

Code Transformations to Improve Memory Parallelism *

Vijay S. Pai[†] and Sarita Adve[‡]

[†]Electrical and Computer Engineering
Rice University
Houston, TX 77005
vijaypai@rice.edu

[‡]Computer Science
University of Illinois
Urbana-Champaign, IL 61801
sadve@cs.uiuc.edu

Abstract

Current microprocessors incorporate techniques to exploit instruction-level parallelism (ILP). However, previous work has shown that these ILP techniques are less effective in removing memory stall time than CPU time, making the memory system a greater bottleneck in ILP-based systems than previous-generation systems. These deficiencies arise largely because applications present limited opportunities for an out-of-order issue processor to overlap multiple read misses, the dominant source of memory stalls.

This work proposes code transformations to increase parallelism in the memory system by overlapping multiple read misses within the same instruction window, while preserving cache locality. We present an analysis and transformation framework suitable for compiler implementation. Our simulation experiments show substantial increases in memory parallelism, leading to execution time reductions averaging 23% in a multiprocessor and 30% in a uniprocessor. We see similar benefits on a Convex Exemplar.

1. Introduction

Current commodity microprocessors improve performance through aggressive techniques to exploit high levels of instruction-level parallelism (ILP). These techniques include multiple instruction issue, out-of-order (dynamic) issue, non-blocking reads, and speculative execution.

Our previous work characterized the effectiveness of ILP processors in a shared-memory multiprocessor [14]. Although ILP techniques successfully and consistently reduced the CPU component of execution time, their impact on the memory (read) stall component was lower and more application-dependent, making read stall time a larger bottleneck in ILP-based multiprocessors than in previous-generation systems. In particular, current and future read

miss latencies are too long to overlap with other instruction types. Thus, an ILP processor needs to overlap multiple read misses with each other to hide a significant portion of their latencies. An out-of-order processor can only overlap those reads held together within its instruction window. Independent read misses must therefore be *clustered* together within a single instruction window to effectively hide their latencies (*read miss clustering*). The applications in our study typically did not exhibit much read miss clustering, leading to poor parallelism in the memory system.

This paper presents code transformations to improve memory parallelism for systems with out-of-order processors, while preserving cache locality. We exploit code transformations already known and implemented in compilers for other purposes, providing the analysis needed to relate them to memory parallelism. The key transformation we use is unroll-and-jam, which was originally proposed for improving floating-point pipelining and for scalar replacement [1, 2, 4, 11]. We develop an analysis that maps the memory parallelism problem to floating-point pipelining.

We evaluate these transformations applied by hand to a latency-detection microbenchmark and five scientific applications running on simulated and real uniprocessor and multiprocessor systems. Our clustering transformations reduce exposed latency by over 80% for the latency-detection microbenchmark. For the scientific applications, the transformations reduce execution time by 9–39% (averaging 23%) in the simulated multiprocessor and 11–48% (averaging 30%) in the simulated uniprocessor. A substantial part of these execution-time reductions arise from improving memory parallelism, particularly as memory stall time becomes more significant. We confirm the benefits of our transformations on a real system (Convex Exemplar), where they reduce application execution time by 9–34%.

An alternative latency tolerating technique is software prefetching, which has been shown to be effective for systems built with simple processors [10]. However, prefetching can be less effective in ILP systems because of increased late prefetches and resource contention [14]. We only consider read miss clustering in this work; our ongoing investigations indicate ways in which clustering transformations can also improve the effectiveness of prefetching [13].

*This work is supported in part by an IBM Partnership award, Intel Corporation, the National Science Foundation under Grant No. CCR-9410457, CCR-9502500, CDA-9502791, and CDA-9617383, and the Texas Advanced Technology Program under Grant No. 003604-025. Sarita Adve is also supported by an Alfred P. Sloan Research Fellowship. Vijay S. Pai was also supported by a Fannie and John Hertz Foundation Fellowship.

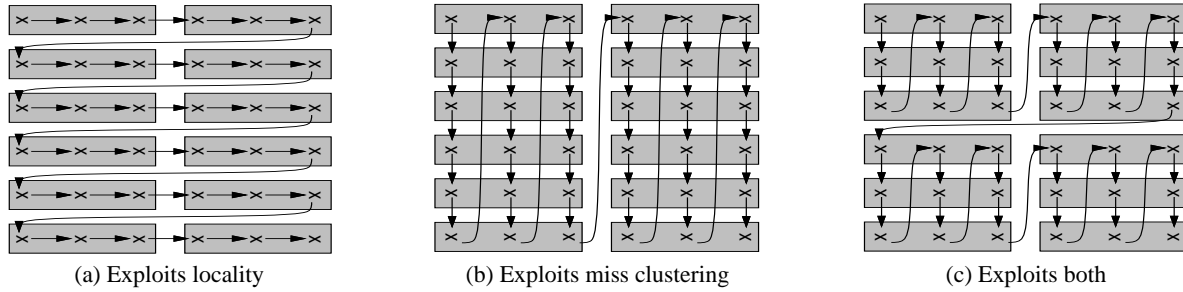


Figure 1. Impact of matrix traversal order on miss clustering. Crosses represent matrix elements, and shaded blocks represent cache lines. The matrix is shown in row-major order.

2. Motivation for Read Miss Clustering

This section discusses the need for read miss clustering, the sources of poor clustering, and code transformations to improve clustering.

2.1. Latency Tolerance in ILP Processors

Instructions in an out-of-order processor’s instruction window (reorder buffer) can issue and complete out-of-order. To maintain precise interrupts, however, instructions commit their results and retire from the window in-order after completion [17]. The only exception is for writes, which can use write-buffering to retire before completion.

Because of the growing gap in processor and memory speeds, external cache misses can take hundreds of processor cycles. However, current out-of-order processors typically have only 32–80 element instruction windows. Consider an outstanding read miss that reaches the head of the window. If all other instructions in the window are fast (e.g., typical computation and read hits) or can be buffered aside (e.g., writes), the independent instructions may not overlap enough latency to keep the processor busy throughout the cache miss. Since the later instructions wait to retire in-order, the instruction window will fill up and block the processor. Thus, this miss remains exposed despite such ILP features as out-of-order issue and non-blocking reads.

Suppose that independent misses from elsewhere in the application could be scheduled into the instruction window behind the outstanding read miss. Then, the later misses are hidden behind the stall time of the first miss. Thus, read miss latencies can typically be effectively overlapped only behind other read misses, and such overlap only occurs if read misses to multiple cache lines appear clustered within the same instruction window. We refer to this phenomenon as *read miss clustering*, or simply *clustering*.

2.2. Increasing Read Miss Clustering

To understand the sources of poor read miss clustering in typical code, we consider a loop nest traversing a 2-D matrix. Figure 1 graphically represents three different matrix

<pre> for(...j++) for(...i++) ...A[j,i] </pre> <p>(a) Base code</p>	<pre> for(...i++) for(...j++) ...A[j,i] </pre> <p>(b) Interchange</p>
<pre> for(...jj+=N) for(...i++) for(j=jj; j<jj+N; j++) ...A[j,i] </pre> <p>(c) Strip-mine and interchange</p>	<pre> for(...j+=N) for(...i++){ ...A[j,i] ...A[j+1,i] A[j+N-1,i] } </pre> <p>(c') Unroll-and-jam</p>

Figure 2. Pseudocode for Figure 1 matrix traversals (row-major notation).

traversals. The matrix is shown in row-major order, with crosses for data elements and shaded blocks for cache lines. Figure 2 relates these matrix traversals to code generation, with pseudocode shown in row-major notation.

Figures 1(a) and 2(a) show a matrix traversal optimized for spatial locality, following much compiler research. In this row-wise traversal, L successive loop iterations access each cache line, where L is the number of data elements per cache line. While this traversal maximizes spatial locality, it minimizes clustering. For example, an instruction window that holds L or fewer iterations never holds read misses to multiple cache lines, preventing clustering. This problem is exacerbated by larger cache lines or larger loop bodies.

Read miss clustering can be maximized by a column-wise traversal, since successive iterations held in the instruction window access different cache lines. Figures 1(b) and 2(b) show such a column-wise traversal, obtained by applying loop interchange to the code in Figure 2(a). Each cache line is now accessed on multiple successive outer-loop iterations. However, the traversal passes through every row before revisiting an older cache line. If there are more rows than cache lines, this traversal could lose all cache locality, potentially overwhelming any performance benefits from clustering.

The above example suggests a tradeoff between spa-

tial locality (favored by current code-generation schemes) and miss clustering. We seek a solution that achieves the benefits of clustering while preserving spatial locality. A column-wise traversal can maximize clustering; however, it must stop before losing locality. In particular, the column-wise traversal can stop as soon as the miss clustering resources are fully utilized. For example, a processor that allows ten simultaneous cache misses sees the maximum memory parallelism when ten independent miss references are clustered. The traversal could then continue in a row-wise fashion to preserve locality. Figure 1(c) shows a matrix traversal that exploits clustering and locality in this way. Figure 2(c) expresses this traversal by applying strip-mine and interchange to Figure 2(a).

Since the column-wise traversal length (N) of Figure 2(c) is based on the hardware resources for overlap (≤ 12 today), the strip size is small, and the innermost loop can be fully unrolled. Figure 2(c') shows the result of that unrolling. Now, the code reflects the transformation of unroll-and-jam applied to Figure 2(a). This transformation unrolls an outer loop and fuses (jams) the resulting inner loop copies into a single inner loop. Previous work has used unroll-and-jam for scalar replacement (replacing array memory operations with register accesses), better floating-point pipelining, or cache locality [1, 2, 3, 4, 11]. Using unroll-and-jam for read miss clustering requires different heuristics, and may help even when the previously studied benefits are unavailable.

We prefer to use unroll-and-jam instead of strip-mine and interchange for two reasons. First, unroll-and-jam allows us to exploit benefits from scalar replacement. Second, unroll-and-jam does not change the inner-loop iteration count. The shorter inner loops of strip-mining can negatively impact techniques that target inner loops, such as dynamic branch prediction. By increasing inner-loop computation without changing the iteration count, unroll-and-jam can also help software prefetching [13].

Unroll-and-jam creates an N -way unrolled steady-state, followed by an untransformed postlude of leftover iterations. To enable clustering in the postlude, we simply interchange the postlude when possible. This should not degrade locality, since the postlude originally has fewer outer-loop iterations than the unroll-and-jam degree.

3. Analysis and Transformation Framework

This section provides a formal framework to apply memory parallelism transformations in a compiler.

3.1. Dependences that Limit Memory Parallelism

We first describe a dependence framework to represent limitations to memory parallelism. As in other domains, dependences here indicate reasons why one operation will

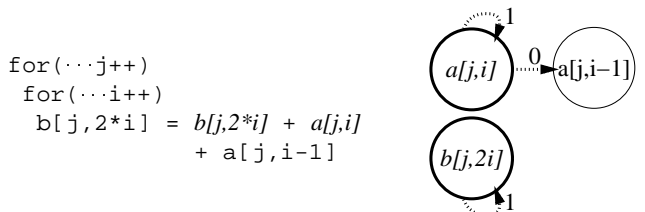
not issue in parallel with another. However, these dependences are not ordinary data dependences, since memory operations can be serialized for different reasons. We build this framework to gauge performance potential, not to specify legality. Thus, we optimistically estimate memory parallelism and specify dependences only when their presence is known. The transformation stages must then use more conventional (and conservative) dependence analysis for legality. For simplicity, we only consider memory parallelism dependences that are either loop-independent or carried on the innermost loop. We can then exploit previous work with the same simplification [2].

Since we focus on parallelism among read misses, we first require locality analysis to determine which static references can miss in the external cache (*leading references*), and which leading references are known to exhibit spatial locality across successive iterations of the innermost loop (*inner-loop self-spatial locality*). Known locality analysis techniques can provide the needed information [19]. Currently, we do not consider cache conflicts in our analysis and transformations.

We use the above information to identify limitations to read miss parallelism. We focus on three kinds of limitations, which we call *cache-line dependences*, *address dependences*, and *window constraints*.

Cache-line dependences. If a read miss is outstanding, then another reference to the same cache line simply coalesces with the outstanding miss, adding no read miss parallelism. Thus, we say that there is a *cache-line dependence* from memory operation A to B if A is a leading reference and a miss on A brings in the data of B. The cache-line dependence is a new resource dependence class, extending input dependences to support multi-word cache lines.

The following code illustrates cache-line dependences. In all examples, leading references known to have inner-loop self-spatial locality will be italicized, while other leading references will be boldfaced. The accompanying graph shows static memory references as nodes and dependences as edges. Each edge is marked with the inner-loop dependence distance, the minimum number of inner-loop iterations separating the dependent operations specified.

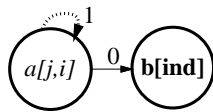


Note that there are no cache-line dependences from one leading reference to another; such a dependence would make the second node a non-leading reference. Additionally, any leading reference with inner-loop self-spatial locality has a cache-line dependence onto itself. That depen-

dence has distance 1 for any stride, since the address of the miss reference will be closer to the instance 1 iteration later than to an instance farther away.

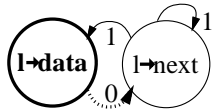
Address dependences. There is an address dependence from memory operation A to B if the result of A is used in computing the address of B, serializing B behind A. Address dependences typically arise for irregular accesses, such as indirect addressing or pointer-chasing. The following code segments show address dependences. The graphs show address dependences as solid lines and cache-line dependences as dotted lines. The first example shows the indirect addressing typical of sparse-matrix applications.

```
for(...j++)
  for(...i++){
    ind = a[j,i]
    sum[j] = sum[j] + b[ind]
```



The above shows one leading reference that exhibits cache-line dependences, connected through an address dependence to another leading reference. The following code shows address dependences from pointer dereferencing.

```
for(...i++){
  l = list[i]
  for(...l=l->next)
    sum[i] += l->data
```



The above assumes that the `data` and `next` fields always lie on the same cache line and that separate instances of `l` are not known to share cache lines. Even though `l->next` is a non-leading reference, it is important since a dependence flows from this node to the leading reference.

Window constraints. Even without other dependences, read miss parallelism is limited to the number of independent read misses in the loop iterations simultaneously held in the instruction window. We do not include these resource limitations in our dependence graphs, since they can change at each stage of transformation. We will, however, consider these constraints in our transformations.

Control-flow and memory consistency requirements may also restrict read miss parallelism. We do not consider these constraints, since their performance impact can be mitigated through well-known static or dynamic techniques such as speculation. However, these dependences may still affect the legality of any code transformations.

Of the three dependence classes that we consider (cache-line, address, and window), only address dependences are true data-flow dependences. Window constraints can be eliminated through careful scheduling of the loop body, possibly enhanced by inner-loop unrolling. Such scheduling would aim to cluster together misses spread over a long loop body. Loop-carried cache line dependences can be converted to loop-invariant dependences through inner-loop unrolling by a multiple of L , where L iterations share each

cache line. Then, no cache line is shared across unrolled loop iterations. The inner-loop unrolling degree may need to go as high as $N \times L$ to provide clustered misses to N separate cache lines. This can be excessive, particularly with long cache lines. We therefore leave these loop-carried cache-line dependences in place and seek to extract read miss parallelism with less code expansion through outer-loop unroll-and-jam.

We will address memory parallelism limitations in loop nests by first resolving recurrences (cycles in the dependence graph), and then handling window constraints. A loop nest may suffer from one or both problems, and recurrence resolution may create new window constraints.

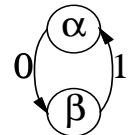
3.2. Resolving Memory Parallelism Recurrences

Unroll-and-jam has previously been used to improve floating-point pipelining in the presence of inner-loop floating-point recurrences [2, 11]. We seek to use unroll-and-jam to target loop nests with memory-parallelism recurrences, which arise for such common access patterns as self-spatial or pointer-chasing leading references. We map memory parallelism to floating-point pipelining, exposing several key similarities and differences between these problems. This section thus shows how to automate the process described in Section 2.2, which used unroll-and-jam to increase miss clustering without degrading locality.

3.2.1. Background on Floating-point Pipelining

Consider an inner loop that carries a floating-point recurrence (a cycle of true dependences). The operations of later iterations can stall for the results of earlier iterations, preventing maximum pipeline throughput. Further, inner-loop unrolling and scheduling cannot help, as later inner-loop iterations are also in the cycle. The following pseudocode has an inner-loop recurrence between statements α and β . The graph shows floating-point true dependences and dependence distances.

```
for(...j++)
  for(...i++){
     $\alpha$ : b[j,i] = a[j,i-1] + c[i]
     $\beta$ : a[j,i] = b[j,i] + d[i]
```



The above recurrence has two floating-point operations, and needs 1 iteration for a complete cycle (the sum of the dependence distances). Thus, the system must serialize 2 floating-point operations (the number in the recurrence) to complete 1 iteration (the length of the cycle), regardless of the pipelining supported. Callahan et al. described floating-point recurrences as follows [2]¹:

¹Their notation was slightly different, with $\#R$, $\tau(R)$, $\rho(R)$, and f_L instead of R , τ , ρ , and f , respectively.

- l_p : number of stages in the floating-point pipeline
- ρ : ratio of the number of nodes (static floating-point operations) in the inner-loop recurrence (R) to the number of iterations to traverse the cycle (τ)
- f : static count of floating-point operations in the innermost loop

Since R floating point operations must be serialized in τ iterations, the recurrence requires at least ρ pipeline latencies ($= l_p \times \rho$ pipeline stages) per iteration. Without dependences, each iteration would require only the time of f pipeline stages. Thus, the recurrence limits pipeline utilization to $\frac{f}{l_p \times \rho}$. Unroll-and-jam introduces independent copies of the recurrence from separate outer loop iterations, increasing f without affecting ρ [2]. To fill the pipeline, unroll-and-jam must be applied until $f \geq l_p \times \rho$. (The maximum ρ should be used for a loop with multiple recurrences, since each recurrence limits pipeline utilization.)

Certain dependences can prevent unroll-and-jam, but they are not directly related to the recurrences targeted. Previous work more thoroughly discusses legality and the choice of outer loops to unroll for deeper nests [2, 4, 11].

3.2.2. Mapping to Memory Parallelism

Above, unroll-and-jam used only the number of pipeline stages, not the latency. The pipeline simply represents the number of floating-point operations that can be processed in parallel. Thus, we can map this algorithm to memory parallelism: the goal is to fully utilize the miss clustering resources, not to schedule for some specific miss latencies. Here, l_p corresponds to the maximum number of simultaneous outstanding misses supported by the processor. The rest of the mapping is more difficult, as not all memory operations utilize the resources for miss parallelism — only those instances of leading references that miss at run-time do. This difference affects ρ and f .

Characterizing recurrences (ρ). We refer to recurrences with only cache-line dependences as *cache-line recurrences* and recurrences with at least one address dependence as *address recurrences*. Recurrences with no leading miss references are irrelevant here and can be ignored, since they do not impact read miss parallelism.

As discussed in Section 3.2.1, ρ is computed from two values: R and τ . We count only leading references in R , as only these nodes can lead to serialization for a miss. We count τ as in Section 3.2.1, since this specifies the number of iterations after a miss instance before serialization. Although our discussion has focused on tolerating read miss latencies, our algorithm must count both read and write miss references in R and f , since writes also require resources. Nevertheless, we will not apply unroll-and-jam on an outer loop if it only adds write misses, since write latencies can be hidden through write-buffering.

Counting memory parallelism candidates (f). For floating-point pipelines, the f parameter counts the static instructions in the innermost loop. We cannot use this same definition here for two key reasons, described below.

Dynamic inner-loop unrolling. An out-of-order instruction window of W instructions dynamically unrolls a loop body of i instructions by $\lceil \frac{W}{i} \rceil$. (For simplicity, we assume no outer-loop unrolling, although this could arise if the inner loop had fewer than $\lceil \frac{W}{i} \rceil$ iterations). Such unrolling exposes no additional steady-state parallelism for loops with address recurrences, since these are analogous to the recurrences of floating-point pipelining. However, this unrolling can actually break cache-line recurrences. In particular, if L_m successive iterations share a cache line for leading reference m , dynamic inner-loop unrolling creates $\lceil \frac{W}{iL_m} \rceil$ independent misses from the original recurrence. Leading references outside recurrences can also contribute multiple outstanding misses ($L_m = 1$, since no cache-line sharing is known). Thus, we define C_m , the number of copies of m that can contribute overlapped misses:

$$C_m = \begin{cases} \lceil \frac{W}{iL_m} \rceil & \text{loop with no address recurrences} \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

Miss patterns. A simple count of leading references can overestimate memory parallelism, since not all leading reference instances miss in the cache. To determine which leading reference instances miss together, we must know the miss patterns (sequences of hits and misses) for the different leading references and their correlation with each other. Such measures can be difficult to determine in general. In this work, we make some simple assumptions, described below.

We split the leading references into two types: regular (arrays indexed with affine functions of the loop indices) and irregular (all others). For regular references, we assume that at least some passes through the inner loop experience misses on each cache line accessed, and that different regular leading references experience misses together. These assumptions lead to maximum estimated parallelism for regular leading references.

For irregulars, the miss pattern is not typically analyzable. We assume no correlation, either among instances of the same reference or across multiple references. Thus, we only need to know the overall miss rate, P_m , for each reference m . P_m can be measured through cache simulation or profiling. These assumptions allow more aggressive transformation than the more common assumption of no locality for irregulars.

We can now estimate the f parameter, accounting for both dynamic inner-loop unrolling and miss patterns:

$$f = f_{reg} + f_{irreg} \quad (2)$$

$$f_{reg} = \sum_{m \in RLR} C_m \quad (3)$$

$$f_{irreg} = \lceil \sum_{m \in ILR} P_m \times C_m \rceil \quad (4)$$

We split f into regular and irregular components, with RLR and ILR the sets of regular and irregular leading references respectively. The terms C_m in Equation 3 and $P_m \times C_m$ in Equation 4 give the maximum expected number of misses to separate cache lines contributed by leading reference m . We round up f_{irreg} to insure that some resources are held for irregular references when they are present.

The floating-point pipelining algorithm applied unroll-and-jam until $f \geq l_p \times \rho$. We should be more conservative for memory parallelism, as the cache can see extra contention when the resources for outstanding misses (MSHRs) fill up. Thus, we aim to apply unroll-and-jam as much as possible while maintaining $f \leq l_p \times \rho$ (using the maximum ρ for the loop).

After applying unroll-and-jam, we must recompute f for two reasons. First, unroll-and-jam can introduce new leading references and increase the iteration size. On the other hand, some leading reference copies may become non-leading references because of scalar replacement or group locality. For similar reasons, we must repeat the locality and dependence analysis passes.

Since f varies as described above, we may need to attempt unroll-and-jam multiple times with different unrolling degrees to reach our desired f . We can limit the number of invocations by choosing a maximum unrolling degree U based on the resources for memory parallelism, code expansion, register pressure, and potential for cache conflicts. If we unroll only one outer loop, we can choose the unrolling degree by binary search, using at most $\lceil \log_2 U \rceil$ passes [4]. Generalized searching for unrolling multiple outer loops can follow the strategies described in previous work [4]. We also refer to previous work for legality issues [2, 4, 11]. We add only that we prefer not to unroll-and-jam loops that only expose additional write miss references, since buffering can hide write latencies.

To revisit the motivating example of Section 2.2, note that the matrix traversal of Figure 2(a) has a cache-line recurrence with $\rho = 1$. L_m typically ranges from 4 to 16 for stride-1 double-word accesses, so $\lceil \frac{W}{iL_m} \rceil$ is most likely 1 for a loop body with a moderate amount of computation and current instruction window sizes. Thus, $f = f_{reg} = 1$ initially. This example has no scalar replacement opportunities, so each recurrence copy created by unroll-and-jam contributes a leading reference to the calculation of f . Assuming U is chosen to be at least l_p , the search algorithm will find that unroll-and-jam by l_p leads to $f = l_p \times \rho$.

3.3. Resolving Window Constraints

We now address memory parallelism limitations from window constraints. These can arise for loops with or without recurrences. Further, recurrence resolution can actually

create new window constraints, since unroll-and-jam can spread its candidates for read miss parallelism over a span of instructions larger than a single instruction window. We proceed in two stages: first using loop unrolling to resolve any inter-iteration window constraints, then using local instruction scheduling to resolve intra-iteration constraints.

As discussed in Section 3.2.2, an instruction window of W instructions dynamically unrolls an inner-loop body of i instructions by $\lceil \frac{W}{i} \rceil$. Inter-iteration window constraints arise when the independent read misses in $\lceil \frac{W}{i} \rceil$ iterations do not fill the resources for memory parallelism (typically because of large loop bodies). Since any recurrences have already been resolved, we can now use inner-loop unrolling to better expose independent misses to the instruction scheduler. We can directly count the maximum expected number of independent misses in $\lceil \frac{W}{i} \rceil$ iterations, using the miss rate P_m to weight the irregular leading references. We then unroll until the resources for memory parallelism are filled, recomputing the exposed independent miss count after each invocation of unrolling.

Now we resolve any intra-iteration window constraints stemming from loop bodies larger than a single instruction window (possibly because of unroll-and-jam or inner-loop unrolling). In such cases, the instruction scheduler should pack independent miss references in the loop body close to each other. The technique of balanced scheduling can provide some of these benefits [6, 7], but may also miss some opportunities since it does not explicitly consider window size. Nevertheless, this heuristic worked well for the code sequences we examined. More appropriate local scheduling algorithms remain the subject of future research.

4. Experimental Methodology

4.1. Evaluation Environments

We perform most of our experiments using RSIM, the Rice Simulator for ILP Multiprocessors [15]. We model both an ILP uniprocessor and an ILP-based CC-NUMA multiprocessor with release consistency. Table 1 summarizes the base configuration. The cache sizes are scaled based on application input sizes according to the methodology of Woo et al. [20]. The memory banks use permutation-based interleaving on a cache-line granularity to support a variety of strides [18]. The simulated system latencies without contention are 1 cycle for L1 hits, 10 cycles for L2 hits, 85 cycles for local memory, 180–260 cycles for remote memory, and 210–310 cycles for cache-to-cache transfers. We also briefly summarize experimental results using a real machine (Convex Exemplar), with more detail in our extended report [12].

Processor parameters	
Clock rate	500 MHz
Fetch rate	4 instructions/cycle
Instruction window	64 instructions in-flight
Memory queue size	32
Outstanding branches	16
Functional unit count	2 ALUs, 2 FPUs, 2 address units
Functional unit latencies (cycles)	1 (addr. gen., most ALU), 3 (most FPU), 7 (int. mult./div.), 16 (FP div.), 33 (FP sqrt.)
Memory hierarchy and network parameters	
L1 D-cache	16 KB, direct-mapped, 2 ports, 10 MSHRs, 64-byte line
L1 I-cache	16 KB, direct-mapped, 64-byte line
L2 cache	64 KB (1 MB for Em3d), 4-way as- sociative, 1 port, 10 MSHRs, 64-byte line, pipelined
Memory banks	4-way, permutation interleaving
Bus	167 MHz, 256 bits, split transaction
Network	2D mesh, 250MHz, 64 bits, flit de- lay of 2 network cycles per hop

Table 1. Base simulated configuration.

4.2. Evaluation Workload

We evaluate our clustering transformations using a latency-detection microbenchmark and five scientific applications. Table 2 summarizes the evaluation workload for the simulated system. The number of processors used for the simulated multiprocessor experiments is based on application scalability, with a limit of 16. The input sizes and processor counts for experiments on the real machine are reported in [12]. Each code is compiled with the Sun SPARC SC4.2 compiler, using the `-xO4` optimization level. We incorporate miss clustering transformations by hand, following the algorithms presented.

Latbench is based on the `lat_mem_rd` kernel of `lm-bench` [8]. `lat_mem_rd` sees inner-loop address recurrences from pointer-chasing. Latbench wraps this loop in an outer loop that iterates over different pointer chains, with no locality in or across chains. The pseudocode, given below, shows code added for Latbench in sans-serif.

```

for (j=0;j<N;j++){
  p = A[j];
  for(i=0;i<I;i++)
    p = p→next // serialized misses
  USE(p)       // keeps p live

```

Latbench is clustered with unroll-and-jam. As in `lat_mem_rd`, looping overhead is minimized by unrolling the innermost loops to include 1000 pointer dereferences in each loop body, for both the base and clustered versions.

Em3d is a shared-memory adaptation of a Split-C application [5], and is clustered using unroll-and-jam. This code has both cache-line and address dependences, but only cache-line recurrences. The dominant loop nest has variable

Microbenchmark	Input Size	Procs.
Latbench	6.4M data size	1
Application	Input Size	Procs.
Em3d	32K nodes, deg. 20, 20% rem.	1,16
Erlebacher	64x64x64 cube, block 8	1,16
FFT	64K points	1,16
LU	256x256 matrix, block 16	1,8
Mp3d	100K particles	1,8

Table 2. Data set sizes and number of processors for simulation experiments.

inner-loop length, so only the minimum length seen in the unrolled copies is fused. Each copy completes its remaining length separately. We assumed that the outer loop was explicitly identified as parallel to enable transformation despite Em3d’s pointer references. Because of its larger working set, Em3d is simulated with 1 MB L2 caches.

Erlebacher is a shared-memory port of a program by Thomas Eidson at the Institute for Computer Applications in Science and Engineering (ICASE). FFT and LU are from SPLASH-2 [20]. For better load balance, LU is modified slightly to use flags instead of barriers. These three regular codes see only cache-line recurrences. Each is clustered with unroll-and-jam and postlude interchanging.

Mp3d is an irregular, asynchronous, communication-intensive SPLASH code [16]. To eliminate false-sharing, key data structures were padded to a multiple of the cache line size. To reduce true-sharing and improve locality, the data elements were sorted by position in the modeled physical world [9]. Mp3d has no recurrences, but sees poor miss clustering because of large loop bodies. Thus, inner-loop unrolling and aggressive scheduling can provide clustering here, as discussed in Section 3.3. We assumed that the dominant move loop was explicitly marked parallel.

5. Experimental Results

5.1. Performance of Latbench

The base Latbench of Section 4.2 exposes an average miss latency of 171 ns on the simulated system (identical to `lat_mem_rd`). Clustering drops the average exposed latency to 32 ns, a speedup of 5.34X. On the Convex Exemplar, clustering reduces the average exposed latency from 502 ns to 87 ns, for a speedup of 5.77X.

These results indicate the potential gains from memory parallelism transformations, but also indicate some bottlenecks, since the speedups are less than 10 (the number of simultaneous misses supported by each processor). Our more detailed statistics for the simulated system show that clustering increases contention, increasing average total latency to 316 ns (from address generation to completion). Further, bus and memory bank utilization both exceed 85% after clustering. Thus, a further increase in speedup would

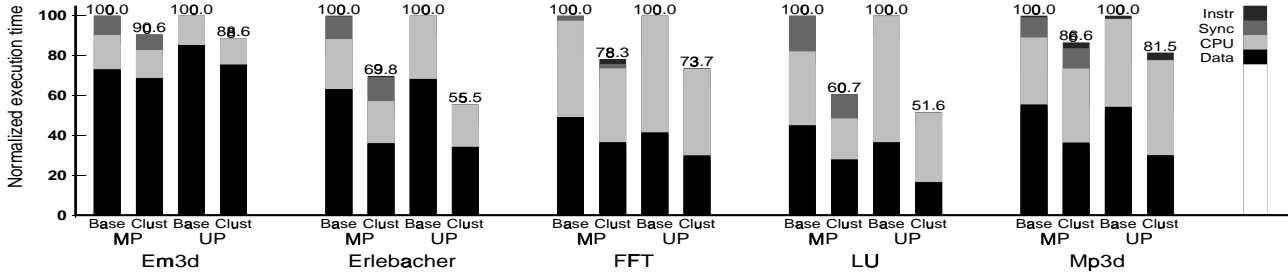


Figure 3. Impact of clustering transformations on application execution time.

require greater bandwidth at both the bus and the memory.

5.2. Impact on Application Performance

Figure 3 shows the impact of the clustering transformations on application execution time for the base simulated system. The graph shows multiprocessor and uniprocessor experiments (MP/UP) before and after clustering (Base/Clust), normalized to the given application and system size without clustering. For analysis, execution time is categorized into data memory stall, CPU, synchronization stall, and instruction memory stall times, following the conventions of previous work (e.g., [14]). Since writes can retire before completing and read hits are fast, nearly all data memory stalls stem from reads that miss in the L2 cache.

Overall, the clustering transformations studied provide from 9–39% reduction in multiprocessor execution time for these applications, averaging 23%. The multiprocessor benefits in Erlebacher and Mp3d come almost entirely from reducing the memory stall time. (Mp3d sees some CPU degradation because of no scalar replacement or pipeline improvement and slightly worse return-address prediction.) Em3d, FFT, and LU see benefits split between memory stall time and CPU time; unroll-and-jam helps the CPU component through better functional unit utilization and through scalar replacement (in FFT and LU). By speeding up the data producers in LU, the clustering transformations also reduce the synchronization time for data consumers. Our more detailed statistics show that the L2 miss count is nearly unchanged in all applications, indicating that locality is preserved and that scalar replacement primarily affects cache hits. All applications see more multiprocessor execution time reduction from the newly exposed benefits in read miss clustering than the previously studied benefits in CPU time.

The uniprocessor sees slightly larger overall benefits from the clustering transformations, ranging from 11–48% (average 30%). The speedup of data memory stalls is greater in the uniprocessor than in the multiprocessor, as the uniform latency and bandwidth characteristics of the uniprocessor better facilitate overlap. However, since the uniprocessor typically spends a smaller fraction of time in data memory stalls, the transformations’ benefits for FFT and LU are predominantly in the CPU component.

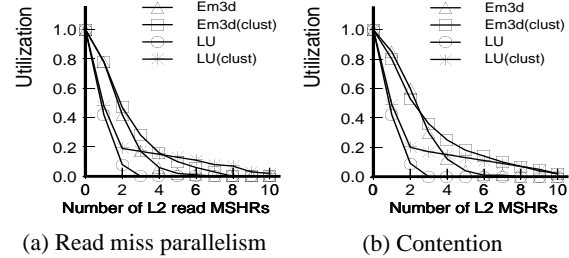


Figure 4. Factors shaping memory parallelism (read L2 MSHR utilization) and contention (total L2 MSHR utilization).

To represent the growing processor-memory speed gap, we simulated a system with 1 GHz processors and all memory and interconnect parameters identical (in ns or MHz) to the base. The total execution time reductions are similar (10–36% in the multiprocessor, averaging 24%; 12–47% in the uniprocessor, averaging 34%). However, the larger fraction of memory stall time in these systems allows memory parallelism to provide more of the total benefits than in the base. Thus, targeting memory parallelism becomes more important for such potential future configurations.

All simulation experiments show few instruction memory stalls. Thus, the code added by our transformations does not significantly impact I-cache locality for these loop-intensive codes.

We also performed experiments on a Convex Exemplar, using larger input sizes appropriate for the real machine [12]. Each HP PA-8000 processor in the Exemplar supports 10 simultaneous misses. The clustering transformations give the Exemplar 9–34% reductions in application execution time for the multiprocessor and uniprocessor experiments. Our extended report provides more details [12].

5.3. Memory Parallelism and Contention

The MSHR utilization graphs of Figure 4 depict the sources of memory parallelism and contention for the multiprocessor runs of Em3d and LU, the two extreme applications with regard to improvement from the transformations. Figure 4(a) indicates read miss parallelism, showing the fraction of total time for which at least N L2 MSHRs are

occupied by read misses for each possible N on the X axis. The clustering transformations only slightly improve read miss parallelism for Em3d, since Em3d's irregular accesses give even the base version some clustering. In contrast, the transformations convert LU from a code that almost never had more than 1 outstanding read miss to one with 2 or more outstanding read misses 20% of the time and up to 9 outstanding read misses at times.

Figure 4(b) shows the total L2 MSHR utilization, including both reads and writes. This indicates contention, measuring how many requests use the memory system at once. Both Em3d and LU see curves similar to their read miss curves, indicating that contention in these applications comes primarily from reads. Thus, for these applications, any negative impact from increased contention is offset by the performance benefits of read miss parallelism.

6. Conclusions and Future Work

This study finds that code transformations can improve memory parallelism in systems with out-of-order processors. We adapt compiler transformations known for other purposes to the new goal of memory parallelism. Our experimental results show substantial improvements in memory parallelism, thus hiding more memory stall time and reducing execution time significantly. As memory stalls become more important (e.g., multiprocessors or future systems with greater processor-memory speed gaps), more execution time reductions come from the transformations' newly exposed benefits in memory stall time than their previously studied benefits in CPU time.

We can extend this work in several ways. For example, we can seek to resolve memory-parallelism recurrences for unnested loops by fusing otherwise unrelated loops. We are also investigating the interactions of miss clustering with software prefetching, as their different approaches to latency tolerance allow each to provide distinct benefits.

Acknowledgments

We thank Vikram Adve, Keith Cooper, Chen Ding, Ken Kennedy, John Mellor-Crummey, Partha Ranganathan, and Willy Zwaenepoel for valuable comments on this work.

References

- [1] F. E. Allen and J. Cocke. A Catalogue of Optimizing Transformations. In *Design and Optimization of Compilers*, pages 1–30. Prentice-Hall, 1972.
- [2] D. Callahan et al. Estimating Interlock and Improving Balance for Pipelined Machines. *Journal of Parallel and Distributed Computing*, 5(4):334–358, Aug. 1988.
- [3] S. Carr. Combining Optimization for Cache and Instruction-Level Parallelism. In *Proc. of the Conf. on Parallel Architectures and Compilation Techniques*, 1996.
- [4] S. Carr and K. Kennedy. Improving the Ratio of Memory Operations to Floating-Point Operations in Loops. *ACM Trans. on Programming Languages and Systems*, 16(6):1768–1810, Nov. 1994.
- [5] D. E. Culler et al. Parallel Programming in Split-C. In *Proc. of Supercomputing*, 1993.
- [6] D. R. Kerns and S. J. Eggers. Balanced Scheduling: Instruction Scheduling When Memory Latency is Uncertain. In *Proc. of the Conf. on Programming Language Design and Implementation*, 1993.
- [7] J. L. Lo and S. J. Eggers. Improving Balanced Scheduling with Compiler Optimizations that Increase Instruction-Level Parallelism. In *Proc. of the Conf. on Programming Language Design and Implementation*, 1995.
- [8] L. McVoy and C. Staelin. Imbench: Portable Tools for Performance Analysis. In *Proc. of the USENIX Technical Conf.*, 1996.
- [9] J. Mellor-Crummey et al. Improving Memory Hierarchy Performance for Irregular Applications. In *Proc. of the 13th Int'l Conf. on Supercomputing*, 1999.
- [10] T. Mowry. *Tolerating Latency through Software-controlled Data Prefetching*. PhD thesis, Stanford University, 1994.
- [11] A. Nicolau. Loop Quantization or Unwinding Done Right. In *Proc. of the 1st Int'l Conf. on Supercomputing*, 1987.
- [12] V. S. Pai and S. Adve. Code Transformations to Improve Memory Parallelism. Technical Report ECE-9909, Rice University, Sep. 1999. <http://www.ece.rice.edu/~rsim/pubs/TR9909.ps>.
- [13] V. S. Pai and S. Adve. Improving Software Prefetching with Transformations to Increase Memory Parallelism. Technical Report ECE-9910, Rice University, Nov. 1999. <http://www.ece.rice.edu/~rsim/pubs/TR9910.ps>.
- [14] V. S. Pai et al. The Impact of Exploiting Instruction-Level Parallelism on Shared-Memory Multiprocessors. *IEEE Trans. on Computers*, 48(2):218–226, Feb. 1999.
- [15] V. S. Pai et al. *RSIM Reference Manual, Version 1.0*. Electrical and Computer Engineering Department, Rice University, Aug. 1997. Technical Report 9705.
- [16] J. P. Singh et al. SPLASH: Stanford Parallel Applications for Shared-Memory. *Computer Architecture News*, 20(1):5–44, Mar. 1992.
- [17] J. E. Smith and A. R. Pleszkun. Implementing precise interrupts in pipelined processors. *IEEE Trans. on Computers*, C-37(5):562–573, May 1988.
- [18] G. S. Sohi. High-Bandwidth Interleaved Memories for Vector Processors – A Simulation Study. *IEEE Trans. on Computers*, 42(1):34–44, Jan. 1993.
- [19] M. E. Wolf and M. S. Lam. A Data Locality Optimizing Algorithm. In *Proc. of the Conf. on Programming Language Design and Implementation*, 1991.
- [20] S. C. Woo et al. The SPLASH-2 Programs: Characterization and Methodological Considerations. In *Proc. of the 22nd Annual Int'l Symp. on Computer Architecture*, 1995.