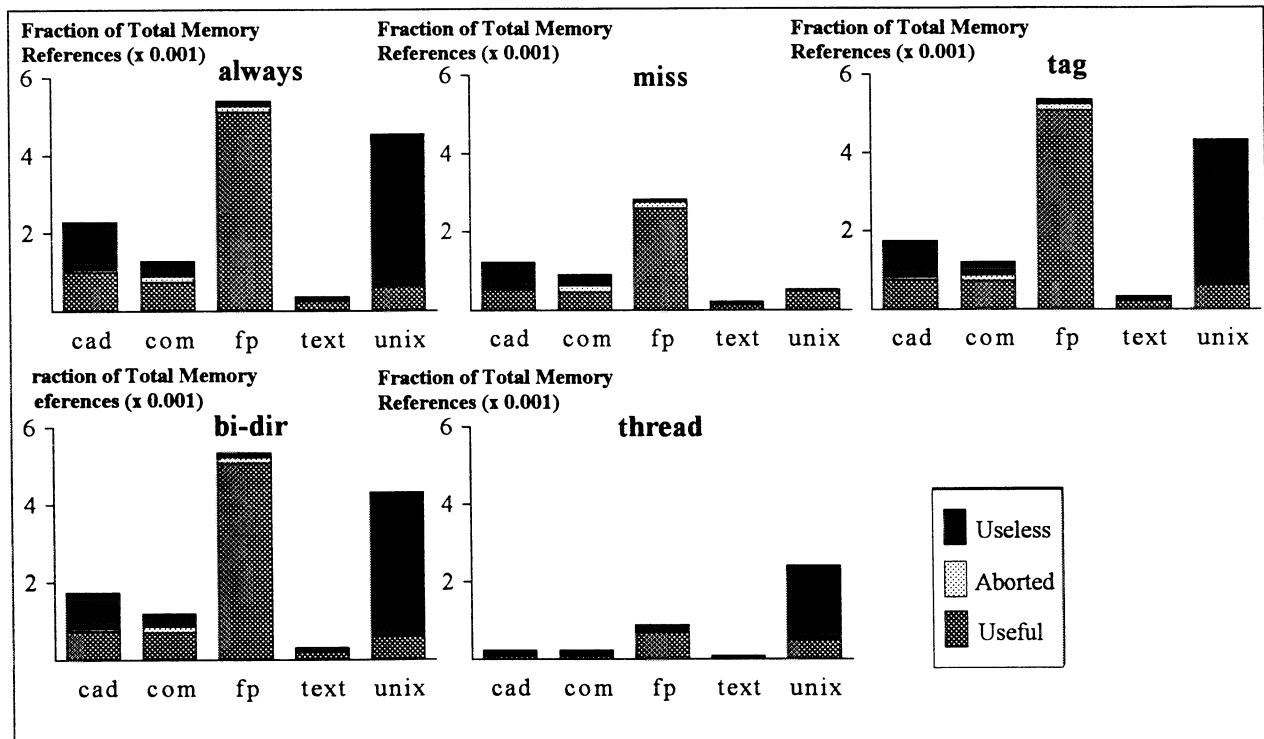


Figure A.1.2. Average distribution of prefetches for the data cache.

System Name	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
CBM	0.7660	0.8247	0.8070	0.7430	0.9855	0.8252	0.6101	0.5901	0.6157	0.6810	0.6992	0.6392
cBM	0.7660	0.8247	0.8070	0.7441	0.9855	0.8255	0.6101	0.5901	0.6157	0.6812	0.6992	0.6393
CbM	0.6219	0.6683	0.6679	0.6221	0.9754	0.7111	0.5957	0.5738	0.6045	0.6648	0.6461	0.6170
CBm	0.7646	0.8247	0.8058	0.7424	0.9834	0.8242	0.5982	0.5888	0.6038	0.6686	0.6846	0.6288
Cbm	0.5982	0.6538	0.6475	0.6040	0.9743	0.6956	0.5923	0.5776	0.6016	0.6556	0.6397	0.6134
cBm	0.7660	0.8330	0.8070	0.7441	0.9855	0.8271	0.6101	0.5896	0.6157	0.6742	0.6992	0.6378
cbM	0.6593	0.6858	0.6993	0.6582	0.9779	0.7361	0.6033	0.5748	0.6104	0.6743	0.6561	0.6238
cbm	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
System Name	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
CBM	0.7384	0.3847	0.6905	0.5707	0.5024	0.5773	0.5442	0.3080	0.5283	0.5151	0.1347	0.4061
cBM	0.7384	0.3847	0.6905	0.5741	0.5024	0.5780	0.5442	0.3080	0.5283	0.5151	0.1347	0.4061
CbM	0.3084	0.2148	0.2713	0.2666	0.2833	0.2689	0.4919	0.2751	0.4703	0.4697	0.0982	0.3610
CBm	0.7380	0.4069	0.6900	0.5721	0.5019	0.5818	0.5399	0.3081	0.5240	0.4938	0.1295	0.3991
Cbm	0.2747	0.2079	0.2460	0.2438	0.2671	0.2479	0.4810	0.2801	0.4610	0.4532	0.0947	0.3540
cBm	0.7384	0.4233	0.6905	0.5723	0.5024	0.5854	0.5442	0.3093	0.5283	0.4957	0.1347	0.4024
cbM	0.3759	0.2353	0.3255	0.3200	0.3183	0.3150	0.5155	0.2791	0.4921	0.4922	0.1039	0.3766
cbm	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612

Table A.2.2. Success ratios and global success ratios in the baseline system.

Bus Utilization Breakdown (%)	none							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
Icache read	2.50	2.50	2.50	2.50	2.50	2.50	2.50	2.50
Dcache read	9.90	9.90	9.90	9.90	9.90	9.90	9.90	9.90
Dcache write	1.78	1.78	1.78	1.78	1.78	1.78	1.78	1.78
prefetch	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
total	14.18	14.18	14.18	14.18	14.18	14.18	14.18	14.18
Bus Utilization Breakdown (%)	always							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
Icache read	0.81	0.76	1.96	0.80	2.04	0.77	1.65	1.89
Dcache read	5.04	4.73	5.28	5.00	5.21	4.74	4.97	4.98
Dcache write	2.00	1.84	1.96	1.99	1.93	1.85	1.85	1.82
prefetch	0.00	6.13	6.08	0.00	11.34	6.17	5.95	10.99
total	7.85	13.46	15.27	7.78	20.53	13.54	14.43	19.68
Bus Utilization Breakdown (%)	miss							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
Icache read	1.58	1.53	2.04	1.51	2.05	1.44	1.94	1.99
Dcache read	7.45	7.08	7.53	7.40	7.37	7.13	7.28	7.17
Dcache write	1.89	1.80	1.87	1.88	1.84	1.81	1.80	1.79
prefetch	0.00	3.63	3.71	0.00	7.08	3.72	3.65	6.97
total	10.92	14.05	15.15	10.79	18.35	14.09	14.68	17.91
Bus Utilization Breakdown (%)	tag							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
Icache read	0.84	0.81	1.93	0.84	2.00	0.81	1.70	1.87
Dcache read	4.96	4.67	5.21	4.92	5.16	4.70	4.94	4.94
Dcache write	1.98	1.82	1.94	1.97	1.92	1.84	1.83	1.81
prefetch	0.00	5.67	5.54	0.00	10.39	5.71	5.42	10.05
total	7.79	12.97	14.63	7.73	19.45	13.05	13.88	18.67
Bus Utilization Breakdown (%)	bidir							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
Icache read	1.27	1.21	2.06	1.25	2.11	1.21	1.85	1.99
Dcache read	5.05	4.92	5.30	5.13	5.36	5.02	5.10	5.22
Dcache write	1.99	1.86	1.96	1.98	1.94	1.86	1.86	1.84
prefetch	0.00	4.44	4.41	0.00	8.30	4.40	4.29	7.95
total	8.30	12.43	13.72	8.36	17.72	12.50	13.09	17.00
Bus Utilization Breakdown (%)	thread							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
Icache read	0.92	0.90	1.61	0.91	1.65	0.90	1.49	1.61
Dcache read	8.99	8.76	9.14	8.95	9.08	8.79	8.99	8.97
Dcache write	1.83	1.79	1.81	1.82	1.79	1.79	1.79	1.78
prefetch	0.00	2.17	1.83	0.00	3.43	2.17	1.88	3.44
total	11.73	13.61	14.40	11.69	15.94	13.65	14.15	15.79

Figure A.2.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
cache misses	99.28	99.27	99.27	99.27	99.27	99.27	99.27	99.27
prefetch (cache)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
prefetch (bus)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
pefetech (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.72	0.72	0.72	0.72	0.72	0.72	0.72	0.72
CPU Stall Reasons	always							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
cache misses	98.80	41.62	89.13	99.07	77.51	45.48	51.75	51.56
prefetch (cache)	0.00	36.10	0.00	0.00	0.00	36.90	43.69	33.97
prefetch (bus)	0.00	21.74	9.74	0.00	21.53	16.71	4.02	13.94
pefetech (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	1.20	0.54	1.12	0.94	0.95	0.92	0.54	0.52
CPU Stall Reasons	miss							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
cache misses	99.06	63.07	92.15	99.14	82.72	65.94	69.89	66.88
prefetch (cache)	0.00	0.00	0.00	0.00	0.00	16.16	25.74	19.68
prefetch (bus)	0.00	36.36	7.05	0.00	16.56	16.88	3.76	12.86
pefetech (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.94	0.57	0.81	0.85	0.72	1.03	0.60	0.58
CPU Stall Reasons	tag							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
cache misses	98.79	43.55	89.73	99.05	78.67	47.27	54.42	53.69
prefetch (cache)	0.00	33.17	0.00	0.00	0.00	35.14	41.01	32.28
prefetch (bus)	0.00	22.72	9.10	0.00	20.31	16.65	3.98	13.47
pefetech (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	1.21	0.55	1.17	0.95	1.02	0.94	0.60	0.56
CPU Stall Reasons	bidir							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
cache misses	99.38	55.25	93.58	99.52	86.33	57.74	60.76	61.07
prefetch (cache)	0.00	41.24	0.00	0.00	0.00	34.41	36.77	30.38
prefetch (bus)	0.00	3.11	5.29	0.00	12.65	7.29	1.88	7.98
pefetech (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.62	0.40	1.14	0.48	1.02	0.56	0.58	0.57
CPU Stall Reasons	thread							
	CBM	cBM	CbM	CBm	Cbm	cBm	cbM	cbm
cache misses	99.14	75.83	95.12	99.27	89.58	80.82	83.39	82.74
prefetch (cache)	0.00	9.60	0.00	0.00	0.00	8.71	13.30	8.23
prefetch (bus)	0.00	13.86	4.13	0.00	9.71	9.72	2.63	8.38
pefetech (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.86	0.71	0.76	0.73	0.70	0.74	0.67	0.65

Table A.2.4. Stall breakdowns in percentage.

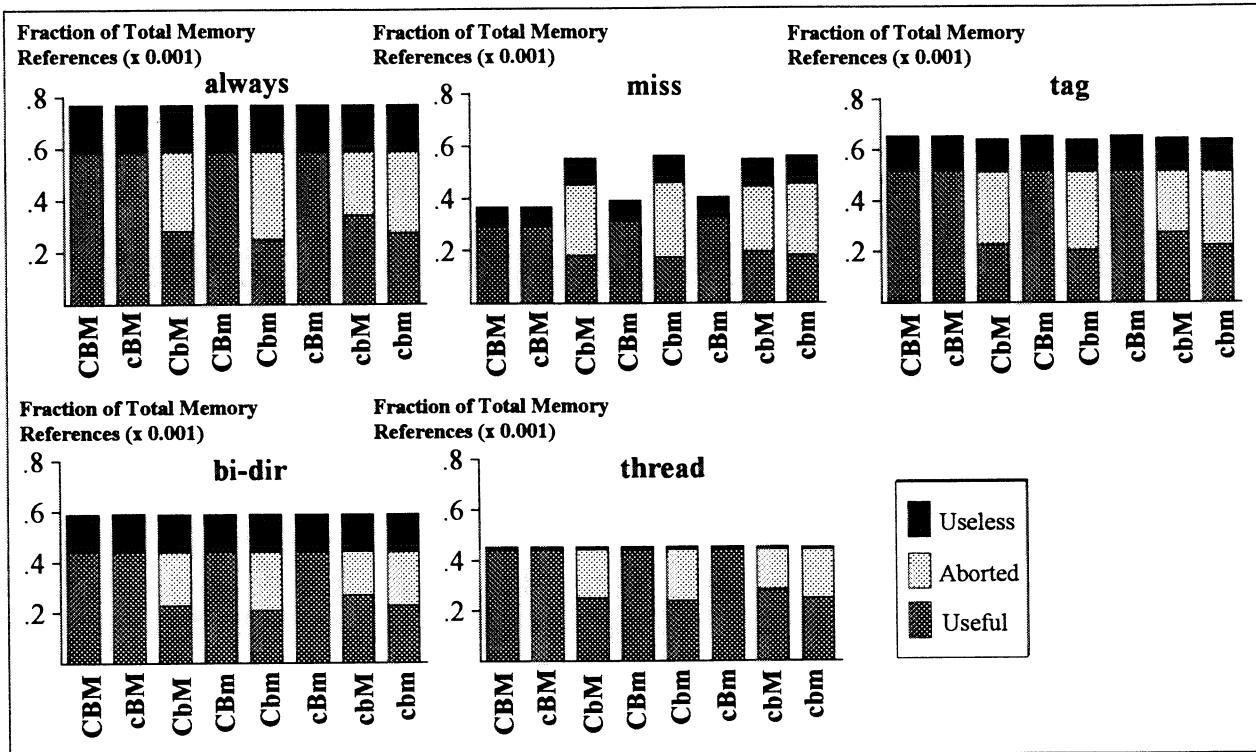
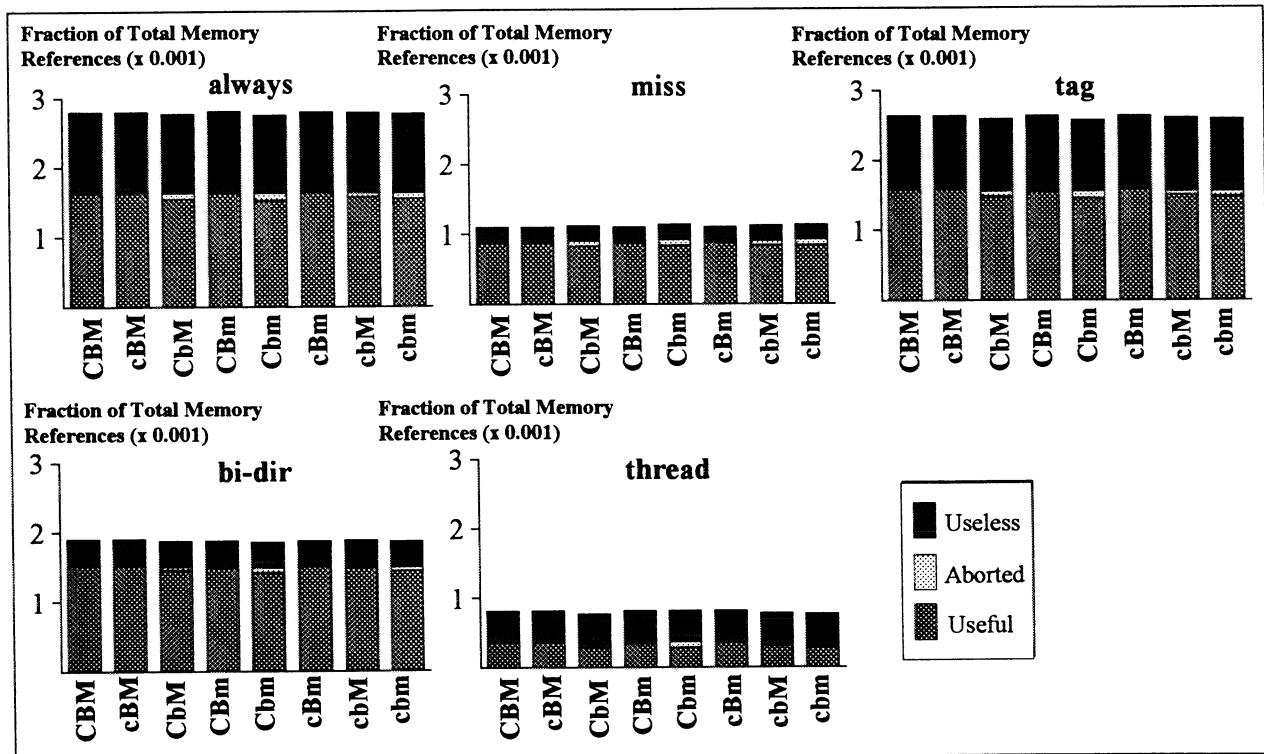


Figure A.2.1. Average distribution of prefetches for the instruction cache.



B.1. Single vs. Double Ported Cache Tag Arrays

Cache Tag Array	True Miss Ratios (Instruction Cache)							True Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
single ported	0.0011	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0109	0.0065	0.0085	0.0065	0.0065	0.0101	0.0076
double ported	0.0011	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0109	0.0065	0.0085	0.0065	0.0065	0.0101	0.0076
Cache Tag Array	Partial Miss Ratios (Instruction Cache)							Partial Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
single ported	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002
double ported	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002
Cache Tag Array	Total Miss Ratios (Instruction Cache)							Total Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
single ported	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078
double ported	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078

Table B.1.1. Instruction and data cache miss ratios.

Cache Tag Array	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
single ported	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
double ported	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
Cache Tag Array	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
single ported	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612
double ported	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612

Table B.1.2. Success ratios and global success ratios.

Bus Utilization Breakdown (%)	none		always		miss		tag		bi-dir		thread	
	single	double										
Icache read	2.50	2.50	0.91	1.89	1.98	1.99	1.86	1.87	0.65	1.99	0.76	1.61
Dcache read	9.90	9.90	2.48	4.98	7.14	7.17	4.92	4.94	1.72	5.22	4.33	8.97
Dcache write	1.78	1.78	0.88	1.82	1.78	1.79	1.80	1.81	0.59	1.84	0.85	1.78
prefetch	0.00	0.00	5.33	10.99	6.94	6.97	10.00	10.05	2.54	7.95	1.68	3.44
total	14.18	14.18	9.59	19.68	17.84	17.91	18.58	18.67	5.50	17.00	7.63	15.79

Table B.1.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none		always		miss	
	single	double	single	double	single	double
cache misses	99.27	99.27	4.85	51.56	64.96	66.88
prefetch (cache)	0.00	0.00	93.71	33.97	22.00	19.68
prefetch (bus)	0.00	0.00	1.40	13.94	12.49	12.86
prefetch (mem)	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.72	0.72	0.04	0.52	0.55	0.58

CPU Stall Reasons	tag		bi-dir		thread	
	single	double	single	double	single	double
cache misses	51.75	53.69	2.67	61.07	6.51	82.74
prefetch (cache)	34.73	32.28	96.98	30.38	92.73	8.23
prefetch (bus)	12.99	13.47	0.33	7.98	0.71	8.38
prefetch (mem)	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.53	0.56	0.02	0.57	0.04	0.65

Table B.1.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory; write buffer = write backs by the write buffer.

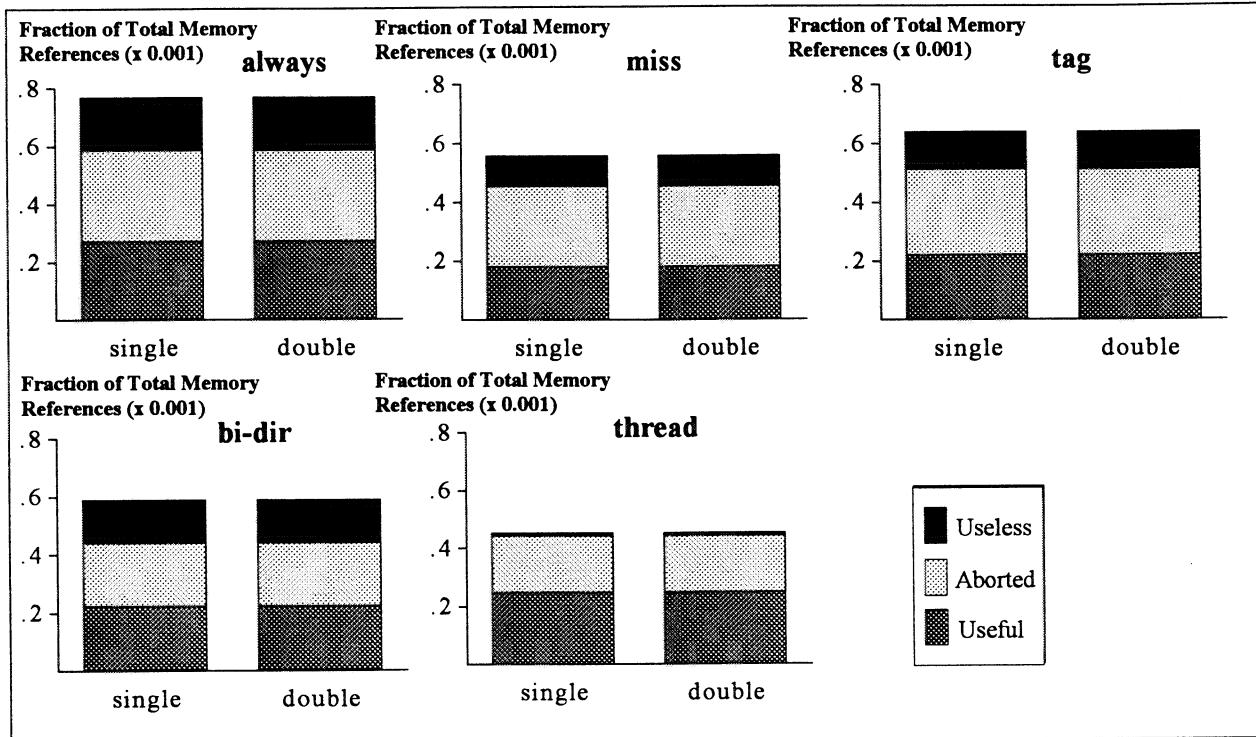


Figure B.1.1. Average distribution of prefetches for the instruction cache.

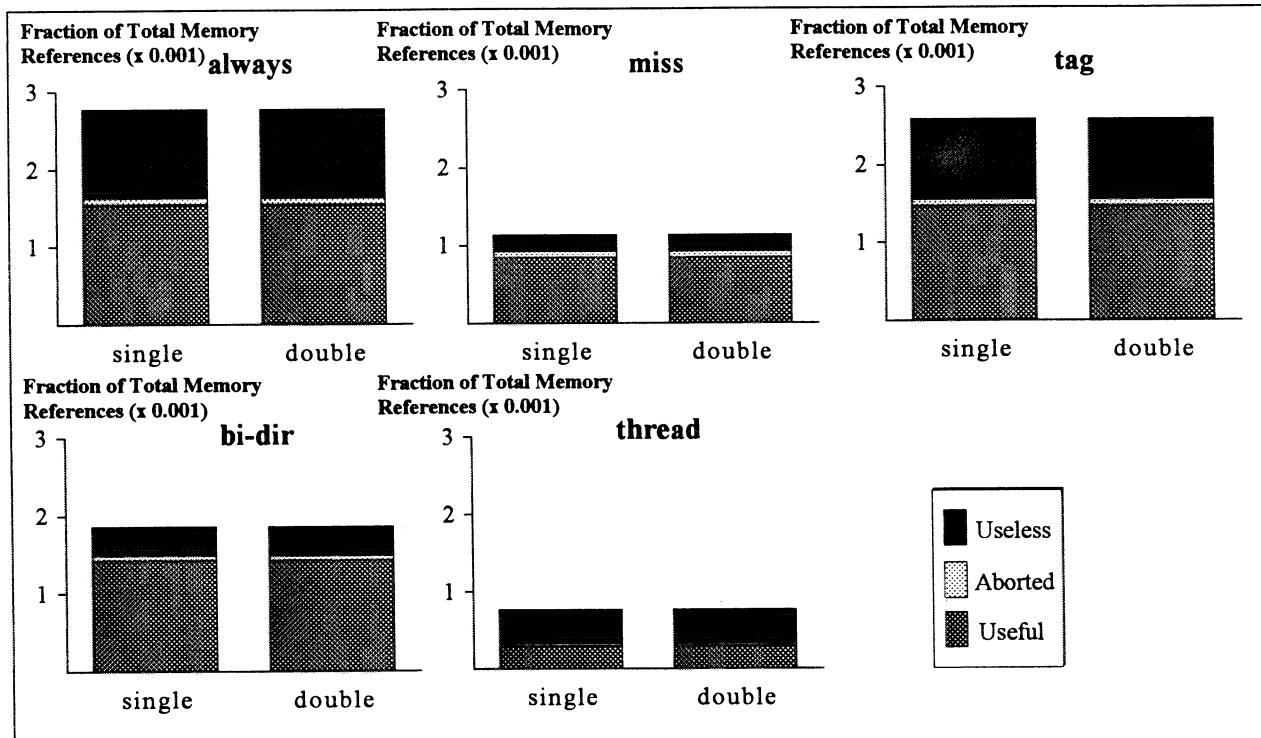


Figure B.1.2. Average distribution of prefetches for the data cache.

B.2. Single vs. Double Ported Cache Data Arrays

Cache Data Array	True Miss Ratios (Instruction Cache)							True Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
single ported	0.0011	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0109	0.0065	0.0085	0.0065	0.0065	0.0101	0.0076
double ported	0.0011	0.0008	0.0009	0.0008	0.0009	0.0007	0.0008	0.0109	0.0066	0.0085	0.0065	0.0066	0.0101	0.0077
single, buffered	0.0011	0.0008	0.0009	0.0008	0.0009	0.0007	0.0008	0.0109	0.0071	0.0086	0.0070	0.0071	0.0101	0.0080
Cache Data Array	Partial Miss Ratios (Instruction Cache)							Partial Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
single ported	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002
double ported	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002
single, buffered	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0003	0.0002	0.0003	0.0002	0.0002	0.0002
Cache Data Array	Total Miss Ratios (Instruction Cache)							Total Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
single ported	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078
double ported	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0068	0.0087	0.0067	0.0067	0.0103	0.0079
single, buffered	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0074	0.0087	0.0073	0.0073	0.0103	0.0082

Table B.2.1. Instruction and data cache miss ratios.

Cache Data Array	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
single ported	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
double ported	0.5982	0.6538	0.6475	0.6040	0.9743	0.6956	0.5923	0.5776	0.6016	0.6556	0.6397	0.6134
single, buffered	0.5956	0.6528	0.6457	0.6011	0.9743	0.6939	0.5899	0.5760	0.5996	0.6531	0.6381	0.6113
Cache Data Array	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
single ported	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612
double ported	0.2747	0.2079	0.2460	0.2438	0.2671	0.2479	0.4810	0.2801	0.4610	0.4532	0.0947	0.3540
single, buffered	0.2711	0.2067	0.2434	0.2404	0.2664	0.2456	0.4490	0.2776	0.4315	0.4236	0.0935	0.3350

Table B.2.2. Success ratios and global success ratios.

Bus Utilization Breakdown (%)	none			always			miss		
	single	double	single, buffered	single	double	single, buffered	single	double	single, buffered
Icache read	2.50	2.50	2.50	1.89	2.04	2.04	1.99	2.05	2.06
Dcache read	9.90	9.90	9.90	4.98	5.21	5.84	7.17	7.37	7.38
Dcache write	1.78	1.78	1.78	1.82	1.93	1.91	1.79	1.84	1.84
prefetch	0.00	0.00	0.00	10.99	11.34	13.04	6.97	7.08	8.59
total	14.18	14.18	14.18	19.68	20.53	22.82	17.91	18.35	19.87

Bus Utilization Breakdown (%)	tag			bi-dir			thread		
	single	double	single, buffered	single	double	single, buffered	single	double	single, buffered
Icache read	1.87	2.00	2.00	1.99	2.11	2.11	1.61	1.65	1.65
Dcache read	4.94	5.16	5.74	5.22	5.36	5.97	8.97	9.08	9.08
Dcache write	1.81	1.92	1.89	1.84	1.94	1.92	1.78	1.79	1.79
prefetch	10.05	10.39	12.04	7.95	8.30	9.82	3.44	3.43	3.88
total	18.67	19.45	21.66	17.00	17.72	19.82	15.79	15.94	16.40

Table B.2.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none			always			miss		
	single	double	single, buffered	single	double	single, buffered	single	double	single, buffered
cache misses	99.27	99.27	99.27	51.56	77.51	76.89	66.88	82.72	82.12
prefetch (cache)	0.00	0.00	0.00	33.97	0.00	0.00	19.68	0.00	0.00
prefetch (bus)	0.00	0.00	0.00	13.94	21.53	22.30	12.86	16.56	17.15
prefetch (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

CPU Stall Reasons	tag			bi-dir			thread		
	single	double	single, buffered	single	double	single, buffered	single	double	single, buffered
cache misses	53.69	78.67	77.95	61.07	86.33	85.24	82.74	89.58	88.96
prefetch (cache)	32.28	0.00	0.00	30.38	0.00	0.00	8.23	0.00	0.00
prefetch (bus)	13.47	20.31	21.18	7.98	12.65	13.90	8.38	9.71	10.34
prefetch (mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table B.2.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory.

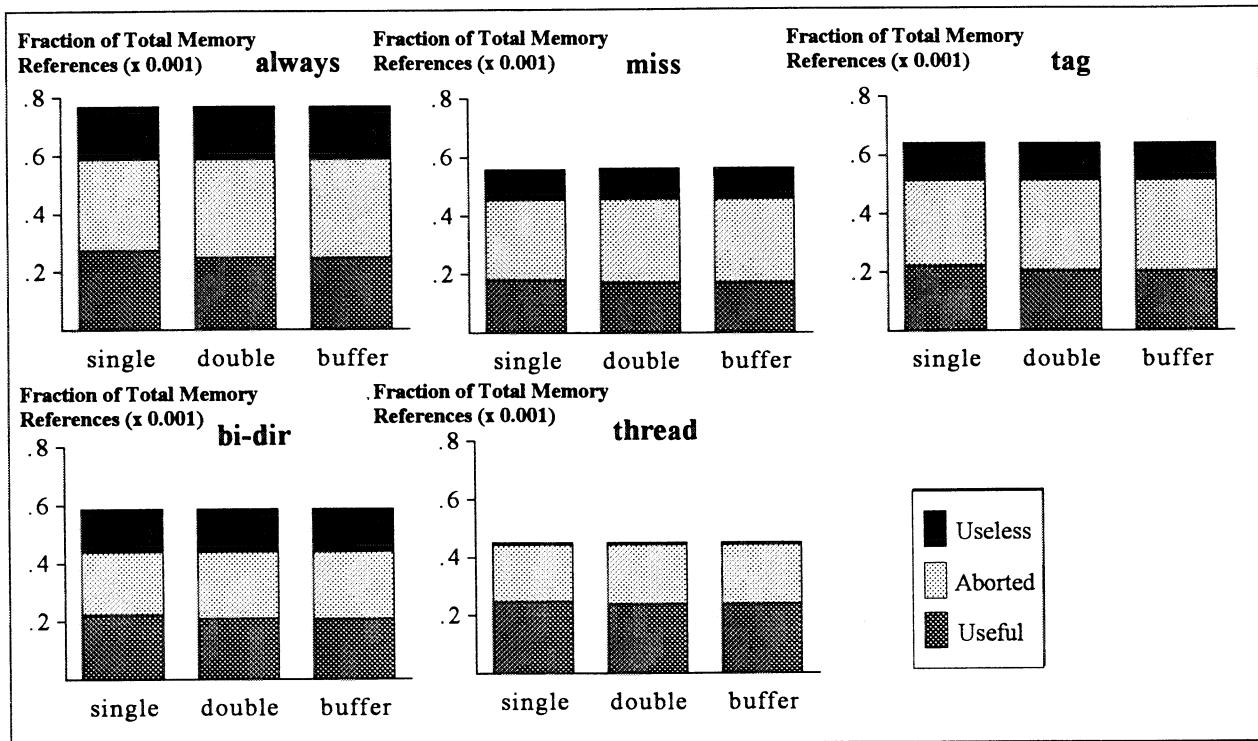


Figure B.2.1. Average distribution of prefetches for the instruction cache.

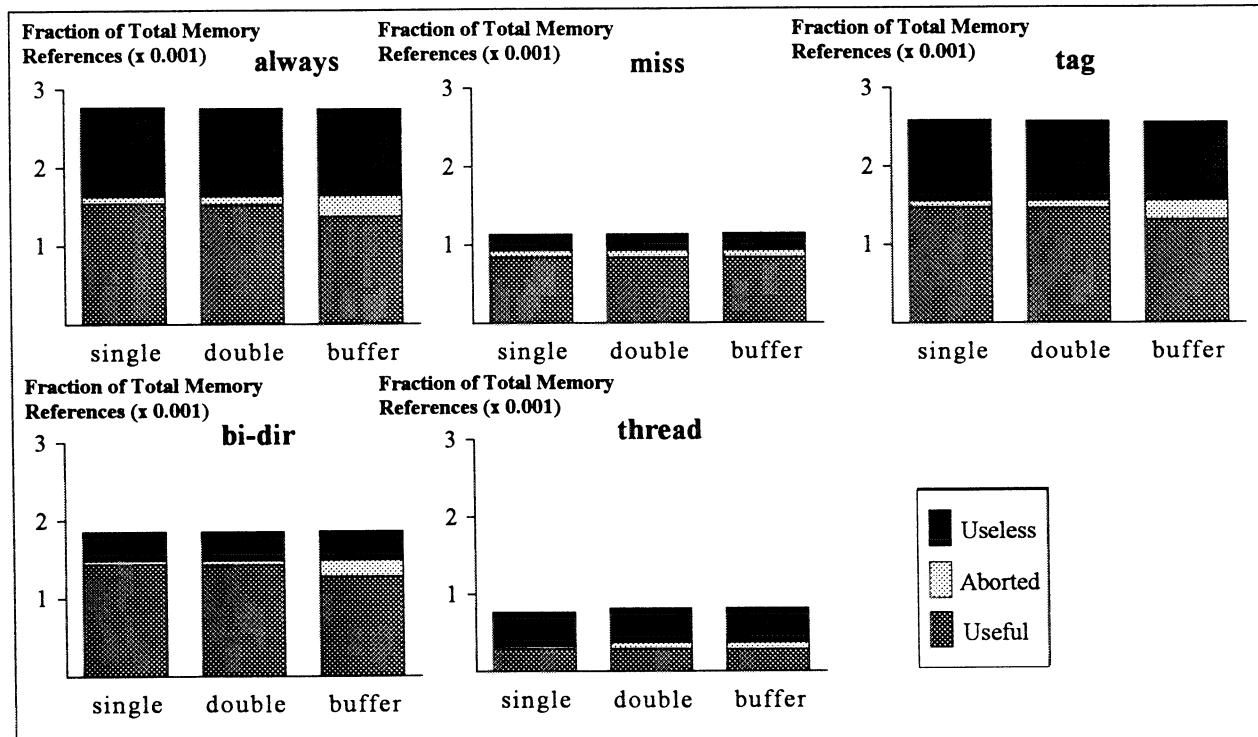


Figure B.2.2. Average distribution of prefetches for the data cache.

Cache Size	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
4K	0.4455	0.5149	0.5173	0.4722	0.8717	0.5643	0.3503	0.3595	0.3952	0.4160	0.4354	0.3913
8K	0.5728	0.6093	0.6161	0.5900	0.9311	0.6639	0.4120	0.4302	0.4616	0.4865	0.4559	0.4492
16K	0.6018	0.6211	0.6260	0.6121	0.9531	0.6828	0.4855	0.5163	0.5298	0.5654	0.5145	0.5223
32K	0.5968	0.6157	0.6242	0.5990	0.9620	0.6796	0.5597	0.5748	0.5867	0.6337	0.5952	0.5900
64K	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
128K	0.5979	0.6808	0.6788	0.5933	0.7796	0.6661	0.6292	0.6213	0.6458	0.6724	0.6669	0.6471
256K	0.5569	0.6919	0.6914	0.5422	0.5922	0.6149	0.6797	0.7011	0.7282	0.7067	0.7429	0.7117
512K	0.5299	0.7410	0.7422	0.5233	0.2000	0.5473	0.7560	0.8139	0.8535	0.7572	0.6789	0.7719
1M	0.5501	0.7650	0.7556	0.5539	0.0000	0.5249	0.7307	0.8652	0.9000	0.7298	0.5445	0.7540
Cache Size	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
4K	0.1532	0.1657	0.1909	0.1503	0.3150	0.1950	0.2613	0.1656	0.2478	0.2648	0.0991	0.2077
8K	0.2686	0.2519	0.2878	0.2456	0.4278	0.2963	0.3066	0.1949	0.2926	0.2952	0.0719	0.2322
16K	0.3438	0.2738	0.3117	0.3076	0.4137	0.3301	0.3524	0.2200	0.3380	0.3360	0.0717	0.2636
32K	0.2903	0.2056	0.2516	0.2466	0.3333	0.2655	0.4728	0.2811	0.4625	0.4403	0.0832	0.3480
64K	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612
128K	0.3204	0.2158	0.2716	0.2737	0.2529	0.2669	0.5269	0.2955	0.4925	0.5037	0.0984	0.3834
256K	0.3684	0.1975	0.2564	0.3086	0.1434	0.2549	0.6095	0.3251	0.5520	0.6085	0.0921	0.4374
512K	0.3910	0.2509	0.3250	0.3304	0.0041	0.2603	0.7171	0.3432	0.6114	0.7472	0.0649	0.4968
1M	0.4555	0.3252	0.3950	0.4047	0.0000	0.3161	0.8156	0.4140	0.7478	0.8297	0.0367	0.5688

Table B.3.2. Success ratios and global success ratios.

Bus Utilization Breakdown %	none								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
Icache read	29.35	18.70	9.04	4.67	2.50	1.19	0.34	0.22	0.22
Dcache read	18.65	16.09	14.09	11.85	9.90	8.35	6.53	4.22	3.38
Dcache write	3.34	2.58	2.38	1.91	1.78	1.53	1.39	0.87	0.59
prefetch	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
total	51.34	37.37	25.52	18.44	14.18	11.07	8.27	5.32	4.18
Bus Utilization Breakdown %	always								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
Icache read	23.63	14.51	6.31	3.49	1.89	0.99	0.32	0.18	0.14
Dcache read	13.02	10.22	8.82	6.06	4.98	3.76	2.37	0.53	0.16
Dcache write	3.55	2.69	2.40	1.91	1.82	1.57	1.48	0.94	0.61
prefetch	27.04	24.55	18.42	14.05	10.99	8.50	6.54	4.27	3.53
total	67.23	51.97	35.95	25.51	19.68	14.82	10.71	5.92	4.45
Bus Utilization Breakdown %	miss								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
Icache read	23.57	14.69	6.74	3.74	1.99	1.04	0.32	0.19	0.14
Dcache read	14.32	12.08	10.75	8.44	7.17	5.90	4.45	2.37	1.77
Dcache write	3.36	2.60	2.36	1.89	1.79	1.55	1.44	0.90	0.60
prefetch	20.77	17.85	12.48	9.34	6.97	5.18	3.34	2.10	1.78
total	62.02	47.22	32.33	23.42	17.91	13.68	9.55	5.56	4.28
Bus Utilization Breakdown %	tag								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
Icache read	22.83	14.02	6.40	3.51	1.87	0.99	0.31	0.18	0.14
Dcache read	12.30	10.16	8.81	6.07	4.94	3.81	2.45	0.56	0.18
Dcache write	3.40	2.62	2.36	1.90	1.81	1.56	1.46	0.91	0.61
prefetch	24.96	21.84	16.05	13.05	10.05	7.94	5.78	4.00	3.40
total	63.49	48.64	33.63	24.53	18.67	14.31	10.01	5.66	4.34
Bus Utilization Breakdown %	bi-dir								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
Icache read	24.33	15.25	6.71	3.73	1.99	1.03	0.32	0.19	0.15
Dcache read	12.98	10.60	9.14	6.52	5.22	3.88	2.34	0.53	0.18
Dcache write	3.52	2.66	2.40	1.95	1.84	1.57	1.47	0.94	0.61
prefetch	22.76	19.31	13.58	10.12	7.95	6.71	6.05	4.27	3.51
total	63.59	47.83	31.83	22.32	17.00	13.19	10.18	5.93	4.45
Bus Utilization Breakdown %	thread								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
Icache read	19.06	10.59	5.01	2.82	1.61	0.80	0.29	0.22	0.22
Dcache read	16.42	14.90	13.37	10.88	8.97	7.55	5.88	3.72	2.98
Dcache write	3.29	2.54	2.33	1.89	1.78	1.54	1.41	0.88	0.60
prefetch	17.68	13.28	8.09	4.85	3.44	2.38	1.23	0.63	0.50
total	56.45	41.32	28.81	20.45	15.79	12.28	8.81	5.46	4.29

Table B.3.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
cache misses	98.69	98.06	98.15	99.30	99.27	99.31	99.21	99.30	99.98
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
prefetch(bus)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	1.30	1.94	1.86	0.70	0.72	0.69	0.79	0.72	0.02
CPU Stall Reasons	always								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
cache misses	64.99	60.73	57.57	52.36	51.56	49.11	44.10	26.15	15.99
prefetch(cache)	15.29	18.80	22.70	31.52	33.97	37.81	47.54	70.05	79.01
prefetch(bus)	18.98	19.16	18.38	15.65	13.94	12.51	7.73	3.15	5.00
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.74	1.31	1.35	0.48	0.52	0.57	0.63	0.69	0.02
CPU Stall Reasons	miss								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
cache misses	71.31	69.28	68.48	66.48	66.88	67.03	68.79	63.36	59.56
prefetch(cache)	10.39	12.45	14.73	18.78	19.68	21.15	23.57	31.27	33.07
prefetch(bus)	17.38	16.74	15.28	14.20	12.86	11.22	6.97	4.66	7.42
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.92	1.53	1.51	0.53	0.58	0.59	0.67	0.69	0.02
CPU Stall Reasons	tag								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
cache misses	66.69	63.69	60.57	54.23	53.69	51.18	48.47	29.12	17.70
prefetch(cache)	14.41	17.34	20.75	30.31	32.28	36.49	44.53	67.17	77.12
prefetch(bus)	18.00	17.49	17.21	14.94	13.47	11.74	6.30	2.94	5.14
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.90	1.48	1.47	0.52	0.56	0.59	0.70	0.75	0.02
CPU Stall Reasons	bi-dir								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
cache misses	69.17	67.79	67.49	62.85	61.07	55.91	45.41	24.78	15.60
prefetch(cache)	13.67	16.74	19.65	27.62	30.38	35.10	47.41	71.81	79.84
prefetch(bus)	16.37	14.08	11.41	9.00	7.98	8.40	6.51	2.74	4.52
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.79	1.38	1.45	0.52	0.57	0.60	0.68	0.69	0.02

CPU Stall Reasons	thread								
	4K	8K	16K	32K	64K	128K	256K	512K	1M
cache misses	77.15	77.51	79.20	81.95	82.74	84.13	87.75	89.20	89.73
prefetch(cache)	8.07	7.84	7.36	8.00	8.23	8.24	8.02	8.44	8.62
prefetch(bus)	13.63	12.84	11.73	9.42	8.38	6.97	3.48	1.66	1.65
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	1.16	1.81	1.71	0.63	0.65	0.64	0.75	0.70	0.02

Table B.3.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory.

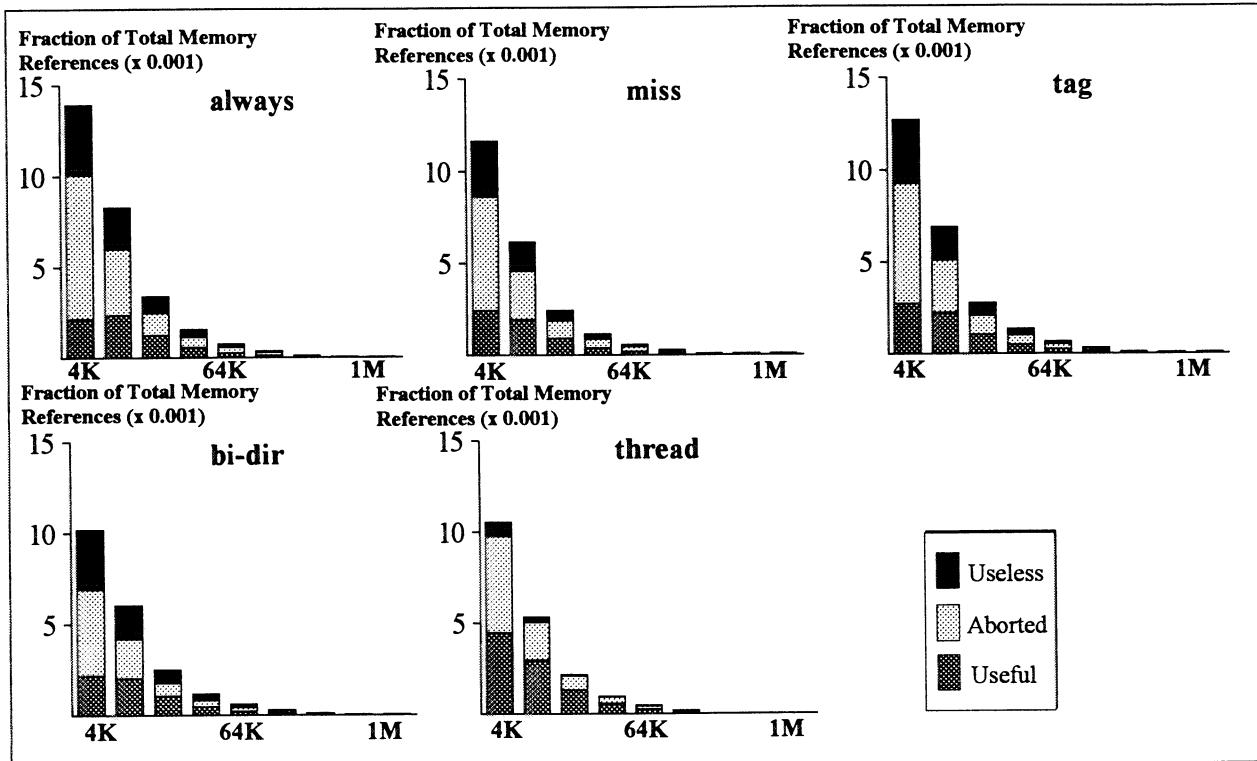


Figure B.3.1. Average distribution of prefetches for the instruction cache.

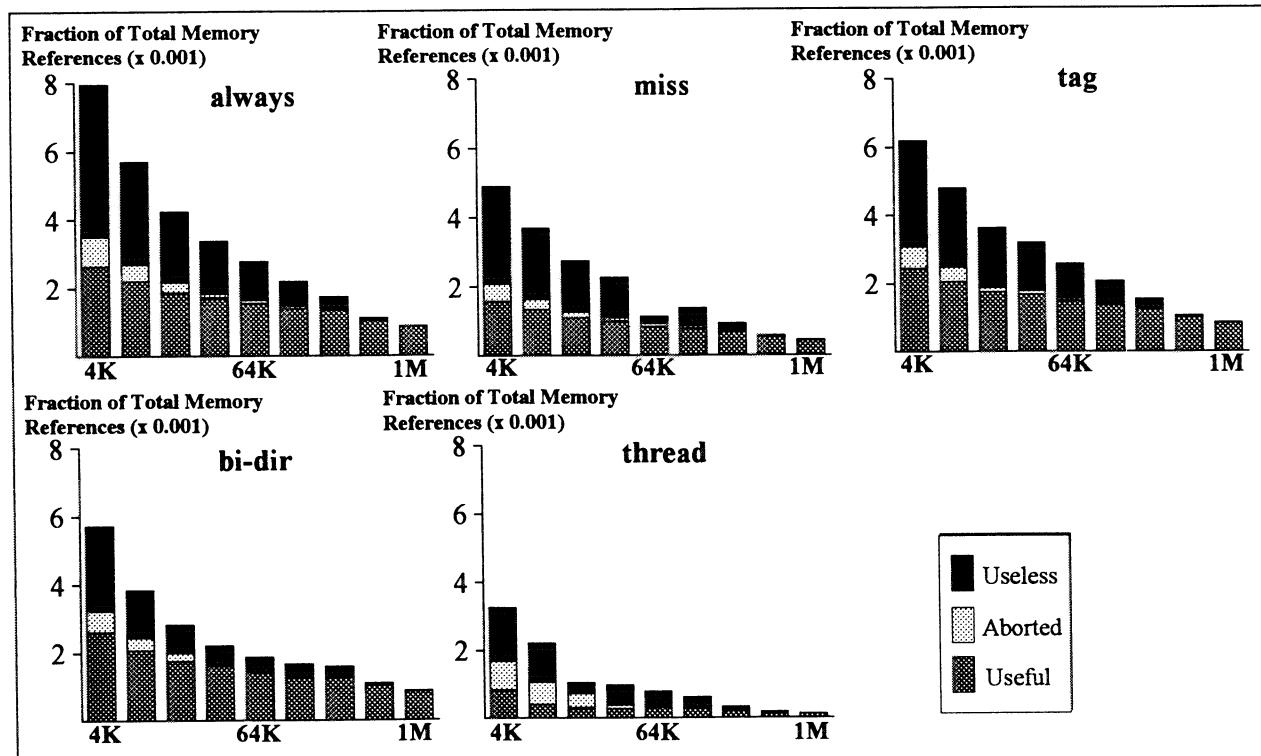


Figure B.3.2. Average distribution of prefetches for the data cache.

B.4. Cache Block Size

Cache Block Size (in bytes)	True Miss Ratios (Instruction Cache)							True Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
16	0.0024	0.0014	0.0015	0.0014	0.0015	0.0012	0.0014	0.0251	0.0107	0.0150	0.0108	0.0111	0.0208	0.0137
32	0.0015	0.0010	0.0011	0.0010	0.0011	0.0009	0.0010	0.0157	0.0070	0.0109	0.0070	0.0073	0.0141	0.0093
64	0.0011	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0109	0.0065	0.0085	0.0065	0.0065	0.0101	0.0076
128	0.0009	0.0007	0.0007	0.0007	0.0007	0.0006	0.0007	0.0086	0.0059	0.0073	0.0061	0.0058	0.0083	0.0067
256	0.0008	0.0007	0.0007	0.0007	0.0007	0.0006	0.0007	0.0072	0.0056	0.0064	0.0057	0.0055	0.0070	0.0060
Cache Block Size (in bytes)	Partial Miss Ratios (Instruction Cache)							Partial Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
16	0.0000	0.0009	0.0008	0.0008	0.0009	0.0010	0.0009	0.0000	0.0053	0.0033	0.0053	0.0051	0.0024	0.0043
32	0.0000	0.0004	0.0003	0.0003	0.0004	0.0004	0.0004	0.0000	0.0004	0.0004	0.0004	0.0003	0.0005	0.0004
64	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002
128	0.0000	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0000	0.0000	0.0001	0.0000	0.0000	0.0001	0.0001
256	0.0000	0.0001	0.0000	0.0001	0.0001	0.0001	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Cache Block Size (in bytes)	Total Miss Ratios (Instruction Cache)							Total Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
16	0.0024	0.0023	0.0023	0.0023	0.0023	0.0022	0.0023	0.0251	0.0160	0.0183	0.0161	0.0163	0.0232	0.0180
32	0.0015	0.0014	0.0014	0.0014	0.0014	0.0013	0.0014	0.0157	0.0074	0.0113	0.0074	0.0076	0.0145	0.0097
64	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078
128	0.0009	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0086	0.0060	0.0074	0.0061	0.0059	0.0083	0.0067
256	0.0008	0.0008	0.0008	0.0007	0.0008	0.0006	0.0007	0.0072	0.0056	0.0064	0.0057	0.0055	0.0070	0.0061

Table B.4.1. Instruction and data cache miss ratios.

Cache Block Size (in bytes)	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
16	0.7508	0.8052	0.7973	0.7583	0.7917	0.7807	0.6547	0.6253	0.6628	0.7499	0.7194	0.6824
32	0.6727	0.7323	0.7272	0.6769	0.7846	0.7188	0.6175	0.6060	0.6275	0.6867	0.6624	0.6400
64	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
128	0.5600	0.6106	0.6035	0.5527	0.9592	0.6572	0.6004	0.5621	0.6112	0.6452	0.6212	0.6080
256	0.5182	0.5201	0.5312	0.5018	0.9360	0.6015	0.5539	0.5312	0.5812	0.5717	0.5984	0.5673
Cache Block Size (in bytes)	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
16	0.4436	0.4005	0.4284	0.4240	0.4058	0.4205	0.5122	0.3392	0.5014	0.4722	0.1453	0.3941
32	0.3521	0.2944	0.3352	0.3260	0.3295	0.3274	0.5168	0.3084	0.5017	0.4752	0.1131	0.3830
64	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612
128	0.2696	0.1749	0.2225	0.2231	0.2260	0.2232	0.4759	0.2545	0.4483	0.4576	0.0764	0.3425
256	0.2348	0.1334	0.1901	0.1790	0.1988	0.1872	0.4087	0.2227	0.3875	0.4035	0.0573	0.2959

Table B.4.2. Success ratios and global success ratios.

Bus Utilization Breakdown %	none					always				
	16	32	64	128	256	16	32	64	128	256
Icache read	3.95	3.05	2.50	2.66	3.33	2.36	2.12	1.89	2.07	2.73
Dcache read	17.70	13.15	9.90	10.93	13.23	6.49	4.72	4.98	6.16	8.93
Dcache write	1.05	1.39	1.78	2.38	3.31	1.16	1.50	1.82	2.32	3.20
prefetch	0.00	0.00	0.00	0.00	0.00	15.97	13.70	10.99	10.95	9.62
total	22.71	17.58	14.18	15.98	19.87	25.97	22.03	19.68	21.51	24.48
Bus Utilization Breakdown %	miss					tag				
	16	32	64	128	256	16	32	64	128	256
Icache read	2.43	2.21	1.99	2.23	2.92	2.37	2.12	1.87	2.08	2.71
Dcache read	8.97	8.18	7.17	8.36	10.85	6.55	4.73	4.94	6.36	9.07
Dcache write	1.12	1.44	1.79	2.33	3.22	1.16	1.49	1.81	2.32	3.20
prefetch	11.79	8.37	6.97	7.11	5.81	15.52	13.08	10.05	9.75	8.17
total	24.30	20.20	17.91	20.04	22.80	25.59	21.42	18.67	20.52	23.14
Bus Utilization Breakdown %	bi-dir					thread				
	16	32	64	128	256	16	32	64	128	256
Icache read	2.44	2.20	1.99	2.25	2.94	1.95	1.77	1.61	1.82	2.42
Dcache read	7.77	5.52	5.22	6.44	8.86	14.04	11.58	8.97	10.08	12.51
Dcache write	1.16	1.51	1.84	2.38	3.23	1.07	1.40	1.78	2.34	3.25
prefetch	13.56	11.04	7.95	7.74	8.37	6.09	4.03	3.44	3.68	2.91
total	24.94	20.27	17.00	18.81	23.39	23.15	18.78	15.79	17.92	21.09

Table B.4.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none					always				
	16	32	64	128	256	16	32	64	128	256
cache misses	89.53	99.57	99.27	98.56	95.65	62.09	59.89	51.56	45.80	49.67
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	4.77	23.73	33.97	43.00	43.03
prefetch(bus)	0.00	0.00	0.00	0.00	0.00	16.18	15.92	13.94	10.48	5.37
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	10.47	0.43	0.72	1.43	4.35	16.97	0.47	0.52	0.71	1.92
CPU Stall Reasons	miss					tag				
	16	32	64	128	256	16	32	64	128	256
cache misses	67.90	74.42	66.88	61.72	64.02	62.68	61.06	53.69	49.34	53.48
prefetch(cache)	2.98	10.48	19.68	27.59	29.04	4.54	22.95	32.28	39.98	39.57
prefetch(bus)	13.31	14.69	12.86	9.80	4.53	15.75	15.49	13.47	9.90	4.83
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	15.80	0.41	0.58	0.87	2.40	17.04	0.50	0.56	0.77	2.12
CPU Stall Reasons	bi-dir					thread				
	16	32	64	128	256	16	32	64	128	256
cache misses	69.65	68.89	61.07	55.01	52.58	80.72	87.71	82.74	78.08	79.00
prefetch(cache)	4.69	22.50	30.38	37.91	40.86	1.07	3.53	8.23	13.12	13.97
prefetch(bus)	9.77	8.11	7.98	6.23	4.31	6.22	8.37	8.38	7.65	3.80
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	15.89	0.49	0.57	0.85	2.26	11.99	0.39	0.65	1.16	3.24

Table B.4.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory; write buffer = write backs by the write buffer.

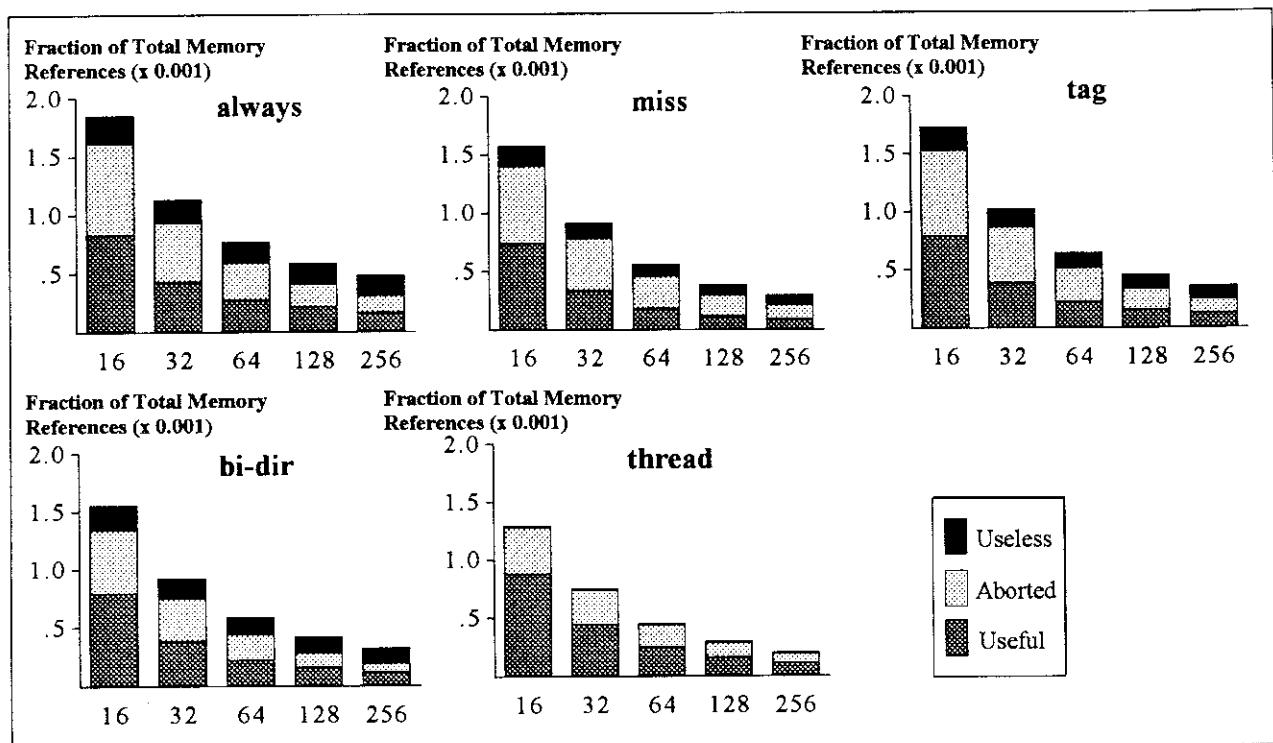


Figure B.4.1. Average distribution of prefetches for the instruction cache.

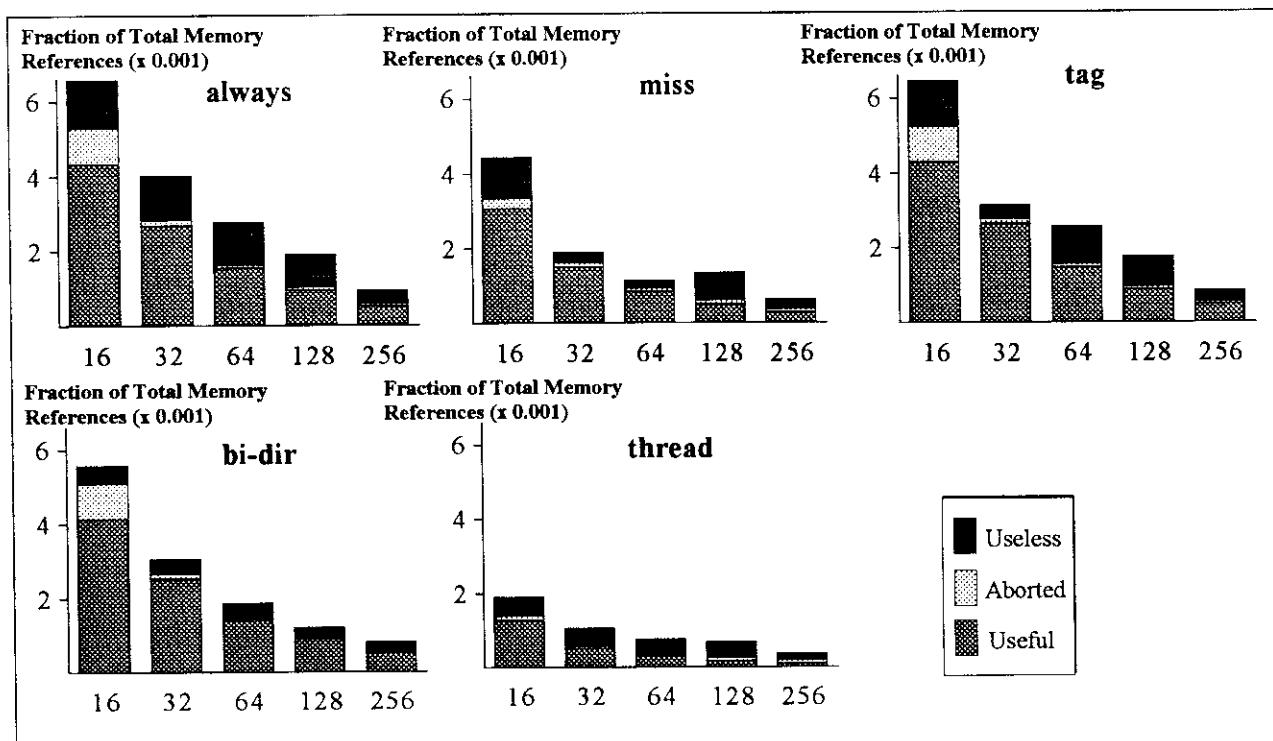


Figure B.4.2. Average distribution of prefetches for the data cache.

Bus Utilization Breakdown %	none					always				
	1	2	4	8	16	1	2	4	8	16
Icache read	5.56	3.34	2.64	2.50	2.43	3.76	2.39	1.96	1.89	1.82
Dcache read	13.28	10.06	9.89	9.90	9.81	8.87	5.18	4.98	4.98	4.97
Dcache write	2.49	1.76	1.76	1.78	1.79	2.80	1.81	1.81	1.82	1.85
prefetch	0.00	0.00	0.00	0.00	0.00	15.63	11.82	11.13	10.99	10.87
total	21.33	15.15	14.29	14.18	14.03	31.06	21.19	19.88	19.68	19.51
Bus Utilization Breakdown %	miss					tag				
	1	2	4	8	16	1	2	4	8	16
Icache read	4.14	2.64	2.12	1.99	1.92	3.91	2.47	2.00	1.87	1.80
Dcache read	10.75	7.35	7.18	7.17	7.10	8.57	5.21	4.99	4.94	4.85
Dcache write	2.67	1.77	1.77	1.79	1.81	2.71	1.79	1.79	1.81	1.84
prefetch	10.55	7.60	7.05	6.97	6.87	13.93	10.69	10.07	10.05	9.91
total	28.11	19.36	18.12	17.91	17.69	29.12	20.17	18.85	18.67	18.39
Bus Utilization Breakdown %	bi-dir					thread				
	1	2	4	8	16	1	2	4	8	16
Icache read	4.23	2.60	2.11	1.99	1.91	2.84	1.91	1.65	1.61	1.61
Dcache read	8.81	5.38	5.23	5.22	5.19	12.04	9.08	8.96	8.97	8.90
Dcache write	2.66	1.83	1.83	1.84	1.87	2.54	1.75	1.75	1.78	1.79
prefetch	11.98	8.60	8.05	7.95	7.85	6.98	4.13	3.56	3.44	3.36
total	27.68	18.42	17.22	17.00	16.81	24.41	16.87	15.93	15.79	15.64

Table B.5.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none					always				
	1	2	4	8	16	1	2	4	8	16
cache misses	99.56	99.33	99.30	99.27	99.24	59.09	52.49	51.66	51.56	51.50
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	26.65	32.79	33.87	33.97	34.08
prefetch(bus)	0.00	0.00	0.00	0.00	0.00	13.96	14.23	13.96	13.94	13.90
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.44	0.67	0.70	0.72	0.76	0.31	0.50	0.50	0.52	0.52
CPU Stall Reasons	miss					tag				
	1	2	4	8	16	1	2	4	8	16
cache misses	70.30	67.25	67.03	66.88	66.72	61.33	55.06	54.19	53.69	53.41
prefetch(cache)	17.25	19.54	19.71	19.68	19.62	25.30	30.93	31.83	32.28	32.29
prefetch(bus)	12.11	12.68	12.69	12.86	13.07	13.05	13.48	13.42	13.47	13.72
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.33	0.53	0.56	0.58	0.60	0.33	0.53	0.56	0.56	0.57
CPU Stall Reasons	bi-dir					thread				
	1	2	4	8	16	1	2	4	8	16
cache misses	66.46	61.79	61.27	61.07	61.00	81.39	81.68	82.54	82.74	82.70
prefetch(cache)	24.57	29.17	30.07	30.38	30.48	9.50	9.18	8.49	8.23	7.99
prefetch(bus)	8.62	8.52	8.10	7.98	7.94	8.71	8.53	8.34	8.38	8.64
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.34	0.53	0.55	0.57	0.58	0.39	0.60	0.63	0.65	0.68

Table B.5.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory; write buffer = write backs by the write buffer.

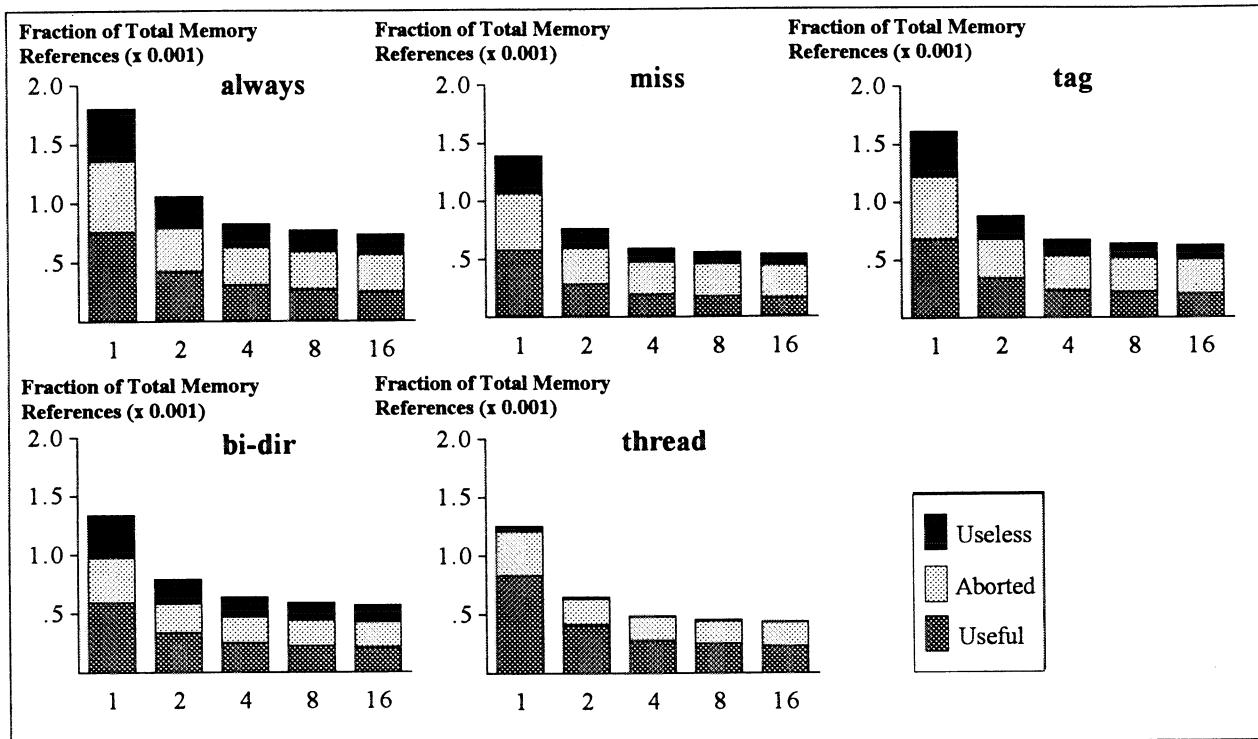


Figure B.5.1. Average distribution of prefetches for the instruction cache.

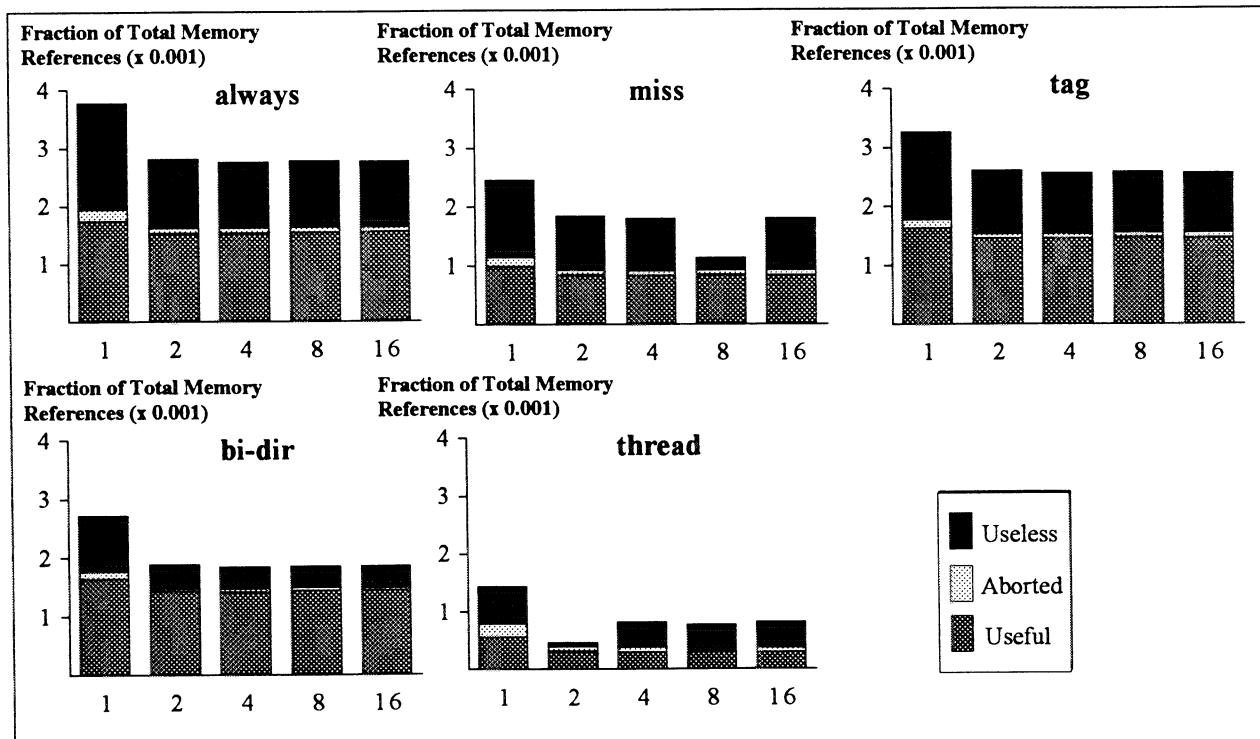


Figure B.5.2. Average distribution of prefetches for the data cache.

B.6. Split vs Non-Split Bus Transaction

Bus Transaction	True Miss Ratios (Instruction Cache)							True Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
non-split	0.0011	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0109	0.0063	0.0085	0.0065	0.0065	0.0101	0.0076
split	0.0011	0.0004	0.0006	0.0004	0.0007	0.0004	0.0005	0.0109	0.0062	0.0084	0.0062	0.0063	0.0098	0.0074
Bus Transaction	Partial Miss Ratios (Instruction Cache)							Partial Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
non-split	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002
split	0.0000	0.0003	0.0002	0.0003	0.0003	0.0003	0.0003	0.0000	0.0001	0.0001	0.0001	0.0001	0.0003	0.0001
Bus Transaction	Total Miss Ratios (Instruction Cache)							Total Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
non-split	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078
split	0.0011	0.0008	0.0008	0.0007	0.0009	0.0007	0.0008	0.0109	0.0064	0.0085	0.0063	0.0064	0.0101	0.0075

Table B.6.1. Instruction and data cache miss ratios.

Bus Transaction	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
non-split	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
split	0.7434	0.8066	0.7862	0.7019	0.9839	0.8044	0.6078	0.5867	0.6128	0.6722	0.6904	0.6340
Bus Transaction	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
non-split	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612
split	0.6160	0.4191	0.5731	0.4300	0.4531	0.4983	0.5373	0.3051	0.5192	0.4914	0.1262	0.3958

Table B.6.2. Success ratios and global success ratios.

Bus Utilization Breakdown (%)	none		always		miss		tag		bi-dir		thread	
	non-sp	split										
Icache read	2.50	1.29	1.89	0.56	1.99	0.75	1.87	0.57	1.99	0.81	1.61	0.55
Dcache read	9.90	5.10	4.98	2.48	7.17	3.71	4.94	2.46	5.22	2.59	8.97	4.57
Dcache write	1.78	1.78	1.82	1.86	1.79	1.82	1.81	1.84	1.84	1.86	1.78	1.80
prefetch	0.00	0.00	10.99	6.41	6.97	4.01	10.05	5.91	7.95	4.46	3.44	2.21
total	14.18	8.17	19.68	11.31	17.91	10.29	18.67	10.79	17.00	9.73	15.79	9.12

Table B.6.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none		always		miss	
	non-split	split	non-split	split	non-split	split
cache misses	99.27	100.00	51.56	51.75	66.88	72.18
prefetch (cache)	0.00	0.00	33.97	40.39	19.68	21.93
prefetch (bus)	0.00	0.00	13.94	7.77	12.86	5.85
prefetch (mem)	0.00	0.00	0.00	0.10	0.00	0.04
write buffer	0.72	0.00	0.52	0.00	0.58	0.00

CPU Stall Reasons	tag		bi-dir		thread	
	non-split	split	non-split	split	non-split	split
cache misses	53.69	53.75	61.07	61.99	82.74	85.88
prefetch (cache)	32.28	38.59	30.38	33.67	8.23	9.71
prefetch (bus)	13.47	7.58	7.98	4.27	8.38	4.37
prefetch (mem)	0.00	0.09	0.00	0.07	0.00	0.03
write buffer	0.56	0.00	0.57	0.00	0.65	0.00

Table B.6.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory; write buffer = write backs by the write buffer.

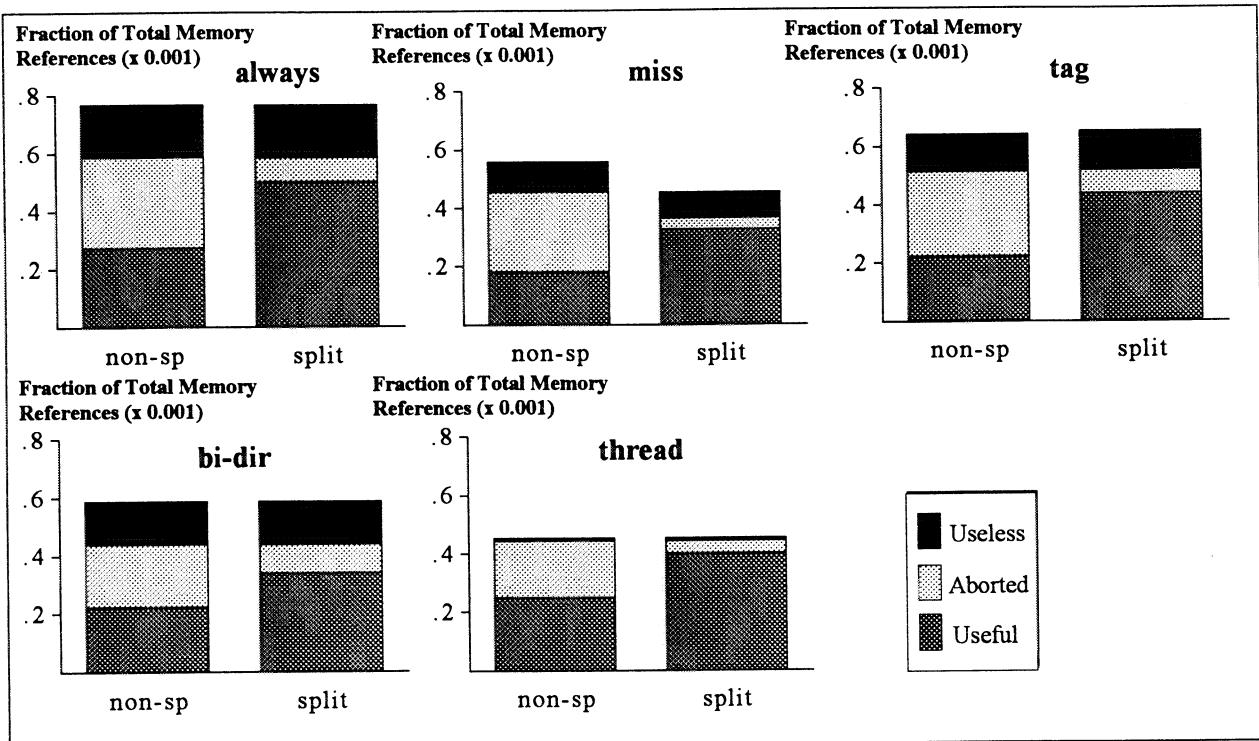


Figure B.6.1. Average distribution of prefetches for the instruction cache.

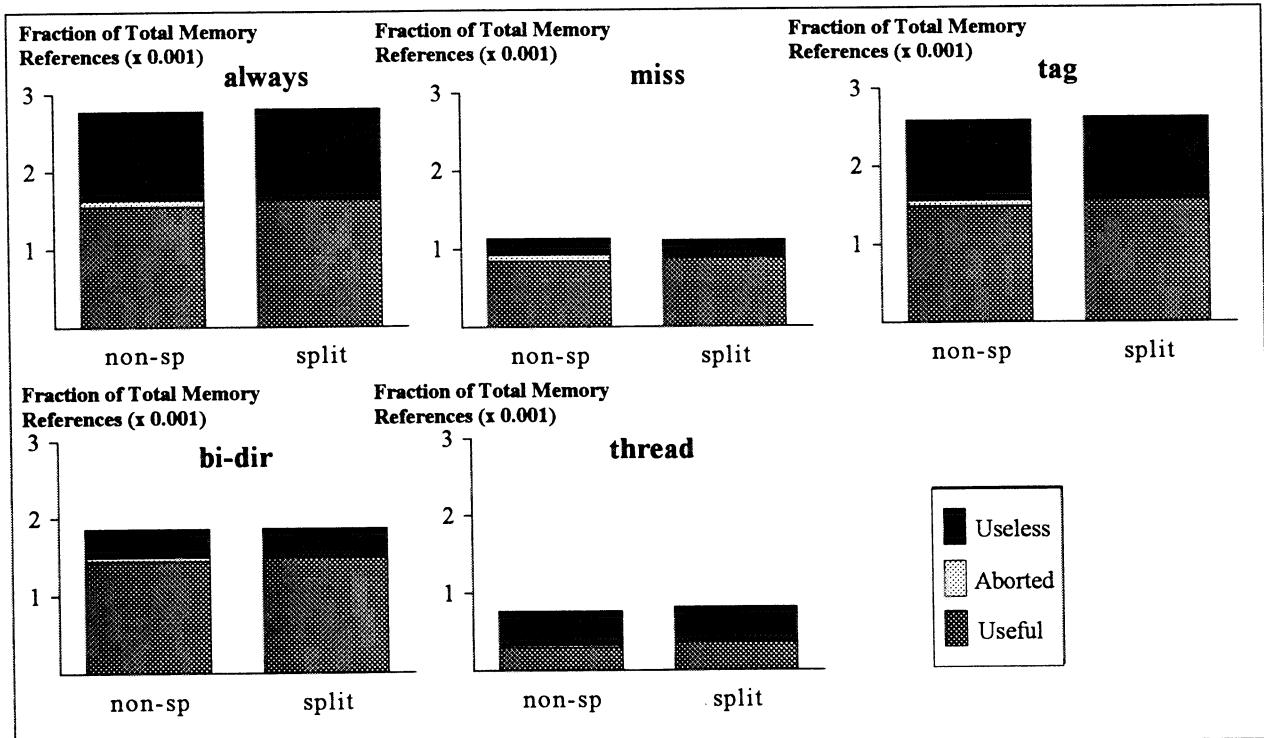


Figure B.6.2. Average distribution of prefetches for the data cache.

B.7. Bus Width

Bus Width (in bytes)	True Miss Ratios (Instruction Cache)							True Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
4	0.0011	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0109	0.0065	0.0085	0.0065	0.0065	0.0101	0.0076
8	0.0010	0.0006	0.0006	0.0006	0.0006	0.0005	0.0006	0.0109	0.0063	0.0084	0.0062	0.0063	0.0098	0.0074
16	0.0009	0.0005	0.0005	0.0005	0.0005	0.0004	0.0005	0.0109	0.0062	0.0083	0.0061	0.0063	0.0097	0.0073
Bus Width (in bytes)	Partial Miss Ratios (Instruction Cache)							Partial Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
4	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002
8	0.0000	0.0003	0.0003	0.0003	0.0003	0.0003	0.0003	0.0000	0.0002	0.0002	0.0002	0.0001	0.0004	0.0002
16	0.0000	0.0004	0.0003	0.0003	0.0003	0.0003	0.0003	0.0000	0.0003	0.0002	0.0003	0.0001	0.0004	0.0003
Bus Width (in bytes)	Total Miss Ratios (Instruction Cache)							Total Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
4	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078
8	0.0010	0.0009	0.0009	0.0008	0.0009	0.0008	0.0009	0.0109	0.0065	0.0086	0.0064	0.0064	0.0102	0.0076
16	0.0009	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0109	0.0065	0.0085	0.0064	0.0064	0.0102	0.0076

Table B.7.1. Instruction and data cache miss ratios.

Bus Width (in bytes)	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
4	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
8	0.6872	0.7558	0.7327	0.6949	0.9813	0.7704	0.6046	0.5874	0.6102	0.6710	0.6772	0.6301
16	0.7219	0.7957	0.7685	0.7197	0.9840	0.7980	0.6060	0.5883	0.6106	0.6718	0.6890	0.6332
Bus Width (in bytes)	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
4	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612
8	0.4788	0.3977	0.4450	0.4376	0.4149	0.4348	0.5283	0.3066	0.5083	0.4892	0.1170	0.3899
16	0.5734	0.4980	0.5426	0.5005	0.4793	0.5188	0.5362	0.3130	0.5166	0.4921	0.1257	0.3967

Table B.7.2. Success ratios and global success ratios.

Bus Utilization Breakdown (%)	none			always			miss		
	4	8	16	4	8	16	4	8	16
Icache read	2.50	2.20	2.11	1.89	1.22	0.89	1.99	1.30	0.96
Dcache read	9.90	9.44	9.28	4.98	4.21	3.66	7.17	6.20	5.39
Dcache write	1.78	0.90	0.45	1.82	0.96	0.49	1.79	0.92	0.47
prefetch	0.00	0.00	0.00	10.99	10.37	9.91	6.97	6.50	6.27
total	14.18	12.54	11.84	19.68	16.75	14.96	17.91	14.93	13.08
Bus Utilization Breakdown (%)	tag			bi-dir			thread		
	4	8	16	4	8	16	4	8	16
Icache read	1.87	1.21	0.89	1.99	1.35	1.10	1.61	1.05	0.81
Dcache read	4.94	4.16	3.62	5.22	4.81	4.66	8.97	8.24	7.85
Dcache write	1.81	0.95	0.48	1.84	0.96	0.49	1.78	0.91	0.46
prefetch	10.05	9.56	9.16	7.95	7.66	7.38	3.44	3.34	3.23
total	18.67	15.88	14.15	17.00	14.79	13.63	15.79	13.54	12.35

Table B.7.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none			always			miss		
	4	8	16	4	8	16	4	8	16
cache misses	99.27	99.89	94.44	51.56	58.42	61.83	66.88	73.24	72.72
prefetch(cache)	0.00	0.00	0.00	33.97	23.27	12.36	19.68	10.94	4.61
prefetch(bus)	0.00	0.00	0.00	13.94	18.21	17.80	12.86	15.73	14.70
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.72	0.10	5.55	0.52	0.09	8.01	0.58	0.09	7.97
CPU Stall Reasons	tag			bi-dir			thread		
	4	8	16	4	8	16	4	8	16
cache misses	53.69	60.05	62.88	61.07	68.01	71.89	82.74	86.30	83.69
prefetch(cache)	32.28	22.14	11.73	30.38	21.82	12.23	8.23	4.03	1.61
prefetch(bus)	13.47	17.72	17.30	7.98	10.07	9.57	8.38	9.57	8.24
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.56	0.10	8.10	0.57	0.10	6.32	0.65	0.11	6.46

Table B.7.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory; write buffer = write backs by the write buffer.

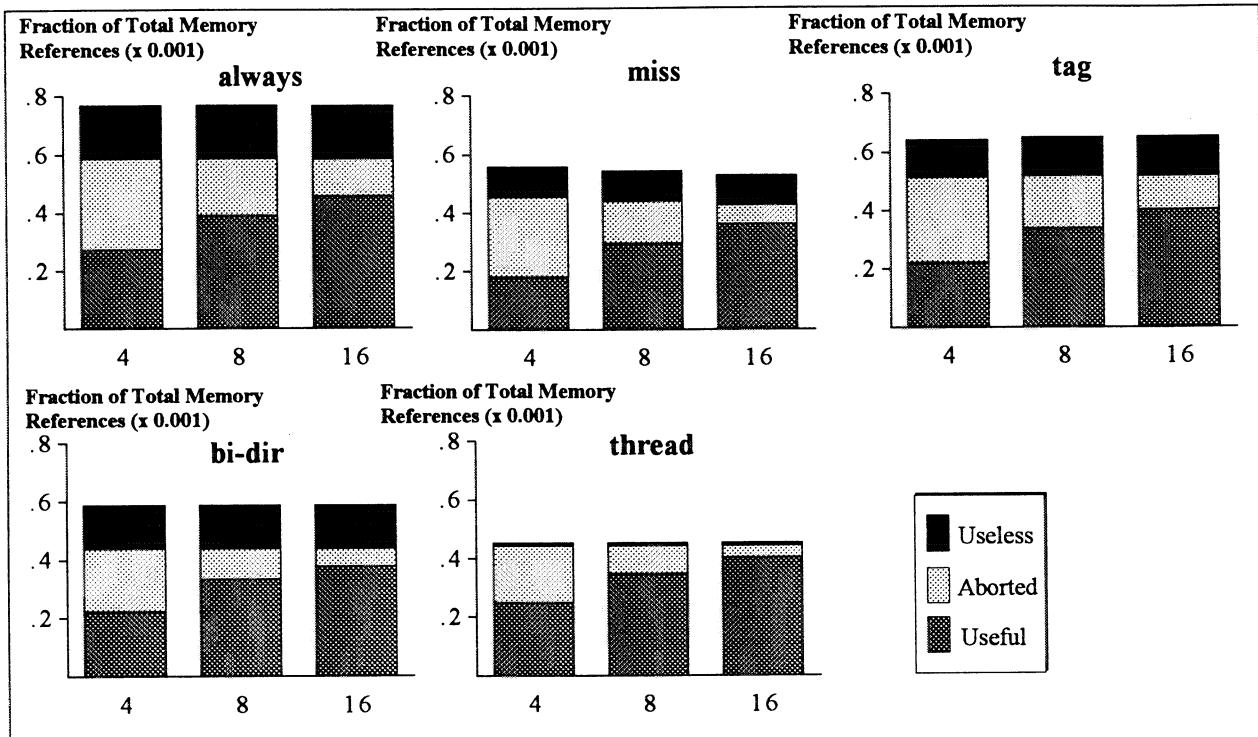


Figure B.7.1. Average distribution of prefetches for the instruction cache.

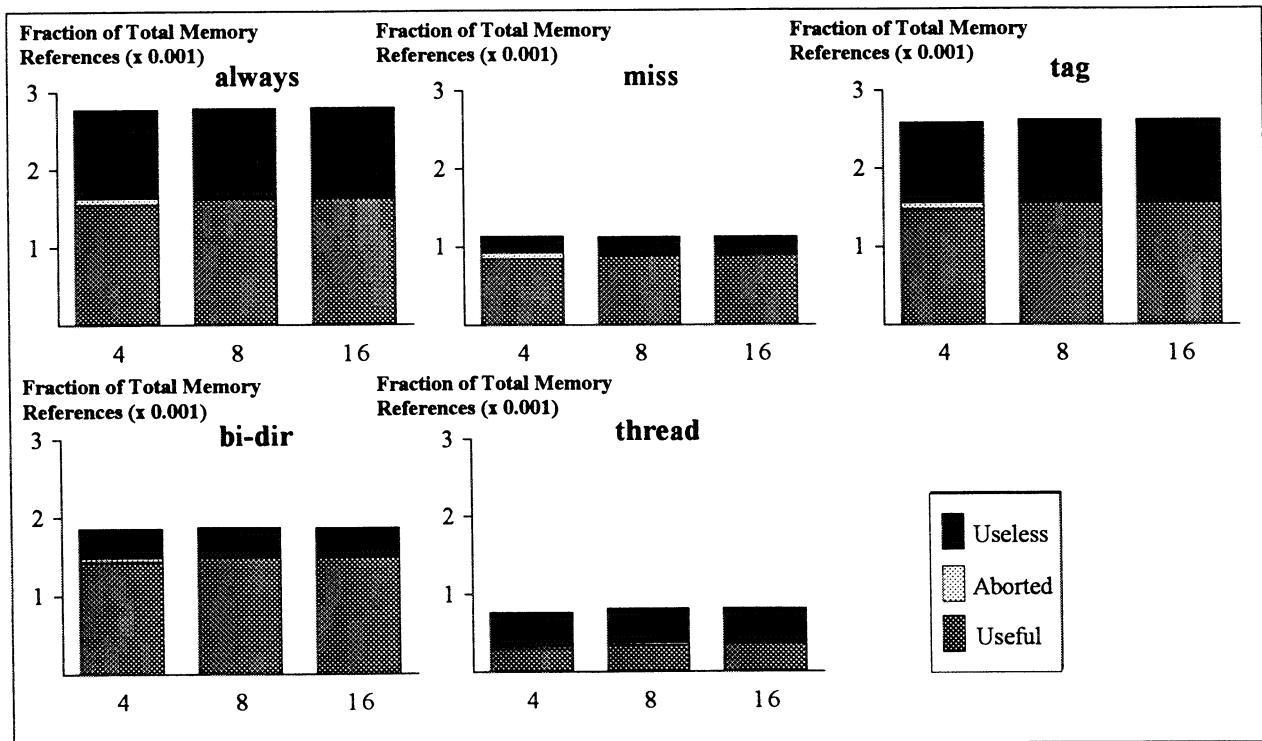


Figure B.7.2. Average distribution of prefetches for the data cache.

Bus Utilization Breakdown %	none					always				
	8	16	32	64	128	8	16	32	64	128
Icache read	1.95	2.50	3.75	5.82	8.89	1.39	1.89	2.91	4.58	7.04
Dcache read	7.81	9.90	16.18	25.54	36.57	3.88	4.98	7.57	11.95	17.46
Dcache write	1.86	1.78	1.64	1.43	1.14	1.86	1.82	1.74	1.55	1.20
prefetch	0.00	0.00	0.00	0.00	0.00	8.66	10.99	16.66	23.42	29.01
total	11.62	14.18	21.57	32.79	46.59	15.80	19.68	28.89	41.50	54.72
Bus Utilization Breakdown %	miss					tag				
	8	16	32	64	128	8	16	32	64	128
Icache read	1.52	1.99	2.95	4.53	6.92	1.39	1.87	2.86	4.48	6.90
Dcache read	5.63	7.17	10.89	15.58	19.40	3.85	4.94	7.53	11.89	17.27
Dcache write	1.85	1.79	1.67	1.47	1.19	1.85	1.81	1.73	1.54	1.21
prefetch	5.47	6.97	10.51	15.92	23.30	7.90	10.05	15.32	21.72	27.08
total	14.48	17.91	26.02	37.51	50.81	14.99	18.67	27.44	39.64	52.45
Bus Utilization Breakdown %	bi-dir					thread				
	8	16	32	64	128	8	16	32	64	128
Icache read	1.49	1.99	3.04	4.75	7.25	1.22	1.61	2.41	3.72	5.72
Dcache read	4.03	5.22	8.53	14.03	20.48	7.05	8.97	14.46	22.46	31.72
Dcache write	1.88	1.84	1.77	1.59	1.25	1.85	1.78	1.64	1.42	1.11
prefetch	6.23	7.95	12.55	18.53	23.68	2.76	3.44	5.03	7.26	9.90
total	13.63	17.00	25.89	38.89	52.66	12.88	15.79	23.54	34.86	48.45

Table B.8.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none					always				
	8	16	32	64	128	8	16	32	64	128
cache misses	99.02	99.27	99.30	99.26	99.22	44.08	51.56	59.87	68.61	73.73
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	44.27	33.97	21.88	10.28	3.38
prefetch(bus)	0.00	0.00	0.00	0.00	0.00	11.11	13.94	17.59	20.37	22.18
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.99	0.72	0.70	0.74	0.78	0.53	0.52	0.66	0.74	0.71
CPU Stall Reasons	miss					tag				
	8	16	32	64	128	8	16	32	64	128
cache misses	59.58	66.88	73.08	76.01	76.48	46.09	53.69	61.66	69.97	75.10
prefetch(cache)	28.80	19.68	10.97	5.24	2.34	42.50	32.28	20.89	9.78	3.25
prefetch(bus)	10.97	12.86	15.32	18.05	20.45	10.82	13.47	16.74	19.45	20.89
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.67	0.58	0.63	0.69	0.75	0.59	0.56	0.71	0.79	0.77
CPU Stall Reasons	bi-dir					thread				
	8	16	32	64	128	8	16	32	64	128
cache misses	53.81	61.07	68.72	76.67	80.76	77.18	82.74	86.05	87.29	87.16
prefetch(cache)	39.04	30.38	19.95	9.45	3.22	14.33	8.23	4.01	1.67	0.56
prefetch(bus)	6.54	7.98	10.63	13.09	15.30	7.67	8.38	9.30	10.34	11.53
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.60	0.57	0.70	0.77	0.73	0.83	0.65	0.65	0.70	0.74

Table B.8.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory; write buffer = write backs by the write buffer.

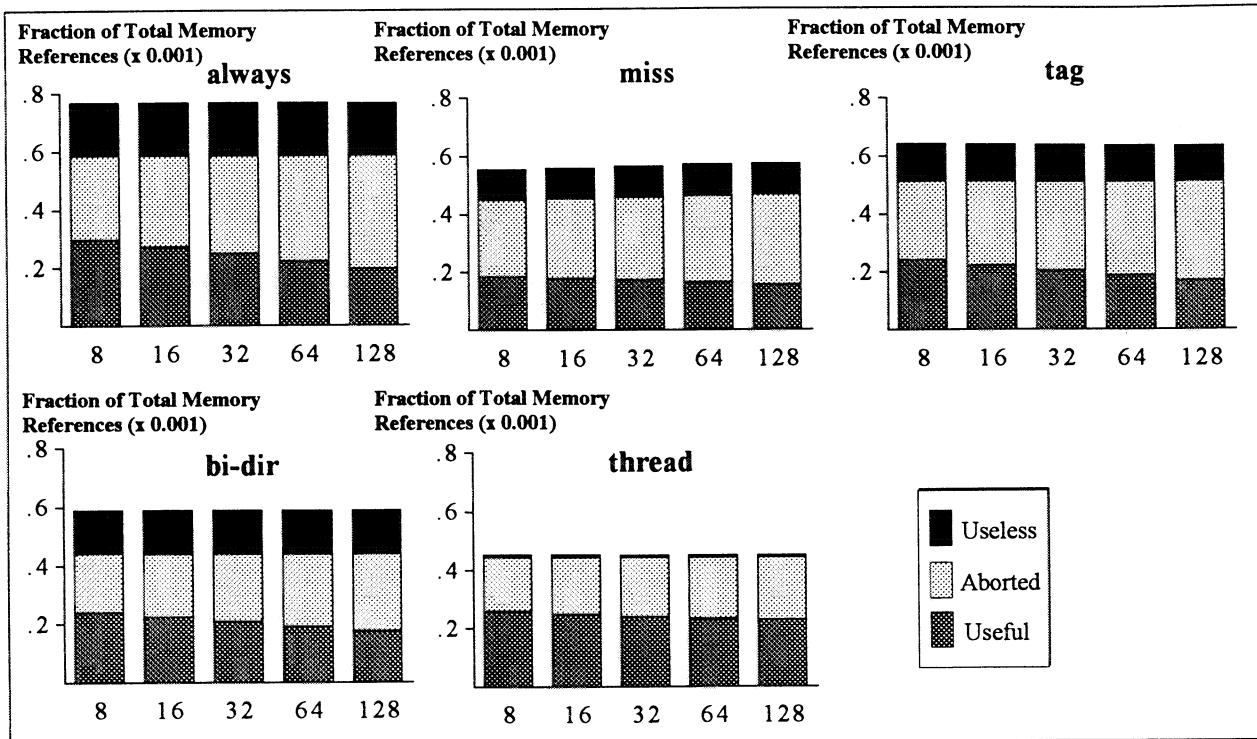


Figure B.8.1. Average distribution of prefetches for the instruction cache.

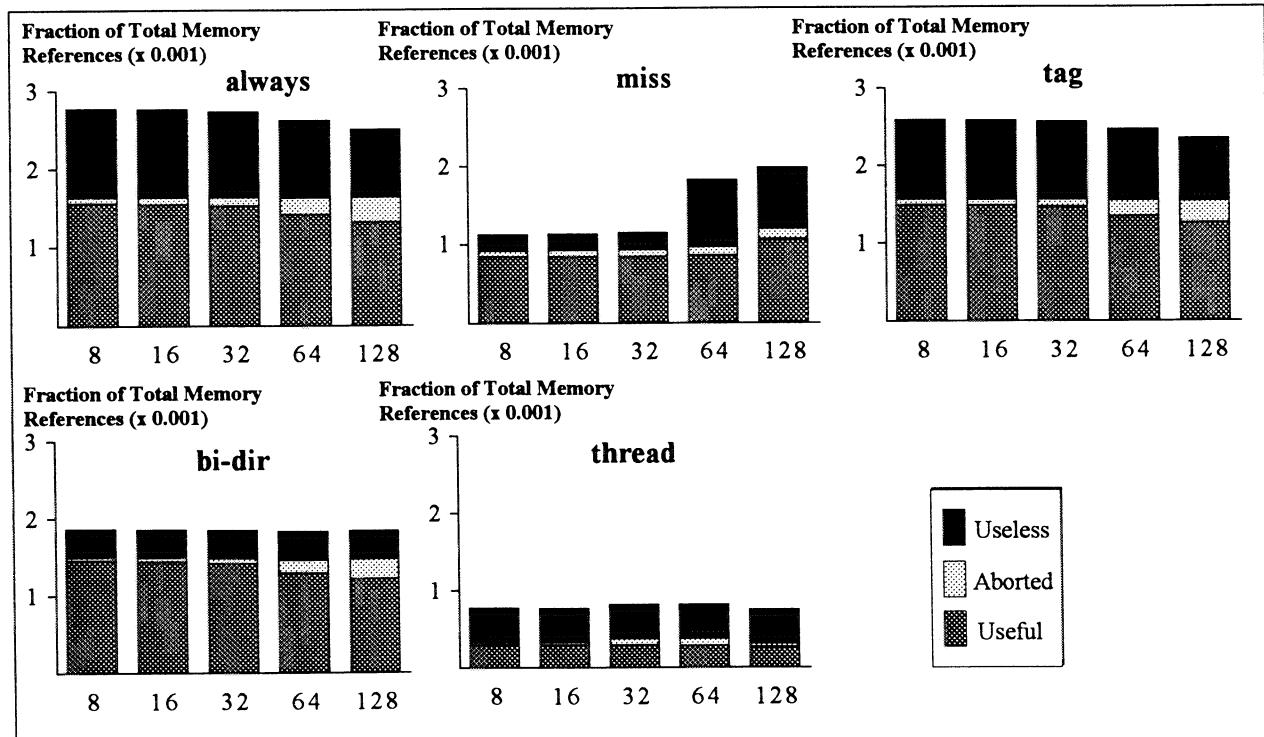


Figure B.8.2. Average distribution of prefetches for the data cache.

Bus Utilization Breakdown (%)	none					always				
	1	2	4	8	16	1	2	4	8	16
Icache read	1.42	1.35	1.29	1.29	1.29	0.56	0.44	0.56	0.55	0.55
Dcache read	5.52	5.34	5.10	5.10	5.10	2.43	2.53	2.48	2.48	2.48
Dcache write	1.64	1.76	1.78	1.78	1.78	1.67	1.82	1.86	1.86	1.86
prefetch	0.00	0.00	0.00	0.00	0.00	6.73	6.77	6.41	6.41	6.41
total	8.57	8.45	8.17	8.17	8.17	11.38	11.57	11.31	11.30	11.30
Bus Utilization Breakdown (%)	miss					tag				
	1	2	4	8	16	1	2	4	8	16
Icache read	0.66	0.62	0.75	0.75	0.75	0.57	0.46	0.57	0.57	0.56
Dcache read	3.69	3.77	3.71	3.71	3.71	2.41	2.51	2.46	2.46	2.46
Dcache write	1.60	1.78	1.82	1.82	1.82	1.66	1.80	1.84	1.84	1.84
prefetch	4.34	4.36	4.01	4.00	4.00	6.26	6.26	5.91	5.92	5.92
total	10.28	10.53	10.29	10.29	10.29	10.90	11.04	10.79	10.79	10.79
Bus Utilization Breakdown (%)	bi-dir					thread				
	1	2	4	8	16	1	2	4	8	16
Icache read	0.79	0.72	0.81	0.82	0.81	0.56	0.50	0.55	0.55	0.55
Dcache read	2.76	2.70	2.59	2.59	2.59	4.79	4.72	4.57	4.57	4.57
Dcache write	1.72	1.84	1.86	1.86	1.86	1.61	1.77	1.80	1.80	1.80
prefetch	4.97	4.80	4.46	4.46	4.46	2.35	2.37	2.21	2.21	2.21
total	10.24	10.05	9.73	9.73	9.73	9.32	9.36	9.12	9.12	9.12

Table B.9.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none					always				
	1	2	4	8	16	1	2	4	8	16
cache misses	100.00	100.00	100.00	100.00	100.00	50.98	46.74	51.75	51.40	51.50
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	15.88	29.85	40.39	39.84	39.91
prefetch(bus)	0.00	0.00	0.00	0.00	0.00	33.05	23.26	7.77	8.71	8.57
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.09	0.16	0.10	0.04	0.02
write buffer	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CPU Stall Reasons	miss					tag				
	1	2	4	8	16	1	2	4	8	16
cache misses	63.00	62.11	72.18	72.04	72.09	52.35	48.39	53.75	53.43	53.51
prefetch(cache)	6.26	15.14	21.93	21.71	21.76	15.35	28.60	38.59	38.12	38.17
prefetch(bus)	30.70	22.68	5.85	6.21	6.14	32.21	22.87	7.58	8.41	8.29
prefetch(mem)	0.05	0.07	0.04	0.02	0.01	0.08	0.14	0.09	0.04	0.02
write buffer	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00
CPU Stall Reasons	bi-dir					thread				
	1	2	4	8	16	1	2	4	8	16
cache misses	63.56	60.07	61.99	62.05	62.06	77.03	78.92	85.88	85.75	85.75
prefetch(cache)	15.99	28.08	33.67	33.70	33.70	3.02	6.64	9.71	9.63	9.61
prefetch(bus)	20.29	11.67	4.27	4.22	4.23	19.88	14.36	4.37	4.60	4.63
prefetch(mem)	0.16	0.18	0.07	0.03	0.02	0.06	0.08	0.03	0.01	0.01
write buffer	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table B.9.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory.

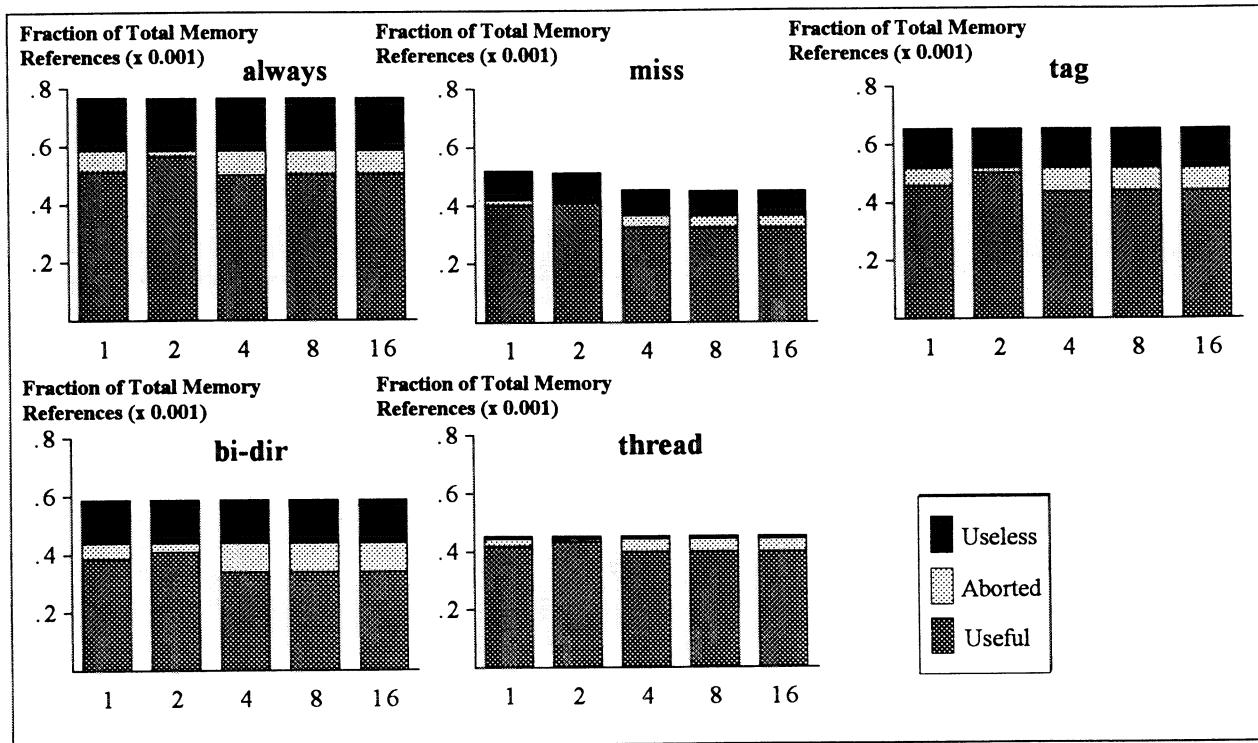
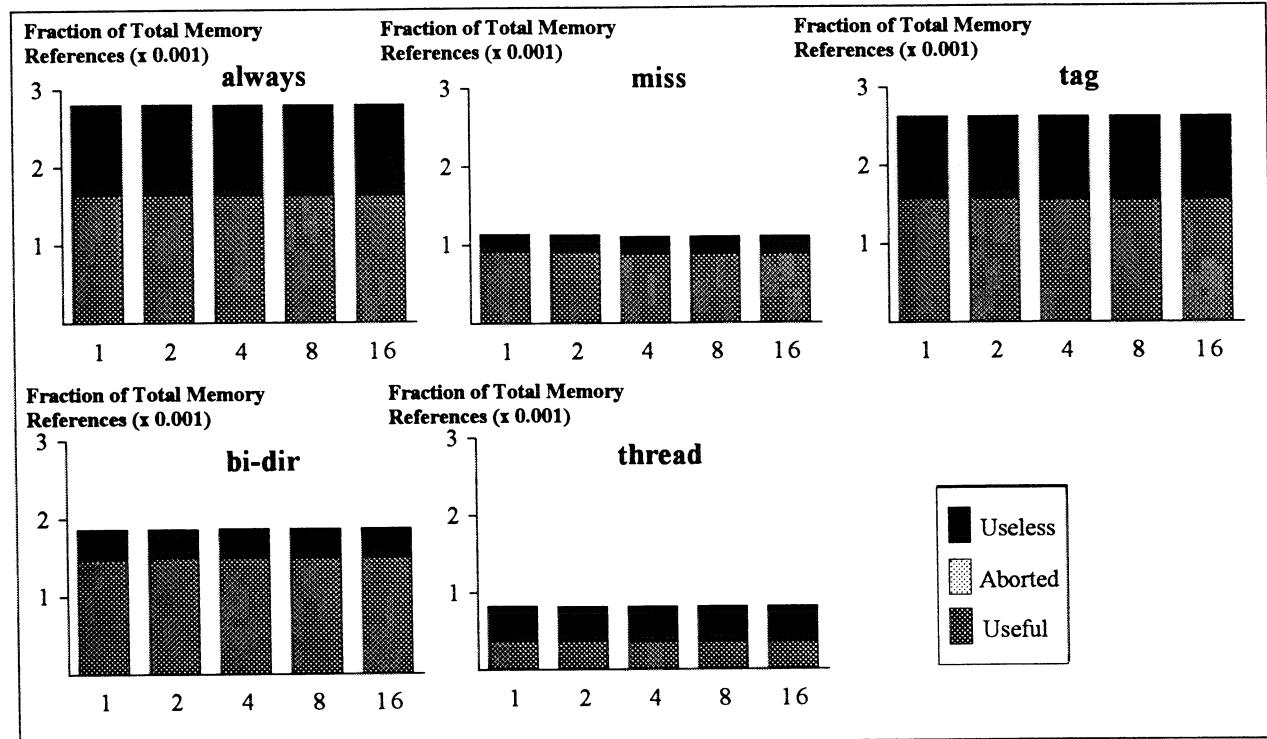


Figure B.9.1. Average distribution of prefetches for the instruction cache.



Bus Utilization Breakdown %	none					always				
	5	17	44	72	78	5	17	44	72	78
Icache read	2.49	2.47	2.41	2.39	2.39	1.97	2.15	2.47	2.44	2.39
Dcache read	9.84	9.75	9.32	9.27	9.26	5.00	5.09	7.13	9.10	9.26
Dcache write	1.77	1.75	1.67	1.66	1.66	1.82	1.80	1.75	1.67	1.66
prefetch	0.00	0.00	0.00	0.00	0.00	10.79	10.10	5.03	0.99	0.00
IO activities	4.83	16.73	43.91	71.87	78.44	4.83	16.72	43.85	71.88	78.44
total	18.93	30.70	57.31	85.18	91.75	24.41	35.87	60.22	86.08	91.75
Bus Utilization Breakdown %	miss					tag				
	5	17	44	72	78	5	17	44	72	78
Icache read	2.04	2.17	2.39	2.41	2.39	1.94	2.10	2.38	2.41	2.39
Dcache read	7.16	7.18	8.17	9.17	9.26	4.97	5.07	7.09	9.06	9.26
Dcache write	1.78	1.77	1.72	1.66	1.66	1.80	1.79	1.74	1.67	1.66
prefetch	6.83	6.37	2.74	0.41	0.00	9.86	9.20	4.14	0.55	0.00
IO activities	4.83	16.73	43.86	71.87	78.44	4.83	16.72	43.85	71.87	78.44
total	22.64	34.22	58.87	85.53	91.75	23.40	34.89	59.19	85.55	91.75
Bus Utilization Breakdown %	bi-dir					thread				
	5	17	44	72	78	5	17	44	72	78
Icache read	2.05	2.20	2.47	2.44	2.39	1.67	1.85	2.18	2.37	2.39
Dcache read	5.24	5.30	7.00	9.10	9.26	8.95	8.94	9.13	9.24	9.26
Dcache write	1.84	1.83	1.76	1.67	1.66	1.76	1.75	1.67	1.66	1.66
prefetch	7.82	7.48	4.79	0.95	0.01	3.33	3.00	0.94	0.11	0.00
IO activities	4.83	16.72	43.84	71.88	78.44	4.83	16.73	43.91	71.87	78.44
total	21.78	33.53	59.86	86.03	91.75	20.54	32.27	57.83	85.25	91.75

Table B.10.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none					always				
	5	17	44	72	78	5	17	44	72	78
cache misses	94.89	86.76	67.94	64.40	64.14	50.89	49.72	56.70	61.62	64.14
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	33.13	31.13	16.04	4.07	0.03
prefetch(bus)	0.00	0.00	0.00	0.00	0.00	12.89	9.79	2.15	0.10	0.00
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.69	0.64	0.49	0.47	0.47	0.52	0.51	0.45	0.42	0.47
IO activities	4.42	12.59	31.57	35.13	35.39	2.57	8.84	24.67	33.80	35.38

CPU Stall Reasons	miss					tag				
	5	17	44	72	78	5	17	44	72	78
cache misses	65.21	61.93	62.35	63.18	64.14	52.97	51.66	58.83	62.88	64.14
prefetch(cache)	19.15	17.90	8.86	1.76	0.00	31.44	29.30	13.63	2.37	0.01
prefetch(bus)	11.74	8.87	1.51	0.05	0.00	12.40	9.45	1.85	0.06	0.00
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.55	0.52	0.47	0.44	0.47	0.56	0.54	0.48	0.45	0.47
IO activities	3.34	10.78	26.81	34.57	35.38	2.63	9.05	25.21	34.25	35.38

CPU Stall Reasons	bi-dir					thread				
	5	17	44	72	78	5	17	44	72	78
cache misses	59.85	56.80	57.08	61.73	64.14	79.86	74.59	65.21	64.04	64.14
prefetch(cache)	29.41	27.26	15.69	3.94	0.02	7.82	7.04	2.72	0.49	0.00
prefetch(bus)	7.44	5.90	2.03	0.10	0.00	7.74	6.09	0.90	0.02	0.00
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.56	0.55	0.45	0.41	0.47	0.63	0.60	0.47	0.46	0.47
IO activites	2.75	9.48	24.75	33.82	35.38	3.96	11.67	30.69	34.99	35.39

Table B.10.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory; write buffer = write backs by the write buffer.

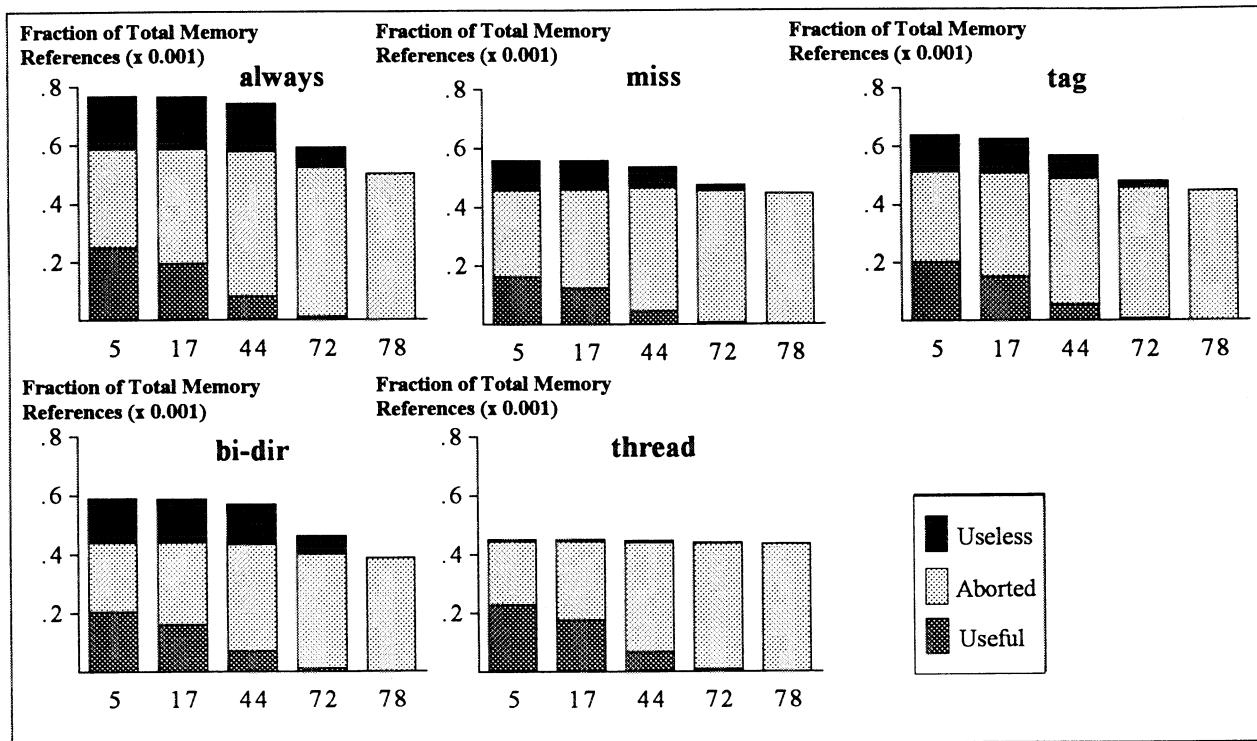


Figure B.10.1. Average distribution of prefetches for the instruction cache.

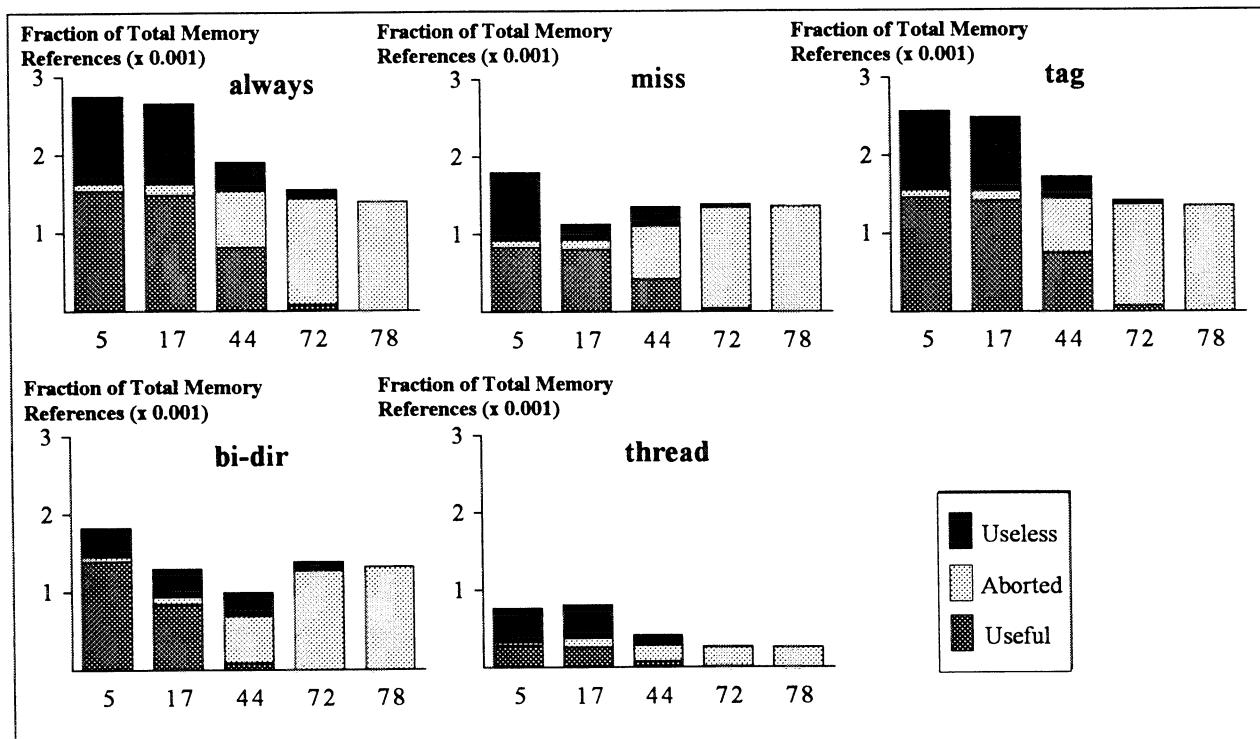


Figure B.10.2. Average distribution of prefetches for the data cache.

B.11. Prefetch Lookahead Distance

Prefetch Lookahead Dist	True Miss Ratios (Instruction Cache)								True Miss Ratios (Data Cache)							
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg		
1	0.0011	0.0008	0.0008	0.0008	0.0008	0.0007	0.0008	0.0109	0.0065	0.0085	0.0065	0.0065	0.0101	0.0076		
2	0.0011	0.0009	0.0009	0.0009	0.0009	0.0007	0.0009	0.0109	0.0069	0.0088	0.0068	0.0069	0.0089	0.0076		
3	0.0011	0.0010	0.0009	0.0009	0.0010	0.0008	0.0009	0.0109	0.0066	0.0087	0.0066	0.0068	0.0099	0.0077		
4	0.0011	0.0011	0.0010	0.0009	0.0010	0.0008	0.0010	0.0109	0.0071	0.0089	0.0070	0.0071	0.0087	0.0077		
Prefetch Lookahead Dist	Partial Miss Ratios (Instruction Cache)								Partial Miss Ratios (Data Cache)							
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg		
1	0.0000	0.0002	0.0001	0.0002	0.0002	0.0002	0.0002	0.0000	0.0002	0.0002	0.0002	0.0001	0.0002	0.0002		
2	0.0000	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0000	0.0003	0.0001	0.0002	0.0001	0.0002	0.0002		
3	0.0000	0.0001	0.0000	0.0000	0.0001	0.0001	0.0001	0.0000	0.0001	0.0001	0.0001	0.0000	0.0001	0.0001		
4	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000	0.0000	0.0001	0.0001	0.0001	0.0000	0.0001	0.0001		
Prefetch Lookahead Dist	Total Miss Ratios (Instruction Cache)								Total Miss Ratios (Data Cache)							
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg		
1	0.0011	0.0010	0.0010	0.0010	0.0010	0.0009	0.0010	0.0109	0.0067	0.0087	0.0066	0.0066	0.0103	0.0078		
2	0.0011	0.0010	0.0010	0.0009	0.0010	0.0009	0.0010	0.0109	0.0071	0.0089	0.0070	0.0070	0.0091	0.0078		
3	0.0011	0.0011	0.0010	0.0010	0.0010	0.0008	0.0010	0.0109	0.0068	0.0088	0.0068	0.0068	0.0100	0.0078		
4	0.0011	0.0011	0.0010	0.0010	0.0011	0.0008	0.0010	0.0109	0.0073	0.0090	0.0071	0.0071	0.0088	0.0078		

Table B.11.1. Instruction and data cache miss ratios.

Prefetch Lookahead Dist	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
1	0.6150	0.6638	0.6613	0.6190	0.9752	0.7069	0.5962	0.5788	0.6052	0.6605	0.6430	0.6167
2	0.4781	0.5330	0.5331	0.4765	0.9610	0.5963	0.5471	0.5180	0.5607	0.6267	0.7028	0.5911
3	0.4158	0.4687	0.4731	0.4162	0.9567	0.5461	0.5572	0.5265	0.5729	0.6242	0.6463	0.5854
4	0.3716	0.4157	0.4212	0.3772	0.9550	0.5081	0.5080	0.4855	0.5227	0.5816	0.7332	0.5662
Prefetch Lookahead Dist	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
1	0.2984	0.2168	0.2649	0.2619	0.2779	0.2640	0.4926	0.2826	0.4718	0.4620	0.0969	0.3612
2	0.2689	0.1855	0.2346	0.2334	0.2345	0.2314	0.4614	0.2624	0.4417	0.4185	0.1829	0.3534
3	0.2553	0.1685	0.2166	0.2192	0.2298	0.2179	0.4621	0.2604	0.4416	0.4156	0.1083	0.3376
4	0.2307	0.1508	0.1925	0.2041	0.2250	0.2006	0.4369	0.2544	0.4170	0.4057	0.2160	0.3460

Table B.11.2. Success ratios and global success ratios.

Bus Utilization Breakdown %	none				always			
	1	2	3	4	1	2	3	4
Icache read	2.50	2.50	2.50	2.50	1.89	2.13	2.30	2.42
Dcache read	9.90	9.90	9.90	9.90	4.98	5.26	5.06	5.46
Dcache write	1.78	1.78	1.78	1.78	1.82	1.84	1.84	1.85
prefetch	0.00	0.00	0.00	0.00	10.99	11.50	11.53	12.14
total	14.18	14.18	14.18	14.18	19.68	20.72	20.73	21.87
Bus Utilization Breakdown %	miss				tag			
	1	2	3	4	1	2	3	4
Icache read	1.99	2.13	2.20	2.26	1.87	2.02	2.12	2.18
Dcache read	7.17	7.36	7.29	7.43	4.94	5.20	5.11	5.40
Dcache write	1.79	1.80	1.79	1.83	1.81	1.82	1.83	1.84
prefetch	6.97	7.23	7.13	7.49	10.05	10.20	10.22	10.47
total	17.91	18.52	18.41	19.00	18.67	19.25	19.28	19.89
Bus Utilization Breakdown %	bi-dir				thread			
	1	2	3	4	1	2	3	4
Icache read	1.99	2.16	2.30	2.39	1.61	1.76	1.78	1.79
Dcache read	5.22	5.54	5.46	5.72	8.97	7.62	8.78	7.46
Dcache write	1.84	1.85	1.86	1.87	1.78	1.78	1.78	1.78
prefetch	7.95	7.88	8.00	8.46	3.44	4.38	3.07	4.17
total	17.00	17.43	17.62	18.44	15.79	15.55	15.41	15.21

Table B.11.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none				always			
	1	2	3	4	1	2	3	4
cache misses	99.27	99.27	99.27	99.27	51.56	50.56	49.46	49.54
prefetch(cache)	0.00	0.00	0.00	0.00	33.97	33.86	34.54	34.38
prefetch(bus)	0.00	0.00	0.00	0.00	13.94	15.11	15.54	15.64
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.72	0.72	0.72	0.72	0.52	0.48	0.47	0.44
CPU Stall Reasons	miss				tag			
	1	2	3	4	1	2	3	4
cache misses	66.88	65.22	64.88	64.15	53.69	53.16	52.53	52.78
prefetch(cache)	19.68	19.65	19.43	20.08	32.28	31.76	32.26	31.70
prefetch(bus)	12.86	14.61	15.17	15.28	13.47	14.56	14.70	15.05
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.58	0.52	0.53	0.49	0.56	0.52	0.52	0.48
CPU Stall Reasons	bi-dir				thread			
	1	2	3	4	1	2	3	4
cache misses	61.07	61.48	60.72	60.10	82.74	77.05	82.67	76.89
prefetch(cache)	30.38	29.88	29.80	30.28	8.23	13.54	8.87	14.84
prefetch(bus)	7.98	8.11	8.96	9.12	8.38	8.76	7.84	7.60
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
write buffer	0.57	0.54	0.52	0.50	0.65	0.66	0.63	0.67

Table B.11.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory.

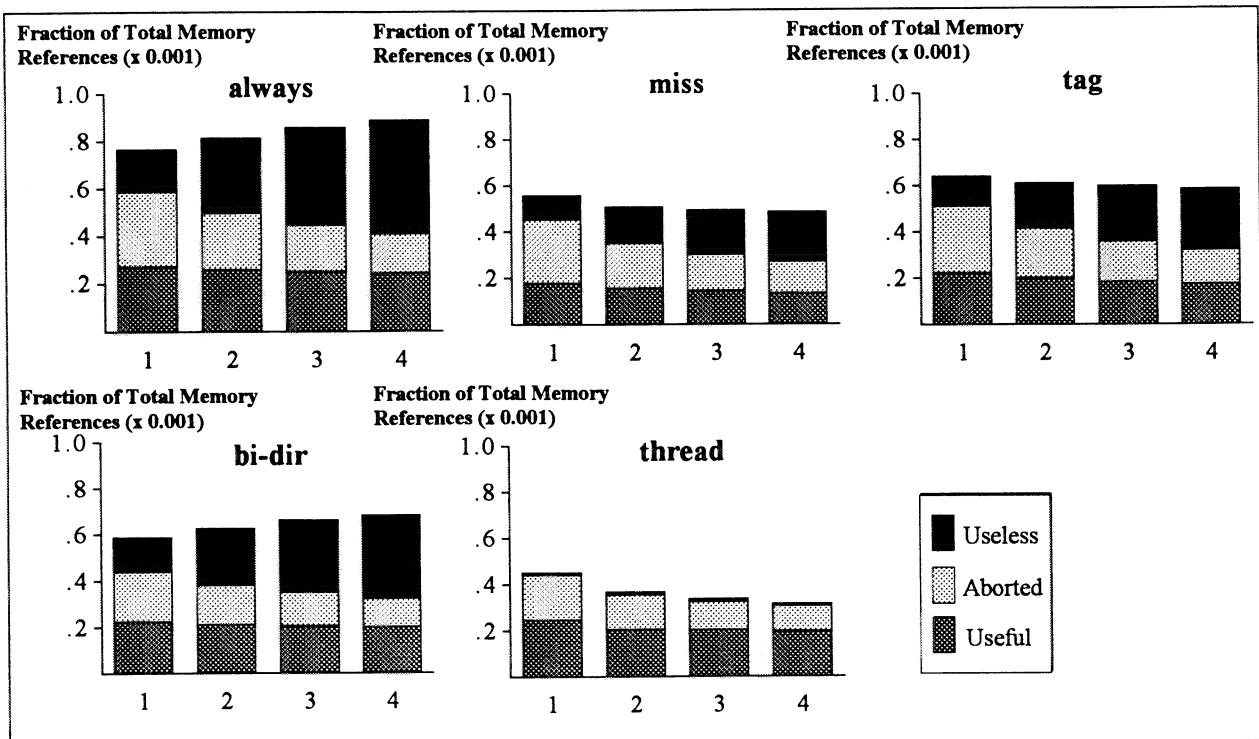


Figure B.11.1. Average distribution of prefetches for the instruction cache.

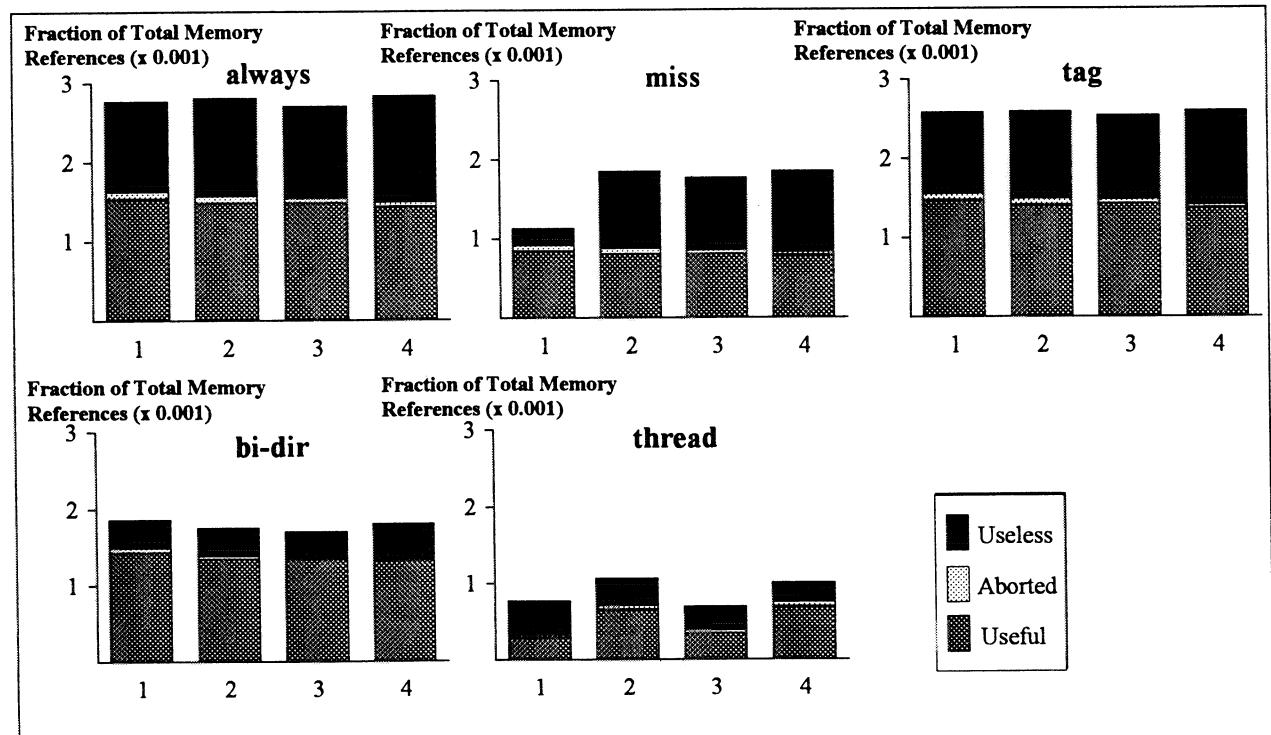


Figure B.11.2. Average distribution of prefetches for the data cache.

C.1. The Best System

Trace Name	True Miss Ratios (Instruction Cache)							True Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
cad	0.0021	0.0004	0.0012	0.0004	0.0007	0.0004	0.0006	0.0132	0.0077	0.0100	0.0076	0.0090	0.0117	0.0092
comp	0.0046	0.0010	0.0026	0.0011	0.0019	0.0010	0.0015	0.0160	0.0071	0.0111	0.0072	0.0078	0.0138	0.0094
fp	0.0035	0.0001	0.0018	0.0001	0.0002	0.0014	0.0007	0.0684	0.0023	0.0348	0.0025	0.0027	0.0527	0.0190
text	0.0017	0.0004	0.0010	0.0004	0.0000	0.0005	0.0006	0.0047	0.0009	0.0028	0.0011	0.0011	0.0037	0.0019
unix	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0235	0.0198	0.0215	0.0198	0.0198	0.0201	0.0202
trace average	0.0024	0.0004	0.0013	0.0004	0.0007	0.0007	0.0007	0.0252	0.0076	0.0160	0.0076	0.0081	0.0204	0.0119
Trace Name	Partial Miss Ratios (Instruction Cache)							Partial Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
cad	0.0000	0.0009	0.0000	0.0008	0.0013	0.0009	0.0008	0.0000	0.0012	0.0000	0.0012	0.0017	0.0003	0.0009
comp	0.0000	0.0017	0.0001	0.0016	0.0022	0.0020	0.0015	0.0000	0.0016	0.0000	0.0016	0.0020	0.0003	0.0011
fp	0.0000	0.0022	0.0001	0.0021	0.0032	0.0012	0.0018	0.0000	0.0260	0.0000	0.0257	0.0418	0.0007	0.0188
text	0.0000	0.0008	0.0000	0.0006	0.0010	0.0007	0.0006	0.0000	0.0008	0.0000	0.0007	0.0009	0.0004	0.0005
unix	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0007	0.0000	0.0007	0.0011	0.0013	0.0008
trace average	0.0000	0.0011	0.0000	0.0010	0.0015	0.0010	0.0009	0.0000	0.0061	0.0000	0.0060	0.0095	0.0006	0.0044
Trace Name	Total Miss Ratios (Instruction Cache)							Total Miss Ratios (Data Cache)						
	none	always	miss	tag	bi-dir	thread	pf avg	none	always	miss	tag	bi-dir	thread	pf avg
cad	0.0021	0.0013	0.0012	0.0013	0.0020	0.0013	0.0014	0.0132	0.0089	0.0100	0.0088	0.0107	0.0121	0.0101
comp	0.0046	0.0027	0.0027	0.0027	0.0041	0.0030	0.0031	0.0160	0.0087	0.0111	0.0088	0.0098	0.0141	0.0105
fp	0.0035	0.0023	0.0018	0.0023	0.0034	0.0027	0.0025	0.0684	0.0283	0.0348	0.0282	0.0445	0.0534	0.0378
text	0.0017	0.0011	0.0010	0.0011	0.0017	0.0012	0.0012	0.0047	0.0016	0.0028	0.0018	0.0020	0.0041	0.0025
unix	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0235	0.0204	0.0215	0.0204	0.0209	0.0214	0.0210
trace average	0.0024	0.0015	0.0013	0.0015	0.0022	0.0016	0.0016	0.0252	0.0136	0.0160	0.0136	0.0176	0.0210	0.0164

Table C.1.1. Instruction and data cache miss ratios in the best system.

Trace Name	Success Ratios (Instruction Cache)						Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
cad	0.8701	0.8879	0.8773	0.8560	0.9959	0.8974	0.5303	0.4503	0.5369	0.5314	0.7764	0.5651
comp	0.8118	0.8419	0.8214	0.7856	0.9916	0.8505	0.7051	0.6319	0.7083	0.7943	0.7706	0.7220
fp	0.9810	0.9825	0.9819	0.9818	0.9994	0.9853	0.9887	0.9879	0.9891	0.9861	0.8878	0.9679
text	0.8417	0.8699	0.8543	0.8210	0.9843	0.8742	0.8920	0.8876	0.9086	0.8770	0.9448	0.9020
unix	0.7500	0.8511	0.8516	0.7220	0.0000	0.6349	0.2365	0.1866	0.2415	0.6522	0.3710	0.3376
trace average	0.8509	0.8867	0.8773	0.8333	0.7942	0.8485	0.6705	0.6289	0.6769	0.7682	0.7501	0.6989
Trace Name	Global Success Ratios (Instruction Cache)						Global Success Ratios (Data Cache)					
	always	miss	tag	bi-dir	thread	avg.	always	miss	tag	bi-dir	thread	avg.
cad	0.8461	0.4494	0.8065	0.7048	0.8032	0.7220	0.4422	0.2676	0.4416	0.3389	0.1178	0.3216
comp	0.7874	0.4478	0.7773	0.5982	0.7789	0.6779	0.5664	0.3224	0.5624	0.5199	0.1467	0.4236
fp	0.9738	0.5000	0.9708	0.9515	0.5947	0.7982	0.9666	0.4929	0.9637	0.9616	0.2321	0.7234
text	0.8162	0.4376	0.7687	0.6467	0.7074	0.6753	0.8252	0.4281	0.7753	0.7896	0.2212	0.6079
unix	0.7927	0.4223	0.7270	0.6237	0.0000	0.5131	0.2191	0.1430	0.2173	0.1666	0.1782	0.1848
trace average	0.8432	0.4514	0.8101	0.7050	0.5768	0.6773	0.6039	0.3308	0.5921	0.5553	0.1792	0.4523

Table C.1.2. Success ratios and global success ratios in the best system.

Bus Utilization Breakdown (%)	none					always				
	cad	comp	fp	text	unix	cad	comp	fp	text	unix
Icache read	0.27	0.54	0.21	0.26	0.00	0.05	0.14	0.01	0.06	0.00
Dcache read	0.58	0.55	1.77	0.22	0.85	0.38	0.30	0.11	0.05	0.74
Dcache write	0.10	0.16	0.39	0.07	0.15	0.12	0.19	0.74	0.08	0.16
prefetch	0.00	0.00	0.00	0.00	0.00	0.90	1.21	3.64	0.57	0.89
total	0.95	1.25	2.37	0.55	1.00	1.44	1.85	4.50	0.76	1.79
Bus Utilization Breakdown (%)	miss					tag				
	cad	comp	fp	text	unix	cad	comp	fp	text	unix
Icache read	0.17	0.35	0.15	0.17	0.00	0.07	0.15	0.01	0.07	0.00
Dcache read	0.47	0.44	1.33	0.14	0.80	0.37	0.30	0.12	0.06	0.75
Dcache write	0.11	0.18	0.58	0.08	0.15	0.11	0.19	0.73	0.08	0.16
prefetch	0.54	0.67	1.47	0.27	0.72	0.86	1.18	3.62	0.51	0.86
total	1.30	1.64	3.53	0.65	1.67	1.42	1.82	4.48	0.72	1.76
Bus Utilization Breakdown (%)	bi-dir					thread				
	cad	comp	fp	text	unix	cad	comp	fp	text	unix
Icache read	0.10	0.25	0.02	0.11	0.00	0.06	0.13	0.10	0.08	0.00
Dcache read	0.41	0.30	0.11	0.05	0.74	0.54	0.53	1.59	0.18	0.75
Dcache write	0.11	0.18	0.62	0.08	0.15	0.11	0.18	0.46	0.07	0.15
prefetch	0.66	0.87	3.07	0.48	0.23	0.33	0.59	0.69	0.25	0.44
total	1.28	1.60	3.81	0.72	1.13	1.04	1.43	2.83	0.59	1.35

Table C.1.3. Breakdowns of bus utilization (in percentage).

CPU Stall Reasons	none						always					
	cad	comp	fp	text	unix	avg	cad	comp	fp	text	unix	avg
cache misses	99.58	99.10	96.93	99.64	99.04	98.86	98.00	96.91	83.12	97.35	98.25	94.77
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
prefetch(bus)	0.00	0.00	0.00	0.00	0.00	0.00	1.59	2.87	16.48	2.37	1.50	4.96
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.02	0.01	0.01
write buffer	0.42	0.90	3.07	0.36	0.96	1.14	0.17	0.21	0.40	0.26	0.24	0.26
CPU Stall Reasons	miss						tag					
	cad	comp	fp	text	unix	avg	cad	comp	fp	text	unix	avg
cache misses	99.39	98.99	94.82	99.38	98.81	98.28	98.00	96.99	83.14	97.72	98.37	94.91
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
prefetch(bus)	0.42	0.81	3.13	0.35	0.94	1.13	1.48	2.79	16.45	1.97	1.39	4.82
prefetch(mem)	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.01	0.00	0.01
write buffer	0.19	0.20	2.05	0.28	0.25	0.59	0.17	0.21	0.40	0.30	0.24	0.26
CPU Stall Reasons	bi-dir						thread					
	cad	comp	fp	text	unix	avg	cad	comp	fp	text	unix	avg
cache misses	98.81	97.81	89.47	98.43	98.61	96.63	99.00	98.41	94.35	98.99	98.44	97.85
prefetch(cache)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
prefetch(bus)	0.78	1.78	10.14	1.15	0.41	2.85	0.55	0.92	2.41	0.68	0.96	1.10
prefetch(mem)	0.01	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.01	0.00	0.00	0.00
write buffer	0.40	0.40	0.38	0.42	0.97	0.51	0.39	0.67	3.23	0.33	0.60	1.04

Table C.1.4. Stall breakdowns in percentage. Key: cache misses = misses in the instruction and data caches; prefetch(cache) = conflicts with prefetches in the caches; prefetch(bus) = conflicts with prefetches in the data bus; prefetch(mem) = conflicts with prefetches in the main memory.

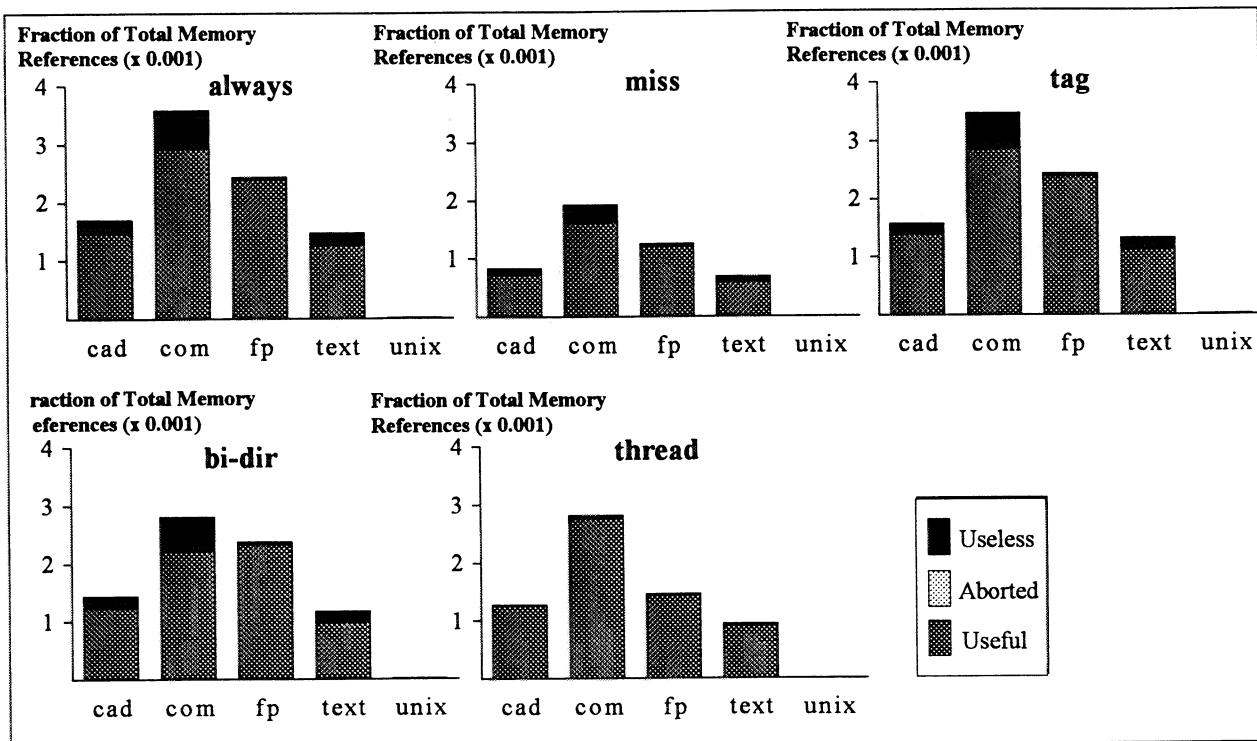


Figure C.1.1. Average distribution of prefetches for the instruction cache.

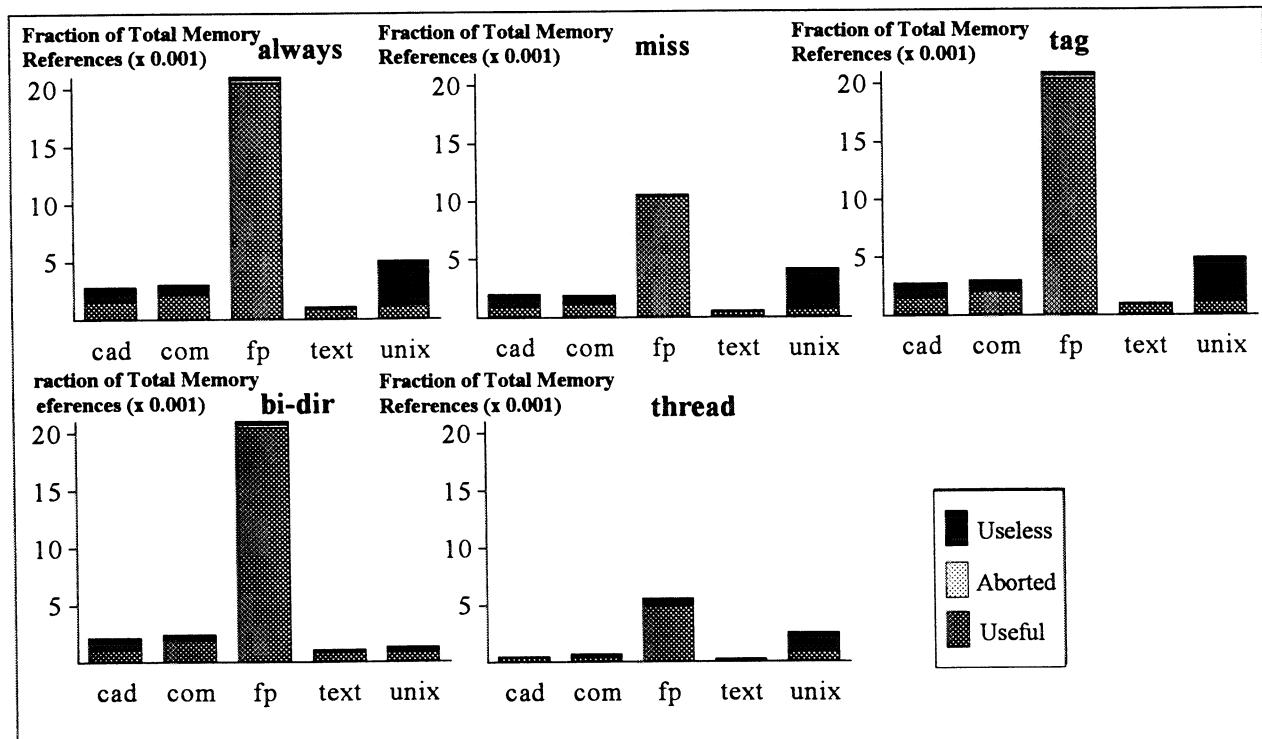


Figure C.1.2. Average distribution of prefetches for the data cache.