

Computation Regrouping: Restructuring Programs for Temporal Data Cache Locality

Venkata K. Pingali
PCEN
Information Sciences Institute
University of Southern California
Los Angeles, CA 90292
pingali@isi.edu

Wilson C. Hsieh
School of Computing
University of Utah
Salt Lake City, UT 84112
wilson@cs.utah.edu

Sally A. McKee
School of Computing
50S Central Campus Drive, Room 3190
University of Utah
Salt Lake City, UT 84112
sam@cs.utah.edu

John B. Carter
School of Computing
University of Utah
Salt Lake City, UT 84112
retrac@cs.utah.edu

ABSTRACT

Data access costs contribute significantly to the execution time of applications with complex data structures. As the latency of memory accesses becomes high relative to processor cycle times, application performance is increasingly limited by memory performance. In some situations it may be reasonable to trade increased computation costs for reduced memory costs. The contributions of this paper are three-fold: we provide a detailed analysis of the memory performance of a set of seven, memory-intensive benchmarks; we describe *Computation Regrouping*, a general, source-level approach to improving the overall performance of these applications by improving temporal locality to reduce cache and TLB miss ratios (and thus memory stall times); and we demonstrate significant performance improvements from applying Computation Regrouping to our suite of seven benchmarks. With Computation Regrouping, we observe an average speedup of 1.97, with individual speedups ranging from 1.26 to 3.03. Most of this improvement comes from eliminating memory stall time.

1. INTRODUCTION

This research was sponsored in part by National Science Foundation award 9806043, and in part by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under agreement number F30602-98-1-0101 and DARPA Order Numbers F393/00-01 and F376/00. The views and conclusions contained herein are those of the authors, and should not be interpreted as representing the official policies or endorsements, either express or implied, of NSF, DARPA, AFRL, or the US Government. This research was conducted while the first author was a student at the University of Utah.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICS'02, June 22-26, 2002, New York, New York, USA.
Copyright 2002 ACM 1-58113-483-5/02/0006 ...\$5.00.

Many applications with performance constraints contain complex data structures, large working sets, and access patterns with poor memory reference locality. Signal processing, databases, and CAD tools are but a few examples of such applications for which memory stalls contribute significantly to execution time. The memory stalls of the applications considered here account for 40% to 82% of execution time. Worse, the contribution as a fraction of total execution time will increase as the performance gap between processors and memory continues to grow. Reducing this memory stall time improves processor resource utilization and will allow an application's performance to scale with improvements in processor speed and technology.

Application memory performance is the focus of much ongoing research [5, 7, 17, 18, 19, 31] for both regular as well as complex pointer-based datastructures. While effective restructuring methods have been demonstrated for modest-sized applications with regular and irregular datastructures [7, 11, 24, 31] memory performance issues of large pointer-based applications, which is our main focus, has received relatively less attention. Most existing approaches [9, 28, 29] try to improve spatial locality through controlled layouts. However, optimal layouts are hard to achieve in applications with complex or dynamic datastructures. In some other cases the complexity in the access patterns renders them ineffective. Our approach tries to improve temporal locality, and is complementary to most spatial locality and local optimizations. Section 6 discusses related work in more detail.

In this paper, we first identify reasons for poor memory performance by qualitatively characterizing benchmark access patterns. We then propose *Computation Regrouping*, extending application control and data structures to modify the execution order of computations. This restructuring reduces the impact of high memory latency by grouping computations based on the data objects they access. The "batched" computations can be executed with higher temporal data locality, thereby improving cache performance. Applications that can benefit from this work are those for

which processor-memory bottlenecks account for a large portion of the runtime. We find Computation Regrouping to be quite effective on our seven memory-bound benchmarks, which are of varying complexity, and which span a number of application domains. For these initial results, optimizations are implemented by hand. The additional programmer effort appears worthwhile for programs that are difficult to optimize by compiler methods: we observe an average application speedup of 1.97, with individual speedups ranging from 1.26 to 3.03. We achieve these speedups by reducing memory stall times by 26-85%.

Section 2 describes in detail our benchmarks and their high-level characteristics. Section 3 presents observations relating the memory performance to the program structure. Section 4 discusses our proposed transformations, and Section 5 presents results obtained on applying these transformations. Section 6 discusses related approaches, and Section 7 summarizes our results and identifies some open problems.

2. BENCHMARK ANALYSIS

We select our benchmarks based on the memory performance of publicly available, reference implementations on large inputs, considering only applications spending 25% or more of their execution time servicing secondary data cache and TLB misses. We use carefully optimized versions of each benchmark as baselines, implementing function-call and layout-related optimizations where appropriate (for instance, we eliminate MST [6] from consideration, since it becomes compute-bound after applying such changes). Applications not meeting this 25% threshold are likely to continue to be compute-bound on next-generation systems, and the small potential performance gains do not justify either the programming effort or run-time overheads of our approach. Many common architecture and compiler benchmarks are therefore excluded. For instance, floating-point costs are sufficiently high to hide the poor memory performance of Barnes-Hut [6], and so even though it could potentially be memory-bound on a different processor configuration, we exclude this benchmark.

Section 5 describes the SGI Power Onyx system on which we perform all our experiments, including those to identify an appropriate benchmark suite. We experiment with a range of inputs to determine the sensitivity of each application to input parameters, and we examine the source code to ensure that poor memory performance is due to access patterns (such as indirect accesses and strided accesses), and not an artifact of our experimental setup.

The seven benchmarks cover a range of application domains (scientific codes, hardware design tools, graphics, and databases), code sizes, and complexity. Table 1 presents cache hit ratios for a sample input and shows how the execution time scales with input size. The table also presents the estimated memory stall time. The specific location of the misses that cause these stalls varies with application. However, each benchmark studied here is built around a core data structure, and we find that performance issues can be traced back to specific operations over this data structure. We discuss the benchmark specifics below. Source code sizes vary from 280 lines of C in IRREG to 60K lines in CUDD. Data structure complexity varies from a simple array in IRREG to a complex, directed acyclic graph (DAG) based on

hash tables in CUDD. We next describe each benchmark and its baseline performance in detail.




Figure 1: The R-Tree from Data Management

R-Tree. R-Trees, originally proposed by Guttman [15], are known to perform well when the number of data dimensions is low. Our R-Tree reference implementation is the DARPA DIS Data Management Benchmark [26]. The R-Tree stores spatial information of objects such as airports and restaurants. Queries usually include spatial constraints such as “all printers in this building” and “landmarks within five miles of the airport”. To execute these operations efficiently, the R-Tree groups objects based on their spatial attributes. Figure 1 shows the structure of the DIS implementation of the R-Tree, a height-balanced tree with a branching factor usually between two and fifteen. Each subtree is associated with a bounding box in four dimensional space, or *hypercube key*, that encompasses the bounding boxes of all the nodes within the subtree. Each internal node contains a set of entries implemented as a linked list, where each list member is a (*child, hypercube key*) pair corresponding to one child subtree of the node. Entries at leaf nodes point to the data objects. Simple tree balancing mechanisms are invoked during node merging and splitting.

Valid matches include all objects whose hypercube keys overlap the search hypercube key. A search key can potentially overlap with multiple objects’ boxes, and hence may trigger multiple tree traversals along different paths. Query and delete tree operations require searching, whereas insert operations access nodes along a single path from root to leaf. Poor performance arises from the many internal nodes touched per search — often up to 50% of all nodes in the tree. For a sample run performing 12K query operations on a tree of 67K nodes, the L1D and L2 cache hit ratios are 97.11% and 57.63%, respectively. Execution time scales almost linearly with the size of the tree and the number of tree operations. R-Tree optimizations and performance are discussed in greater detail elsewhere [27].

```
for (i = 0; i < n_edges ; i++) {
    n1 = left[i] - 1;
    n2 = right[i] - 1;
    rr = (x[n1] - x[n2]) * 0.25;
    y[n1] = y[n1] + rr;
    y[n2] = y[n2] - rr;
}
```

Figure 2: IRREG Main Loop

IRREG. The IRREG kernel is an irregular-mesh PDE solver taken from computational fluid dynamics (CFD) codes [16]. The input mesh’s nodes correspond to bodies (molecules) and its edges to interactions, represented by arrays x and y , respectively, in Figure 2. The forces felt by the bodies are updated repeatedly throughout execution. The largest standard input mesh, *MOL2*, has 442,368 nodes and 3,981,312

Benchmark	Lines of Code	Sample Input	Hit Ratio (%)		Stall Time (%)	Execution Time Scaling
			L1D	L2		
R-Tree	7K	dm23.in	95.4	49.6	71	linearly with tree size
IRREG	300	MOL2, 40 iterations	84.5	54.5	74	linearly with number of edges
Ray Trace	5K	balls 256 × 256	95.8	70.1	58	linearly with viewing window size and input scene
FFTW	35K	10K × 32 × 32	92.4	60.0	65	linearly with dimensions
CUDD	60K	C3540.blif	84.1	45.9	69	scaling properties unclear
EM3D	700	128K, 10, 1	85.4	47.9	66	linearly with nodes, sub-linearly with degree
Health	600	6, 500	71.9	68.6	70	exponentially with depth

Table 1: Benchmark Characteristics.

edges. When the loop in Figure 2 is run on it for 40 iterations, the L1D and L2 hit ratios are 81.81% and 53.86%. As in EM3D, described below, indirect accesses to remote nodes (array x) are expensive. For large data arrays with adequately random index arrays, execution time scales with the number of edges and iterations.

Ray Trace. In Ray Trace, part of the Data Intensive Systems (DIS) benchmark suite [26], rays are emitted from a viewing window into a 3D scene, and each ray is traced through its reflections. A ray is defined by an origin (position) in the viewing window and a direction. The algorithm checks each ray for intersection with all scene surfaces, determining the surface closest to the emission source. Scene objects are allocated dynamically and stored in a three-level hierarchy. The number of objects is fixed per scene, whereas viewing window size, and hence the number of rays, is specified by the user.

When the object structure is significantly larger than L2 cache size, data objects are evicted before reuse, and successive rays suffer the same (large) number of misses. We report on the *balls* input, which is the only DIS-specified input that has high miss ratios. It has 9,828 surfaces and 29,484 vertices, and has a memory footprint of about 6.5MBytes, or roughly three times cache size. The L1D and L2 cache hit ratios for a window size of 1024×1024 are 97.6% and 70.15%. Almost the entire object data structure is accessed in processing each ray, and strided accesses and pointer chasing are the dominant access patterns. Execution time scales almost linearly with viewing window (*i.e.*, the number of emitted rays) and input scene size.

FFTW. The highly tuned DIS FFTW benchmark [26] outperforms most other known FFT implementations [12]. The core computation consists of alternating passes along the X, Y, and Z dimensions of a large array. Accesses along the Z dimension account for 50-90% of the total execution time of FFT, depending on the input size. The Z stride is typically much larger than a virtual memory page, and when this dimension is sufficiently large, there is little or no reuse of cache lines before they are evicted. Accesses during this phase of the program suffer very high cache and TLB miss ratios. For an input array of size $10240 \times 32 \times 32$, with a memory footprint of 320MBytes, the L1D and L2 hit ratios are 97.01% and 64.62%, respectively. Execution time scales linearly with increasing input dimensions.

CUDD. The CUDD Binary Decision Diagram (BDD) manipulation package [34] is widely used in the hardware design community. Based on Shannon’s boolean-expression evaluation rules, BDDs compactly represent and manipulate

logic expressions. The core structure is a large DAG whose internal nodes represent boolean expressions corresponding to the subgraphs of lower nodes. Each DAG level is associated with a boolean variable, and the large structure and nature of the manipulations requires that each level of internal nodes be stored in a separate hash table. The DAG depth is the number of boolean variables. Since variable ordering is important to obtaining a compact structure, the package provides several methods for dynamically changing the ordering. The main operation among all these methods is a variable swap with three steps: extracting nodes from a hash table, updating a second hash table, and garbage-collecting nodes. Each of these touches many nodes, causing similar cache behavior to FFT. For a sample input circuit, *C3540.blif*, with random dynamic reordering of variables, the L1D and L2 hit ratios are 94% and 45%, respectively.

EM3D. EM3D models the propagation of electromagnetic waves through objects in three dimensions [10]. The core data structure is a bipartite graph in which nodes represent objects and edges represent the paths through which the waves propagate. The user specifies the number of nodes and degree at each, and the program randomly generates graph edges to fit these constraints. The primary computation updates each node’s state based on those of its neighbors. We use a slightly constrained version of EM3D in which the number of nodes is a power of two and the out-degree is fixed (arbitrarily) at 10. For an input size of 128K nodes and a degree of 10, the total memory footprint is about 30MBytes. The corresponding L1D and L2 hit ratios are 94.96% and 47.77%. Our experimentation confirms Culler *et al.*’s observation that the primary performance determinant is the cost of remote node accesses (from nodes in one half of the bipartite graph to nodes in the other half) [10]. Execution time scales linearly with the number of nodes, and scales sub-linearly with increasing degree.

Health. From the Olden benchmark suite [6], Health simulates the Columbian health care system. The core data structure is a quad-tree with nodes representing hospitals at various logical levels of capacity and importance. Each node is associated with personnel and three lists of patients in various stages of treatment. Patients are first added to the waiting list, and waiting time depends on the numbers of personnel and patients. When hospital personnel are available, a patient enters the treatment phase, which takes a fixed amount of time. After treatment, the patient exits the system or goes to a bigger hospital where the process repeats. Each simulation cycle visits every quad-tree node, updating the patient lists. For a sample input of depth six,

Benchmark	Critical Working Set	Estimated Threshold	Logical Operation
R-Tree	Tree size	15K nodes	one tree operation (insert, delete, or query)
Ray Trace	3D scene	—	scan of the input scene by one ray
FFTW	Cache line-size \times dimension	Y or Z dim $>$ 16K	one column walk of a 3D array
IRREG	Array size	256K	group of accesses to a set of remote nodes
CUDD	$2 \times$ hash-table size	—	one variable swap
EM3D	$(4 \times \text{degree} + 12) \times \text{nodes}$	40K nodes (degree=10)	group of accesses to a set of remote nodes
Health	Quad-tree size + lists size	3K (listsize=5)	simulation of one time step


Table 2: Benchmark Logical Operations

after 500 simulation cycles the L1D and L2 hit ratios are 87.47% and 68.65%, respectively. Execution time increases linearly with tree size and number of simulation cycles.


3. LOGICAL OPERATIONS

Computation Regrouping is targeted at applications whose program structure is expressible at the granularity of *logical operations*, which we define to be short streams of nearly-independent computations. Table 2 presents the list of logical operations in our benchmarks. Applications typically contain many instances of these logical operations, but are not intentionally programmed that way. Many data objects are accessed within each logical operation, but with little or no reuse. In contrast, there is significant reuse across logical operations. The key insight behind Computation Regrouping is that application structuring in terms of logical operations results in poor memory performance due to the large temporal separation between computations accessing the same data: bringing these computations temporally closer improves memory performance. Computation Regrouping identifies a set of data objects that fit into L2 cache, and then reorganizes the computations from multiple logical operations to maximize utilization of these cached objects. In this context, we define the *critical working set* to be the average size of the data objects accessed by a single logical operation. Table 2’s *Estimated Threshold* column gives estimated minimum parameter values for the critical working set to exceed L2 cache size (given a 2MByte cache and a 128Byte line size), the point at which application performance begins to degrade significantly.

Consider an R-Tree on which we perform insert, delete and query operations requiring tree traversals whose accesses are data-dependent and unpredictable at compile time. Each traversal potentially contains many hypercube-key comparisons. In particular, query and delete operations traverse multiple tree paths and have large working sets. The average memory cost of the comparisons dominates the processor cycle cost. To better visualize R-Tree’s memory behavior, consider the graphs in Figure 3. Figure 3(a) shows a small window of queries versus node identifiers in our R-Tree baseline on a test input. Each dot represents an access made to the node identified by the x -coordinate by the query identified by the y -coordinate. Each row of dots identifies all accesses by a single query, and each column identifies queries accessing a given node. This simple plot identifies two important characteristics of the R-Tree benchmark. First, the set of nodes accessed by successive queries rarely overlaps — there are significant gaps along the vertical axis. In gen-



(a) Access Matrix for Queries versus Nodes



(b) Distribution of the Number of Nodes Accessed by Queries

Figure 3: R-Tree Performance for about 70K Nodes

eral, queries for which there is significant overlap in accessed data are widely separated in time. Second, the set of unique nodes accessed between successive visits to a given node is the union of the rows of dots whose row (query) number lies between those of the two dots, and the number of unique other nodes accessed is the size of this union. Figure 3(b) shows the distribution of this union’s size with query number. In our observations, the size of this union is typically large — sometimes including half the total tree nodes. With large trees compared to cache sizes, and with the high number of nodes accessed (with repeat accesses to a node spaced widely in time), nodes are evicted from cache before they can be reused.

R-Tree’s access behavior is representative of any application consisting of many logical operations where each of those potentially accesses many data objects, but with few accesses to any given object and little computation performed per object. Logical operations execute in a program-dependent order, and each runs to completion before another is started. Critical working set sizes exceed cache size. This set of conditions gives rise to poor temporal locality, preventing applications from exploiting data reuse even when it exists. This is strikingly visible in the FFTW benchmark, which consists of a series of column walks having 100% overlap between cache lines touched, but for which the reuse cannot be exploited. Our other benchmarks behave similarly.

Computing the cost of the memory accesses in any computation requires determining the number of computations within a logical operation, the relationships among computations in different logical operations, and the cache sizes. This information is usually unavailable to the compiler, and thus it is difficult to design generally effective static techniques to take advantage of the data reuse we target here.

Technique	Data Structure Modifications	Control Structure Modifications	Generality
Early Execution	mid	high	mid
Deferred Execution	high	high	high
Computation Merging	low	low	low
Filtered Execution	low	mid	low

Table 3: Characteristics of Regrouping Techniques

The approaches in the next section represent preliminary steps towards developing widely applicable techniques, but much work remains in terms of producing tools to automate or semi-automate their application.

4. COMPUTATION REGROUPING

This section presents four simple optimization techniques that realize Computation Regrouping, summarizes their properties, and discusses the tradeoffs involved.

4.1 Transformations

We present four related implementation techniques: *early execution*, *deferred execution*, *computation merging*, and *filtered execution*. The ideas behind each transformation are similar, but the implementations, tradeoffs, and effects on application source are sufficiently different that we treat each separately. Applying the techniques requires identifying computations suitable for regrouping, choosing application extensions to implement the regrouping, and choosing mechanisms to integrate results from regrouped computations and non-regrouped computations. Table 3 summarizes the general costs and applicability of each technique.

Early Execution. Early execution brings forward future computations that touch data accessed by the current computation. In some instances, early execution can be viewed as a generalization of traditional tiling approaches. For example, FFT’s column walks access elements represented as 16Byte complex numbers. Assuming 128Byte cache lines, each line loaded during one walk contains data accessed by exactly seven other column walks. Serially executing the walks fails to exploit this overlap. Early execution performs computations from all eight walks on a resident cache line. Identifying future computations and storing their results is straightforward. The optimized CUDD moves forward the reference counting for garbage collection, performing it along with data-object creation and manipulation. The ideas bear some similarity to incremental garbage collection, as in Appel, Ellis, and Li’s virtual memory-based approach [1], but we optimize for a different level of the memory hierarchy, and do not strive for real-time performance. In both FFT and CUDD, code changes are largely isolated to control structures. Data structure changes are minimal, and integration of partial results is straightforward.

Deferred Execution. When there is uncertainty in the number and/or order of logical operations that will be performed, computations can be explicitly or implicitly delayed (as in *lazy evaluation*) until sufficiently many accumulate to amortize the costs of fetching their data. The deferred computations can then be executed with high temporal locality. In the R-Tree from Figure 1, deferring queries, inserts, and deletes requires changes to both control and data structures: we associate a small operation queue with a subset of tree nodes, and we enqueue tree operations until some dependency or timing constraint must be satisfied, at which time we “execute” the queues. Since all queued queries access nodes within the subtree, their accesses have good cache lo-

```
for (i = 0; i < max; i++) {
    sum += A[ix[i]];
}
```

Figure 4: Indirect Access Original Source Code

cality. An operation might reside in multiple queues over the period of its execution, and so we trade increased individual operation latency for increased throughput. Synchronization overhead for queries is insignificant, since they are largely read-only. Supporting deletes and inserts gives rise to larger overheads, due to additional consistency checks. Deferred execution can also be applied to CUDD, for which code changes are largely to the control structure, and thus partial result integration is simpler.

Computation Merging. Computation merging is a special case of deferred execution in which the deferred computations are both cumulative and associative. Two or more deferred computations are replaced by one that achieves the end effect. For instance, the most common operation in Health increments a counter for each patient in a waiting list at each node in the tree. These increments can be deferred until an update is forced by the addition or deletion of list elements. This optimization is more specialized than deferred execution, and therefore less generally applicable. However, computation merging can be implemented efficiently, and it improves both computation and memory performance. EM3D benefits from similar optimizations.

```
#define SHIFT(x) (x + window_size)
...
iters = ...
window_size = N/iters;
for (win = start; iters > 0 ; win = SHIFT(win), iters--) {
    for (i = 0; i < imax; i++) {
        if (ix[i] >= win && ix[i] < (win + window_size))
            sum += A[ix[i]];
    }
}
```

Figure 5: Indirect Access Filtered Execution Code

```
for ( k = 0 ; k < n_nodes; k += blocksize ) {
    int n1, n2;
    for (i = 0; i < n_edges; i++) {
        n1 = left(i) - 1;
        n2 = right(i) - 1;
        if ( n2 < k ||| n2 >= (k + blocksize) )|
            continue;
        rr = (x[n1] - x[n2]) * 0.25;
        y[n1] = y[n1] + rr;
        y[n2] = y[n2] - rr;
    }
}
```

Figure 6: Filtered Execution Applied to IRREG

Filtered Execution. Filtered Execution places a sliding window on a traversed data structure, allowing accesses only to the part of the structure visible through this window. Figures 4 and 5 show this for a simple, indirect-access kernel. In each iteration, accesses to locations outside the window are deferred. Like computation merging, filtered execution is a special case of deferred execution in which the deferring is achieved implicitly by the modified loop control structure. The window improves temporal locality in each iteration, but at the cost of added computational and source-code complexity. Figure 6 shows the transformation




Figure 7: Speedup vs. Window Size for IRREG

applied to the IRREG main kernel from Figure 2. Note that increased computational costs may be higher than the savings from reduced cache misses; Figure 7 illustrates the cost/performance tradeoffs for varying window sizes (with respect to data array sizes) in IRREG. For appropriately large windows, the reduction in miss stalls dominates control overhead, but for windows smaller than a (cache-size determined) threshold value, control overhead dominates. Ray Trace and EM3D behave similarly.

4.2 Tradeoffs

Regrouping can increase complexity. First, control and data structure extensions make it more difficult to code and debug the application. The extent of the modifications depends on the application being optimized. Second, reordering computations can reorder the output. If the correctness of the application is defined in terms of the program output, then either the definition of the correctness must be modified, or some data structure extensions must be incorporated to ensure correct output order. Third, regrouping is a general approach, and different applications require different implementation strategies. Choosing appropriate values for optimization parameters requires experimentation (*e.g.*, choosing queue length in R-Trees, or filter window sizes).

The increased computational overhead from regrouping is justified in many cases. Improvements in overall execution time can be significant, and most of these come from reducing memory stall times. Few generic techniques exist to optimize for memory performance in complex applications. In our experience, the control and data structure changes required are usually small compared to overall code sizes. Optimizations are architecture-independent, modulo the cache model, and we expect improvements to hold across generations of processor technology. Finally, most existing compiler-based memory optimizations cannot be directly applied to complex data structures, but regrouping achieves some of the desired effects. For example, early execution for FFTW simulates the effect of tiling. Overall, Computational Regrouping often represents a reasonable tradeoff between added overhead and improved memory performance.

5. RESULTS

We study the impact of Computation Regrouping on the benchmarks discussed in Section 2. We incorporate the appropriate regrouping optimizations into the reference implementations, and execute them all on the same sample inputs. We perform all experiments on a 14-processor SGI Power Onyx with 194MHz MIPS R10000 processors and 4GBytes of main memory. Each processor has 32KByte instruction and data L1 caches, and a 2MByte unified, write-back, two-way set-associative L2 cache. The R10000 processor is four-way superscalar, and implements dynamic in-

struction scheduling, out-of-order execution, and non-blocking caches. We use the SpeedShop performance tool Suite [33] and the perfex [37] command line interface to processor performance counters. We compile with the MIPSpro C compiler at the -O3 level of optimization. To remove artifacts due to OS scheduling or page allocation and to cover a variety of system load conditions, we execute the baseline and optimized versions of the applications ten times over a period of few days. Data presented represent the mean values obtained for these ten experiments.

Table 4 presents a summary of the specific optimization techniques used and their high-level impact on each application. Complex applications such as CUDD can benefit from multiple regrouping techniques. The table shows the extent of modifications performed to the application. We find that they are small compared to the overall application size. In case of FFTW, the same 40-line modification had to be replicated in 36 files, which is reflected in the large change reported. The table also presents details of the number of instructions issued and graduated. We note that in some cases such as IRREG and R-Tree the number of instructions issued and graduated is higher in the restructured application compared to the original application. This fact demonstrates the tradeoff identified in Section 4.2. Increased computation can result in improved performance.

5.1 Performance Summary

Table 5 summarizes the effectiveness of Computation Regrouping for improving the memory performance of our benchmark applications. The mean speedup is 1.97, although individual speedups vary significantly. Improvements are dependent on inputs, and the speedups shown represent points in a spectrum of (occasionally negative) values. Nonetheless, our experiences have been positive in terms of the applicability and effectiveness of Computation Regrouping. We experiment with a range of inputs for each benchmark and observe positive speedups for most cases (more details are given elsewhere [27]). Reducing memory stall time yields almost all the measured performance improvements. In some cases computation costs are reduced, too. In cases where there is a *decrease* in the computation, especially in applications optimized using computation merging, we could not accurately compute overhead.

We present the *typical* perfex estimate values for miss ratios and miss-handling times. Perfex derives these estimates by combining processor performance counter data with execution models. In our observation, perfex performance numbers are generally optimistic, but they provide fairly good estimates of relative costs of performance parameters.

Table 5 presents perfex-estimated times for total execution and for servicing cache misses, along with the computational overhead of regrouping and the execution time savings as a percentage of the original memory stall time. The *Saved* column shows fractions of original memory stall time recovered. The *Overhead* column shows percentages of original memory stall time spent on additional computation. Some entries are blank, either because we observe a savings in the computation, as well, or because of inconsistencies in the perfex estimates. The decrease in memory stall time is 26-85%, and the computation cost incurred to achieve this reduction is 0.3-10.9% of the original stall time.

Benchmark	Access Pattern	Implementation Technique	Code		Issued Instructions		Graduated Instructions	
			Total	Change	Baseline	Restructured	Baseline	Restructured
R-Tree	Pointer Chasing	Deferred Execution with Queues	7K	600	102.5B	103.9B	89.4B	92.7B
IRREG	Indirect Accesses	Filtered Execution	300	40	15.7B	15.9B	8.2B	13.3B
Ray Trace	Strided Accesses Pointer Chasing	Filtered Execution	5K	300	108.5B	108.5B	90.8B	96.7B
FFTW	Strided Accesses	Early Execution	35K	1.5K	6.3B	5.7B	5.11B	5.07B
CUDD	Pointer	Combination Early & Deferred Execution	60K	120	13.6B	13.9B	6.9B	6.9B
EM3D	Indirect Accesses Pointer Chasing	Combination Computation Merging & Filtered Execution	700	200	646M	617M	428M	542M
Health	Pointer Chasing	Computation Merging	600	30	1.6B	0.43B	0.83B	0.39B

Table 4: Regrouping Techniques used and their Impact.

Benchmark	Input	Hit Ratio (%)				Time (sec)				% Stall Time Saved	Overhead (% Stall Time)	Speedup
		Baseline		Restructured		Baseline		Restructured				
		L1D	L2	L1D	L2	Total	Stall	Total	Stall			
R-Tree	dm23.in	95.4	49.6	97.1	77.9	1312	1194	718	474	60	10.5	1.87
IRREG	MOL2	84.5	54.5	84.9	84.3	253	238	145	104	56.3	10.9	1.74
Ray Trace	balls 256 × 256	95.8	70.1	96.1	99.6	905	531	457	81	84.7	0.3	1.98
FFTW	10K × 32 × 32	92.4	60.0	92.3	93.3	111	82	44	20	—	—	2.53
CUDD	C3540.blif	84.1	45.9	88.8	46.0	307	294	241	217	26	3.7	1.26
EM3D	128K, 10, 1	85.4	47.9	86.6	71.8	13.4	11.24	9.4	6.7	40	4.6	1.43
Health	6, 500	71.9	68.6	85.2	46.8	46	46	15	16	—	—	3.03

Table 5: Perfex-Estimated Performance Statistics

R-Tree. Section 4 explains how we extend the R-Tree data structure with queues to support deferring of tree query and delete operations. Our experimental input, *dm23.in*, consists of a million insert, delete, and query operations. The execution style of these queries is similar to batch processing, so increasing the latency of individual operations to improve throughput is an acceptable tradeoff. Overall speedup for the *dm23.in* input is 1.87. The average tree size is about fifty thousand nodes that span approximately 16MBytes. Cache hit ratios improve significantly in the optimized version of the program: from 95.4% to 97.1% for the L1D cache, and from 49.6% to 77.9% for the L2 cache. TLB performance improves somewhat, but the effect is small. The overhead incurred by using the queue structure to defer query and delete operations is 10.5% of the original memory stall time, but this cost is overwhelmed by the 60% reduction in stall time.

IRREG. IRREG is the simplest code — the main loop is fewer than 10 lines. We can easily apply filtered execution, as shown in Figure 6. A filter window size of one-fourth the array size yields a speedup of 1.74 for the *MOL2* input. Cache hit ratios for the L1D and L2 caches are 84.5% and 54.5%, respectively, for the baseline application, and 84.9% and 84.3%, respectively, for the optimized version. The performance improvement comes primarily from the increase in the L2 hit ratio. Table 5 shows that regrouping eliminates 56% of the memory stall time, and that the corresponding overhead is about 11%. The actual savings in memory stall time is higher than the estimated 56% because actual execution time is higher than the perfex estimate by about 50s, whereas the execution time of the optimized version is lower than the perfex estimate by 15s. The large data memory footprint, randomness in the input, and small amount of computation in the innermost loop are important to obtaining good performance improvement. Other standard inputs, such as *MOL1* and *AUTO*, are small enough to fit into L2 cache or are compute-bound on our experimental machine.

Ray Trace. We apply filtered execution to Ray Trace, choosing a window size equal to half the total number of objects. The straight-line code of ray processing is thus replaced by a two-iteration loop similar to the modified IRREG code shown in Figure 6. Our choice of filter window size depends on both the size of the input data structure, and to a lesser extent on the computational overhead of filtered execution. The critical working set of the baseline Ray Trace is about 4MBytes, and a filter window size of approximately half that allows most data to reside in L2 cache in the optimized version. The overheads incurred are due to the filtering control structure and the array that temporarily stores reflected rays. We defer processing reflected rays until all other rays have been processed. Our speedup is 1.98 for an input viewing window of 256 × 256 (64K rays). The L1D and L2 cache ratios improve from 95.8% and 70.1% to 96.1% and 99.6%, respectively. The optimized application effectively becomes compute-bound. 84.7% of stall time is recovered at a computational cost of 0.3% of execution time. We experiment with only one input scene, but our experience suggests that results should scale with input size.

FFTW. FFTW translates the multidimensional FFT computation into a series of calls to highly optimized leaf-level functions called *codelets*, each of which contains a predetermined number of floating-point operations and memory accesses. Each codelet’s memory accesses are in the form of short column walks on the input and output arrays. The critical working set of any codelet is no more than 64 cache lines, fitting easily into the L1D and L2 caches. The early execution optimization for FFT requires that when a column walk is executed, future column walks accessing the same data be executed at the same time. In this way, we simulate the effect of a short, synchronized multicolumn walk. The first iteration loads a fixed number of cache that are reused by the remaining column walks. This eliminates cache misses that would have occurred if the complete walks

Note that we have subsequently learned that the FFTW authors are experimenting with a similar approach.

had been executed serially. The majority of the performance improvement is obtained during the *compute* multicolumn walk, but there is also significant gain from the *copy* multicolumn walk. The regrouping overhead arises from the necessary control extensions (additional function calls and loops). There is no memory overhead. Table 5 shows that the application speedup is 1.55. Although the L1D hit ratio decreases slightly, the nearly 50% increase in the L2 hit ratio offsets the impact. Memory stall time decreases from 73.8% of the baseline application execution time to about 45% of execution time after optimization. The same optimization can be applied to the y -dimension to further improve performance.

CUDD. We experiment with the *nanotrav* tool from the CUDD package [34]. Nanotrav stores input circuits as binary decision diagrams (BDDs) and executes a series of operations to reduce the BDD sizes. The core operation is a variable swap that manipulates large hash tables. Each hash table element is accessed at least three times, and these accesses are distributed among the extraction, insertion, and garbage collection stages of a swap. The many hash table elements processed in each stage introduce a temporal separation between successive accesses to a data object, which results in little or no reuse before eviction.

We implement two regrouping optimizations in *nanotrav*. First, we execute the reference counting step of the garbage collection function early — when the nodes become garbage — rather than during a separate garbage collection phase. Second, we defer the sorting of newly inserted objects until the garbage collection phase. We observe an application speedup of 1.26 for the *C3540.blif* input. Improved L1D performance is the primary factor. Memory stall time still accounts for much of the execution time, and could be further reduced. Our experimentation suggests an alternative implementation of the BDD structures based on partitioned hash tables, but we have performed only a preliminary study of the feasibility of this solution.

EM3D The EM3D application has a graph construction phase and a computation phase. We apply computation merging to the graph construction phase, which involves incrementing counters at remote nodes. We use an array of integers to collect the increments to these counters, and access the remote nodes once at the end. We also apply filtered execution to the compute phase, but with limited impact, due to high control overhead. We use a filter window size of one-third that of the set of remote nodes. Our optimizations are effective only when the graph construction cost dominates, *i.e.*, when the number of nodes is large, and the number of compute-phase iterations is small. For 128K nodes and one compute iteration, we observe a speedup of 1.43, corresponding to a speedup of about 1.5 in the construction phase and 1.12 in the compute phase. Memory stalls still account for about 71% of the optimized application’s run time, down from 84% in the baseline. Note that a significant source of improvement is the reduced computation in the first phase.

Health. We obtain one of our best results with computation merging in Health. We defer patient waiting-list timestamp updates until the list requires modification, reducing the number of list traversals and eliminating many cache misses. The application experiences a slowdown at the beginning when the waiting lists are short and require frequent modification, inducing high deferral overhead. This is offset by the reduced number of accesses during later execution,

when there are long waiting lists and infrequent modifications. The simple structure of Health enables an efficient implementation of the regrouping optimization. Overall application speedup is 3.03. The L1D hit ratio increases from 71.9% to 85.2%, whereas the L2 hit ratio decreases from 68.6% to 46.8%. The perfex estimates are inconsistent for this application, and therefore we could not compute overheads or reductions in stall times. However, even after the optimization, most of the execution time is spent servicing cache misses. The L2 hit ratios and time estimates do not reflect the significant decrease in raw computation due to computation merging. This change in total computation made it impossible to accurately calculate regrouping overhead. To improve memory performance further, we can implement a more aggressive variation of the regrouping optimization that defers update operations from other lists and tree traversals.

6. RELATED WORK

Application memory performance has been the focus of significant research. We limit our discussion here to existing software approaches for increasing memory performance. These may be classified as cache-conscious algorithms, compiler optimizations, or application-level techniques.

Algorithmic Approaches. Cache-conscious algorithms [4, 8, 13, 23, 32]. have been developed for specific applications, such as sorting [23], query processing [32], and matrix multiplication [8]. These algorithms modify an application’s control flow, and sometimes the data structure layout, based on an understanding of the application’s interactions with memory subsystem. Such cache-conscious algorithms significantly out-perform conventional algorithms, but the existence of multiple dominant access patterns within a single program complicates algorithm design significantly. As a result, cache-conscious algorithms are few in number, and are usually domain-specific. Computation regrouping also requires substantial understanding of access patterns, but is a more generic approach. Changes required by regrouping are low-level, but fairly architecture-independent. Regrouping does not require radical application changes, and therefore we believe that regrouping can be applied in more scenarios. We demonstrate this by optimizing a variety of applications.

Compiler Approaches. Compiler-based restructuring techniques for improving spatial and temporal locality are well known [7, 17, 18, 20, 22, 24, 36]. They are usually applied to regular data structures within loops and to nearly perfect loop nests, and thus are driven by analytic models of loop costs [14]. For instance, Kodukula, Ahmed, and Pingali present a comprehensive approach to *data shackling* for regular arrays [21]. They, like we, strive to restructure computation to increase temporal data locality. They focus specifically on the problem domain of dense numerical linear algebra, and have developed rigorous methods for automatically transforming the original source code. In contrast, our approach applies to a broader class of applications and data structures, but making it automatic instead of ad hoc is part of future work.

Complex structures or access patterns are not usually considered for compiler-based restructuring due to the difficulty in deriving accurate analytic models. Further, it is not clear how these diverse techniques can be combined efficiently. Computation regrouping can be viewed as a heuristic-based, application-level variation of traditional blocking algorithms

that cannot be directly applied to complex data structures. Regrouping requires modest changes to control and data structures. Performance improvements from regrouping should be smaller than from compiler-based restructuring techniques, because of the expected higher control overhead. Profitability analysis based on cost models would be helpful to determine optimization parameters, but we do not expect them to be as detailed or accurate as in previous work.

Others have taken intuitions similar to those underlying our restructuring approach and have applied them in other arenas. For instance, to reduce accesses to non-local data, Rogers and Pingali [31] develop a multiprocessor compiler approach in which the application programmer specifies the data-structure layout among the processors, and then describes how the compiler can generate code consistent with this specification, along with some possible optimizations. Our Computation Regrouping can be viewed as a variation in which spatial decomposition appropriate for the multiprocessors is used, but with the multiple processors simulated serially in time on a uniprocessor. We extend this previous work by showing that the approach is useful for complex datastructures and access patterns. In one case, IRREG, the optimized code for both the approaches has structural similarity. Although we started at a different point in the application and machine configuration space, the approaches are similar in that as the memory cost increases, even main memory accesses start looking like remote memory accesses of distributed shared memory machines. Using computation regrouping on multiprocessors is possible, but is likely to require solving more correctness issues and to incur higher costs for integrating partial results.

Prefetching hides high data access latencies for data with poor locality [5, 19, 25], and can significantly improve performance for some kinds of programs. The effectiveness of prefetching on some complex data structures is limited because prefetching does not address the fundamental reason why the data accesses are expensive, *i.e.*, large temporal separations between successive accesses to a given object. Computation Regrouping addresses that problem by modifying the application’s control flow, and thus can complement prefetching.

Application approaches. Application designers can use data restructuring for complex structure. This can be useful when the compiler is unable to identify appropriate legal and profitable transformations. Some techniques, *e.g.*, clustering and coloring [9], take this approach, considering more complex linked data structures like trees and graphs for memory optimizations [9, 28, 29]. Spatial blocking is effective for B-Trees and variants [28, 29]. Improving spatial locality is useful, but has limited impact when the access patterns are complicated. Computation regrouping is similar to these approaches in terms of the level of user involvement, but complements them with modifications to enhance temporal locality.

Computation Regrouping. Some types of computation regrouping have been considered in other contexts. A new loop tiling technique for a class of imperfect nests, for example, tries to take advantage of accesses spread across multiple time-steps [35]. In compiler-based, reuse-driven loop fusion and multi-level data regrouping work by Ding and Kennedy [11], the authors identify opportunities to fuse array references across loop nests. In both cases, it is not clear as how the approach will scale to handle more complex

data structures. Here we extend previous work by considering applications with complex data structures and access patterns.

A queue-enhanced R-Tree called a *Buffer Tree* was previously proposed in the context of external memory algorithms [2, 3]. However, unlike earlier studies that focus on improving I/O performance, in this paper we focus on a memory-resident variation of R-Trees, and we identify the queuing extension as a specific instance of computation regrouping.

7. CONCLUSIONS

Application performance is increasingly limited by memory performance. The growing CPU-memory speed gap only exacerbates the problem. We present a software approach to attack the problem, trading extra computation for reduced memory costs. *Computation Regrouping* executes computations accessing the same data closer together in time, which significantly improves temporal locality (and thus performance) for applications with poor locality. The programmer overhead for identifying appropriate logical operations and the execution overhead of the modified application are the costs we trade for increased temporal access locality. We present a few implementation techniques to realize the regrouping approach, and demonstrate that Computation Regrouping successfully eliminates a significant fraction of memory stall time. Our hand-coded optimizations improve performance by a factor of 1.26 to 3.03 on a variety of applications with a range of sizes, access patterns, and problem domains. Our initial results are promising, and we believe further research in this direction to be warranted.

We are currently exploring a runtime library-based approach to extend work crews [30] with locality information [27]. Preliminary results are encouraging, but much work remains in identifying an abstraction suitable for compiler-based analysis.

Our conclusion from this brief investigation is that program-level access patterns, in addition to low-level access patterns, can interact in meaningful ways with the cache-based memory model of modern processors, impacting application performance. Understanding the nature of these interactions can help us to design techniques to improve application memory performance. Our logical operations-based characterization and optimizations designed to use such techniques reflect our current understanding of these program-level patterns. Going forward, the hard problems to be solved include identifying exact or statistical schemes by which non-local effects of memory accesses can be modeled efficiently, and automating the process of using these models.

ACKNOWLEDGMENTS

The authors thank the anonymous referees for their valuable comments. Keshav Pingali’s input has been invaluable, as has that of the other members of the Impulse Adaptable Memory Controller project. We also thank Robert Braden for his support of Venkata Pingali during the final stages of publication of this work.

8. REFERENCES

- [1] A. Appel, J. Ellis, and K. Li. Real-time concurrent collection on stock multiprocessors. In *Proceedings of the 1988 ACM SIGPLAN Conference on Programming*

- Language Design and Implementation*, pages 11–20, June 1988.
- [2] L. Arge. The Buffer Tree : A New Technique for Optimal I/O-Algorithms. In *Fourth Workshop on Algorithms and Data Structures*, pages 334–345, August 1995.
 - [3] L. Arge, K. Hinrichs, J. Vahrenhold, and J. S. Vitter. Efficient Bulk Operations on Dynamic R-Trees. In *1st Workshop on Algorithm Engineering and Experimentation*, pages 328–348, January 1999.
 - [4] M. A. Bender, E. D. Demaine, and M. Farach-Colton. Cache-Oblivious B-Trees. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 399–409, November 2000.
 - [5] D. Callahan, K. Kennedy, and A. Porterfield. Software prefetching. In *Proceedings of the 4th Symposium on Architectural Support for Programming Languages and Operating Systems*, pages 40–52, April 1991.
 - [6] M. Carlisle, A. Rogers, J. Reppy, and L. Hendren. Early experiences with olden. In *Proceedings of the Sixth Workshop on Languages and Compilers for Parallel Computing*, pages 1–20, August 1993.
 - [7] S. Carr, K. McKinley, and C.-W. Tseng. Compiler Optimizations for Improving Data Locality. In *Proceedings of the 6th Symposium on Architectural Support for Programming Languages and Operating Systems*, pages 252–262, October 1994.
 - [8] S. Chatterjee, A. Lebeck, P. Patnala, and M. Thottethodi. Recursive Array Layouts and Fast Matrix Multiplication. In *Proceedings of the 11th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 222–231, June 1999.
 - [9] T. M. Chilimbi, M. D. Hill, and J. R. Larus. Cache-Conscious Structure Layout. In *Proceedings of the 1999 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 1–12, May 1999.
 - [10] D. E. Culler, A. Dusseau, S. C. Goldstein, A. Krishnamurthy, S. Lumetta, T. von Eicken, and K. Yelick. Parallel Programming in Split-C. In *Proceedings of Supercomputing '93*, pages 262–273, November 1993.
 - [11] C. Ding and K. Kennedy. Improving Effective Bandwidth through Compiler Enhancement of Global Cache Reuse. In *2001 International Parallel and Distributed Processing Symposium*, April 2001.
 - [12] M. Frigo and S. Johnson. FFTW: An adaptive software architecture for the FFT. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 1381–1384, May 1998.
 - [13] M. Frigo, C. Leiserson, H. Prokop, and S. Ramachandran. Cache-Oblivious Algorithms. In *40th Annual Symposium on Foundations of Computer Science*, pages 285–297, October 1999.
 - [14] S. Ghosh, M. Martonosi, and S. Malik. Precise Miss Analysis for Program Transformations with Caches of Arbitrary Associativity. In *Architectural Support for Programming Languages and Operating Systems*, pages 228–239, October 1998.
 - [15] A. Guttmann. R-Trees : A Dynamic Index Structure for Spatial Searching. In *Proceedings of the 1984 International Conference on Management of Data*, pages 47–57, August 1984.
 - [16] H. Han and C. Tseng. Improving Compiler and Run-Time Support for Irregular Reductions. In *Proceedings of the 11th Workshop on Languages and Compilers for Parallel Computing*, Chapel Hill, NC, August 1998.
 - [17] H. Han and C.-W. Tseng. Improving locality for adaptive irregular scientific codes. Technical Report CS-TR-4039, University of Maryland, College Park, September 1999.
 - [18] M. T. Kandemir, A. N. Choudhary, J. Ramanujam, and P. Banerjee. Improving locality using loop and data transformations in an integrated framework. In *International Symposium on Microarchitecture*, pages 285–297, November-December 1998.
 - [19] M. Karlsson, F. Dahlgren, and P. Stenstrom. A Prefetching Technique for Irregular Accesses to Linked Data Structures. In *Proceedings of the Sixth Annual Symposium on High Performance Computer Architecture*, pages 206–217, January 2000.
 - [20] I. Kodukula, N. Ahmed, and K. Pingali. Data-Centric Multi-level Blocking. In *Proceedings of the SIGPLAN '97 Conference on Programming Language Design and Implementation*, pages 346–357, June 1997.
 - [21] I. Kodukula, N. Ahmed, and K. Pingali. Data-Centric Multi-level Blocking. In *Proceedings of the 1997 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 346–357, June 1997.
 - [22] M. S. Lam, E. E. Rothberg, and M. E. Wolf. The cache performance and optimizations of blocked algorithms. In *Proceedings of the 4th ASPLOS*, pages 63–74, April 1991.
 - [23] A. G. LaMarca. *Caches and Algorithms*. PhD thesis, University of Washington, 1996.
 - [24] S. Leung and J. Zahorjan. Optimizing Data Locality by Array Restructuring. Technical Report UW-CSE-95-09-01, University of Washington Dept. of Computer Science and Engineering, September 1995.
 - [25] C.-K. Luk and T. C. Mowry. Compiler-Based Prefetching for Recursive Data Structure. In *Proceedings of the 7th Symposium on Architectural Support for Programming Languages and Operating Systems*, pages 222–233, October 1996.
 - [26] J. W. Manke and J. Wu. *Data-Intensive System Benchmark Suite Analysis and Specification*. Atlantic Aerospace Electronics Corp., June 1999.
 - [27] V. Pingali. Memory performance of complex data structures: Characterization and optimization. Master's thesis, University of Utah, August 2001.
 - [28] J. Rao and K. A. Ross. Cache Conscious Indexing for Decision-Support in Main Memory. In *Proceedings of the 25th VLDB Conference*, pages 78–89, 1999.
 - [29] J. Rao and K. A. Ross. Making B+-Trees Cache Conscious in Main Memory. In *Proceedings of the 26th VLDB Conference*, pages 475–486, 2000.
 - [30] E. S. Roberts and M. T. Vandevoorde. WorkCrews: An Abstraction for Controlling Parallelism. Technical Report SRC-042, Digital Systems Research Center, April 1989.
 - [31] A. Rogers and K. Pingali. Process decomposition through locality of reference. In *Proceedings of the 1989 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 69–80, June 1989.
 - [32] A. Shafdal, C. Kant, and J. Naughton. Cache Conscious Algorithms for Relational Query Processing. In *Proceedings of the 20th VLDB Conference*, pages 510–521, September 1994.
 - [33] Silicon Graphics Inc. *SpeedShop User's Guide*. 1996.
 - [34] F. Somenzi. CUDD: CU Decision Diagram Package Release 2.3.1, 2001.
 - [35] Y. Song and Z. Li. New Tiling Techniques to Improve Cache Temporal Locality. In *Proceedings of the SIGPLAN '99 Conference on Programming Language Design and Implementation*, pages 215–228, May 1999.
 - [36] D. N. Truong, F. Bodin, and A. Sez nec. Improving Cache Behavior of Dynamically Allocated Data Structures. In *Proceedings of the 1998 International Conference on Parallel Architectures and Compilation Techniques*, pages 322–329, October 1998.
 - [37] M. Zagha, B. Larson, S. Turner, and M. Itzkowitz. Performance Analysis Using the MIPS R10000 Performance Counters. In *Proceedings of Supercomputing '96*, November 1996.