DECEMBER 1998

WRL Research Report 98/7

Comparative Evaluation of Fine- and Coarse-Grain Approaches for Software Distributed Shared Memory

Sandhya Dwarkadas
Kourosh Gharachorloo
Leonidas Kontothanassis
Daniel J. Scales
Michael L. Scott
Robert Stets

digital **Western Research Laboratory** 250 University Avenue Palo Alto, California 94301 USA

The Western Research Laboratory (WRL) is a computer systems research group that was founded by Digital Equipment Corporation in 1982. Our focus is computer science research relevant to the design and application of high performance scientific computers. We test our ideas by designing, building, and using real systems. The systems we build are research prototypes; they are not intended to become products.

There are two other research laboratories located in Palo Alto, the Network Systems Lab (NSL) and the Systems Research Center (SRC). Another Digital research group is located in Cambridge, Massachusetts (CRL).

Our research is directed towards mainstream high-performance computer systems. Our prototypes are intended to foreshadow the future computing environments used by many Digital customers. The long-term goal of WRL is to aid and accelerate the development of high-performance uni- and multi-processors. The research projects within WRL will address various aspects of high-performance computing.

We believe that significant advances in computer systems do not come from any single technological advance. Technologies, both hardware and software, do not all advance at the same pace. System design is the art of composing systems which use each level of technology in an appropriate balance. A major advance in overall system performance will require reexamination of all aspects of the system.

We do work in the design, fabrication and packaging of hardware; language processing and scaling issues in system software design; and the exploration of new applications areas that are opening up with the advent of higher performance systems. Researchers at WRL cooperate closely and move freely among the various levels of system design. This allows us to explore a wide range of tradeoffs to meet system goals.

We publish the results of our work in a variety of journals, conferences, research reports, and technical notes. This document is a research report. Research reports are normally accounts of completed research and may include material from earlier technical notes. We use technical notes for rapid distribution of technical material; usually this represents research in progress.

Research reports and technical notes may be ordered from us. You may mail your order to:

Technical Report Distribution
DEC Western Research Laboratory, WRL-2
250 University Avenue
Palo Alto, California 94301 USA

Reports and technical notes may also be ordered by electronic mail. Use one of the following addresses:

Digital E-net:	`JOVE::WRL-TECHREPORTS`
Internet:	`WRL-Techreports@decwrl.pa.dec.com`
UUCP:	`decpa!wrl-techreports`

To obtain more details on ordering by electronic mail, send a message to one of these addresses with the word ‘‘`help`’’ in the Subject line; you will receive detailed instructions.

Reports and technical notes may also be accessed via the World Wide Web: `http://www.research.digital.com/wrl/home.html`.

Sandhya Dwarkadas
Kourosh Gharachorloo
Leonidas Kontothanassis
Daniel J. Scales
Michael L. Scott
Robert Stets

December 1998

digital **Western Research Laboratory** 250 University Avenue Palo Alto, California 94301 USA

Abstract

Symmetric multiprocessors (SMPs) connected with low-latency networks provide attractive building blocks for software distributed shared memory systems. Two distinct approaches have been used: the *fine-grain* approach that instruments application loads and stores to support a small coherence granularity, and the *coarse-grain* approach based on virtual memory hardware that provides coherence at a page granularity. Fine-grain systems offer a simple migration path for applications developed on hardware multiprocessors by supporting coherence protocols similar to those implemented in hardware. On the other hand, coarse-grain systems can potentially provide higher performance through more optimized protocols and larger transfer granularities, while avoiding instrumentation overheads. Numerous studies have examined each approach individually, but major differences in experimental platforms and applications make comparison of the approaches difficult.

This paper presents a detailed comparison of two mature systems, Shasta and Cashmere, representing the fine- and coarse-grain approaches, respectively. Both systems are tuned to run on the same commercially available, state-of-the-art cluster of AlphaServer SMPs connected via a Memory Channel network. As expected, our results show that Shasta provides robust performance for applications tuned for hardware multiprocessors, and can better tolerate fine-grain synchronization. In contrast, Cashmere is highly sensitive to fine-grain synchronization, but provides a performance edge for applications with coarse-grain behavior. Interestingly, we found that the performance gap between the systems can often be bridged by program modifications that address coherence and synchronization granularity. In addition, our study reveals some unexpected results related to the interaction of current compiler technology with application instrumentation, and the ability of SMP-aware protocols to avoid certain performance disadvantages of coarse-grain approaches.

Comparative Evaluation of Fine- and Coarse-Grain Approaches for Software Distributed Shared Memory

Sandhya Dwarkadas[*], Kourosh Gharachorloo[†], Leonidas Kontothanassis[‡], Daniel J. Scales[†], Michael L. Scott[*], and Robert Stets[*]

[*]Dept. of Comp. Science
University of Rochester
Rochester, NY 14627

[†]Western Research Lab
Compaq Computer Corp.
Palo Alto, CA 94301

[‡]Cambridge Research Lab
Compaq Computer Corp.
Cambridge, MA 02139

Abstract

Symmetric multiprocessors (SMPs) connected with low-latency networks provide attractive building blocks for software distributed shared memory systems. Two distinct approaches have been used: the *fine-grain* approach that instruments application loads and stores to support a small coherence granularity, and the *coarse-grain* approach based on virtual memory hardware that provides coherence at a page granularity. Fine-grain systems offer a simple migration path for applications developed on hardware multiprocessors by supporting coherence protocols similar to those implemented in hardware. On the other hand, coarse-grain systems can potentially provide higher performance through more optimized protocols and larger transfer granularities, while avoiding instrumentation overheads. Numerous studies have examined each approach individually, but major differences in experimental platforms and applications make comparison of the approaches difficult.

This paper presents a detailed comparison of two mature systems, Shasta and Cashmere, representing the fine- and coarse-grain approaches, respectively. Both systems are tuned to run on the same commercially available, state-of-the-art cluster of AlphaServer SMPs connected via a Memory Channel network. As expected, our results show that Shasta provides robust performance for applications tuned for hardware multiprocessors, and can better tolerate fine-grain synchronization. In contrast, Cashmere is highly sensitive to fine-grain synchronization, but provides a performance edge for applications with coarse-grain behavior. Interestingly, we found that the performance gap between the systems can often be bridged by program modifications that address coherence and synchronization granularity. In addition, our study reveals some unexpected results related to the interaction of current compiler technology with application instrumentation, and the ability of SMP-aware protocols to avoid certain performance disadvantages of coarse-grain approaches.

1 Introduction

Clusters of symmetric multiprocessors (SMP) provide a powerful platform for executing parallel applications. To ease the burden of programming such clusters, software distributed shared memory (S-DSM) systems support the illusion of shared memory across the cluster via a software run-time layer between the application and the hardware. This approach can potentially provide a cost-effective alternative to larger hardware shared memory systems for executing certain classes of workloads.

Most S-DSM systems use virtual memory hardware to detect access to data that is not available locally. Hence, data is communicated and kept coherent at the coarse granularity of a page (e.g., 4-16KB). Early page-based systems [14] suffered from *false sharing* that arises from fine-grain sharing of data within a page. More recent page-based systems [2, 3, 10, 13] address this issue by employing relaxed memory consistency models that enable protocol optimizations such as delaying coherence operations to synchronization points and allowing

multiple processors to concurrently write to a page. Page-based systems may still experience overheads due to frequent synchronization or sharing at a fine granularity. Furthermore, the aggressive relaxed memory models and the required use of predefined synchronization primitives limit portability for certain applications developed on hardware multiprocessors [17].

As an alternative, a few S-DSM systems [16, 19] have explored supporting data sharing and coherence at a finer granularity (e.g., 64-256 bytes). Fine-grain access is supported by instrumenting the application binary at loads and stores to check if the shared data is available locally. Such systems provide the highest degree of transparency since they can correctly run all programs (or even binaries) developed for hardware multiprocessors by virtue of supporting similar memory models [17]. In addition, this approach reduces false sharing and the transmission of unnecessary data, both of which are potential problems in page-based systems. Nevertheless, page-based systems can potentially benefit from more optimized protocols and larger transfer granularities, without incurring the software checking overheads associated with fine-grain systems.

The recent prevalence of low-cost SMP nodes has led to extensions to software DSM designs for supporting shared memory across SMP clusters. The key advantage of using SMP nodes comes from supporting data sharing within a node directly via the cache coherence hardware, and only invoking the software protocol for sharing across nodes. Several studies have demonstrated significant gains from exploiting SMP-aware protocols in both coarse-grain [7, 15, 21] and fine-grain [18] S-DSM systems.

As described above, there are important trade-offs between *coarse-grain, page-based* and *fine-grain, instrumentation-based* S-DSM systems, both in the *performance* and the *generality* of the shared-memory programming model. Even though there are a large number of papers that study each approach individually, a direct comparison is difficult due to major differences in the experimental platforms, applications, and problem sizes used in the various studies. Furthermore, only a few studies are actually based on SMP-aware protocols.

This paper presents a detailed comparison of the fine- and coarse-grain approaches based on the same hardware platform and applications. We use two mature systems, Shasta [18] and Cashmere [21], both of which are highly efficient and tuned to run on a state-of-the-art cluster of Digital AlphaServer multiprocessors connected through the Memory Channel network [9]. We study a total of thirteen applications: eight SPLASH-2 applications [22] that have been developed for hardware multiprocessors and five applications that were developed for page-based S-DSM systems. The first part of our study compares the performance of the *unmodified* applications on the two systems. This part allows us to evaluate the portability of applications developed for hardware multiprocessors and to measure the performance gap on applications developed for page-based systems. The second part of the study analyzes the performance of the same applications after modifications that improve their performance on either system. To ensure a fair comparison, we undertook this study as a collaborative effort between the Cashmere and Shasta groups.

Our study quantifies a number of expected trends. Shasta provides robust and often better performance for applications written for hardware multiprocessors and is better able to tolerate fine-grain behavior. Cashmere is highly sensitive to the presence of fine-grain synchronization, but provides a performance edge for applications with coarse-grain behavior. However, we found that the performance gap between the systems can be bridged by program modifications that take coherence granularity into account.

Our study also presents several unexpected results. One interesting result is that Cashmere, due to its SMP-aware implementation, shows very good performance on certain applications known to have a high degree of write-write false sharing at the page level. The regular data layout in these applications leads to page-aligned data boundaries across nodes, thus confining the write-write false sharing to processes on the same SMP node and avoiding the expected software overheads. Another interesting result is that fine-grain false sharing can sometimes favor Cashmere relative to Shasta, due to Cashmere's ability to delay and aggregate coherence operations. Finally, the instrumentation overheads in Shasta were more of a determining factor than we expected in a few cases. This effect is partly due to continued improvements in the Alpha compiler that lead to more efficient

code, thus increasing the relative overhead of instrumentation code in some cases.

The only relevant study that we are aware of is by Zhou et al. [23], which also examines performance tradeoffs between fine- and coarse-grain software coherence. However, several critical differences between the studies lead to differing performance results and a number of novel observations in our work. Section 5 contains a detailed comparison of the two studies.

The remainder of this paper is organized as follows. Section 2 presents an overview of the two systems that we compare in this paper. The experimental environment is described in Section 3. Section 4 presents and analyzes the results from our comparison. Finally, we present related work and conclude.

2 Overview of Cashmere and Shasta

This section presents a brief overview of Shasta and Cashmere, and also discusses some portability issues for the two systems. More detailed descriptions of the systems can be found in previous papers [13, 16, 17, 18, 21].

2.1 Shasta

Shasta is a fine-grain software DSM system that relies on inline checks to detect misses to shared data and service them in software. Shasta divides the shared address space into ranges of memory called *blocks*. All data within a block is always fetched and kept coherent as a unit. Shasta inserts code in the application executable at loads and stores to determine if the referenced block is in the correct state and to invoke protocol code if necessary. A unique aspect of the Shasta system is that the block size (i.e. coherence granularity) can be different for different application data structures. To simplify the inline code, Shasta divides the blocks into fixed-size ranges called *lines* (typically 64-256 bytes) and maintains state information for each line. Each inline check requires about seven instructions. Shasta uses a number of optimizations to eliminate checks, reduce the cost of checking loads, and to batch together checks for neighboring loads and stores [16]. Batching can reduce overhead significantly (from a level of 60-70% to 20-30% overhead for dense matrix codes) by avoiding repeated checks to the same line.

Coherence is maintained using a directory-based invalidation protocol. A *home* processor is associated with each block and maintains a *directory* for that block, which contains a list of the processors caching a copy of the block. The Shasta protocol exploits the release consistency model [8] by implementing non-blocking stores and allowing reads and writes to blocks in pending states.

When used in an SMP cluster, Shasta exploits the underlying hardware to maintain coherence within each node [18]. The SMP-aware protocol avoids race conditions by obtaining locks on individual blocks during protocol operations. However, such synchronization is not used in the inline checking code, since it would greatly increase the instrumentation overhead. Instead, the protocol selectively sends explicit messages between processors on the same node for a few protocol operations that can lead to race conditions involving the inline checks. Because Shasta supports programs with races on shared memory locations, the protocol must correctly handle various corner cases that do not arise in protocols (such as Cashmere's) that only support race-free programs [17].

Messages from other processors are serviced through a polling mechanism in both Shasta and Cashmere because of the high cost of handling messages via interrupts. Both protocols poll for messages whenever waiting for a reply and on every loop back-edge. Polling is inexpensive (three instructions) on our Memory Channel cluster because the implementation arranges for a single cachable location that can be tested to determine if a message has arrived.

2.2 Cashmere

Cashmere is a page-based software DSM system that has been designed for SMP clusters connected via a remote-memory-write network such as the Memory Channel [21]. It implements a multiple-writer, release consistent protocol and requires applications to adhere to the data-race-free or properly-labeled programming model [1]. Cashmere requires shared memory accesses to be protected by high-level synchronization primitives such as locks, barriers, or flags that are supported by the run-time system. The consistency model implementation

lies in between those of TreadMarks [2] and Munin [3]. Invalidations in Cashmere are sent during a release and take effect at the time of the next acquire, regardless of whether they are causally related to the acquired lock.

Cashmere uses the broadcast capabilities of the Memory Channel network to maintain a replicated directory of sharing information for each page (i.e., each node maintains a complete copy of the directory). Initially, shared pages are mapped only on their associated home nodes. A page fault generates a request for an up-to-date copy of the page from the home node. A page fault triggered by a write access results in either the current writer becoming the new home node (*home node migration*) or in the creation of a pristine copy of the page (a *twin*). Home node migration occurs if the current home node is not actively modifying the page. Otherwise, a twin is created. Twins then are only created when multiple nodes are concurrently modifying a page, or in other words when a page is falsely shared between nodes. As the final step in servicing a write fault, the page is added to a per-processor *dirty list* (a list of all pages modified by a processor since the last release). As an optimization, Cashmere moves the page into *exclusive* mode if there are no other sharers, and avoids adding the page to the dirty list.

At a release, each page in the dirty list is compared to its *twin*, and the differences are flushed to the home node. After flushing the differences, the releaser sends write notifications to the sharers of each dirty page, as indicated by the page's directory entry. Finally the releaser downgrades write permissions for the dirty pages and clears the list. At a subsequent acquire, a processor invalidates all pages for which notifications have been received, and which have not already been updated by another processor on the node.

The protocol exploits hardware coherence to maintain consistency within each SMP node. All processors in the node share the same physical frame for a shared data page and hence see all local modifications to the page immediately. The protocol is also designed to avoid synchronization within a node whenever possible. For instance the protocol avoids the need for TLB shootdown on incoming page updates by comparing the incoming page to the twin if one exists, thereby detecting and applying only the modifications made by remote nodes. This allows concurrent modifications to the page by other processes within the node without the need for synchronization. The correctness of this approach depends on the assumption that programs are race-free.

2.3 Portability Issues

This section briefly discusses differences between Shasta and Cashmere with respect to two important portability issues: (a) the portability of applications to each software system, and (b) the portability of the underlying software system to different hardware platforms.

Application Portability. One of the key goals in the Shasta design is to support transparent execution of applications (or binaries) developed for hardware multiprocessors [17]. Shasta achieves this transparency by supporting memory consistency models that are similar to hardware systems. On the other hand, Cashmere (like virtually all other page-based systems) opts for a departure from the standard hardware shared-memory programming model in order to achieve better performance. By requiring the use of predefined high-level synchronization primitives to eliminate shared-memory races and using aggressive relaxed memory models, Cashmere can exploit numerous protocol optimizations that are especially important for page-based systems. However, this approach may require extra programming effort to achieve a correct and efficient port of applications that depend on the more general hardware shared-memory programming model.

System Portability. The Cashmere protocol makes heavy use of Memory Channel features, including broadcasting and guarantees on global message ordering. For example, broadcasting is used to propagate directory changes to all nodes. In addition, during a release operation, the processor sending write notifications does not wait for acknowledgements before releasing the lock. Rather, it relies on global ordering of messages to guarantee that causally related invalidations are seen by other processors before any later acquire operation. It is difficult to estimate the performance impact if the protocol were changed to eliminate reliance on broadcasting and total ordering, since the protocol design assumed these network capabilities. In contrast, Shasta was designed for a network that simply offers fast user-level message passing and is therefore more portable to different network

architectures.

On the other hand, Shasta is tuned for the Alpha processor and requires detailed knowledge of both the compiler and the underlying processor architecture for efficient instrumentation. It is again hard to estimate the performance impact of moving the system to a significantly different processor architecture (e.g., Intel x86) where potentially a large variety of instructions can access memory.

3 Experimental Methodology

This section describes our prototype SMP cluster and the applications used in our study.

3.1 Prototype SMP Cluster

Our SMP cluster consists of four DEC Alpha Server 4100 multiprocessors connected by a Memory Channel network. Each AlphaServer 4100 has four 400MHz 21164 processors with 512MBytes of shared local memory. Each processor has 8K on-chip instruction and data caches, a 96K on-chip second-level cache (3-way set associative), and a 4MByte board-level cache (direct-mapped with 64-byte lines). The individual processors are rated at 12.1 SpecInt95 and 17.2 SpecFP95, and the system bus has a bandwidth of 1 Gbyte/s.

The Memory Channel is a memory-mapped network that allows a process to transmit data to a remote process without any operating system overhead via a simple store to a mapped page [9]. The one-way latency from user process to user process over Memory Channel is about 3.5 microseconds, and each network link can support a bandwidth of 70 MBytes/sec (with an aggregate bandwidth of 100MBytes/sec). For Shasta, the roundtrip latency to fetch a 64-byte block from a remote node (two hops) via the Memory Channel is 18 microseconds, and the effective bandwidth for large blocks is about 35 MBytes/s. For Cashmere, the roundtrip latency for fetching an 8K page is less than 600 microseconds.

3.2 Applications

We present results for thirteen applications. The first eight are taken from the Splash-2 [22] suite and have been tuned for hardware shared memory multiprocessors. Five Splash-2 applications are not used: four (Cholesky, FFT, Radiosity, Radix) do not perform well on S-DSM and one (FMM) has not been modified to run under Cashmere (but gets good speedup on Shasta).[1]

Barnes-Hut: an N-body simulation using the hierarchical Barnes-Hut method. The computation has two distinct phases. The first phase builds a shared octree data structure based on the relative positions of the bodies. The second phase computes forces between bodies and updates the relative locations of bodies based on the force calculation. This application requires an extra flag synchronization in the parallel tree-building phase to work correctly (i.e., to make it race-free) under Cashmere (discussed in next section).

Lu: a kernel that finds a factorization for a given matrix. The matrix is divided into square blocks that are distributed among processors in a round-robin fashion.

Contiguous Lu: another kernel that is computationally identical to Lu, but allocates each block contiguously in memory.

Ocean: an application that studies large-scale ocean movements based on eddy and boundary currents. The application partitions the ocean grid into square-like subgrids (tiles) to improve the communication to computation ratio (on a hardware DSM).

Raytrace: a program that renders a three-dimensional scene using ray tracing. The image plane is partitioned among processors in contiguous blocks of pixels, and load balancing is achieved by using distributed task queues with task stealing.

[1] FMM is not "race-free" and requires additional synchronization to work correctly under Cashmere.

Volrend: an application that renders a three-dimensional volume using a ray casting technique. The partitioning of work and the load balancing method are similar to those of Raytrace.

Water-nsquared: a fluid flow simulation. The shared array of molecule structures is divided into equal contiguous chunks, with each chunk assigned to a different processor. The bulk of the interprocessor communication occurs during a phase that updates intermolecular forces using per-molecule locks, resulting in a migratory sharing pattern.

Water-spatial: a fluid flow simulation that solves the same problem as Water-nsquared. It imposes a uniform 3-D grid of cells on the problem domain and uses a linear algorithm that is more efficient for a large number of molecules. Processors that own a cell need only look at neighboring cells to find molecules that might be within the cutoff radius of molecules in the owned box. The movement of molecules in and out of cells causes cell lists to be updated, resulting in additional communication.

The remaining five applications we use are programs that have been tuned and studied in the context of software DSM systems in the past and have been shown to have good performance on such systems [13, 21].

Em3d: a program that simulates electromagnetic wave propagation through 3D objects [4]. The major data structure is an array that contains the set of magnetic and electric nodes. Nodes are distributed among processors in a blocked fashion.

Ilink: a genetic linkage analysis program from the FASTLINK 2.3P package that locates disease genes on chromosomes [5]. The main shared data is a pool of sparse arrays of genotype probabilities. Load balancing is achieved by assigning non-zero elements to processors in a round-robin fashion. The computation is master-slave, with one-to-all and all-to-one data communication. Scalability is limited by an inherent serial component and inherent load imbalance.

Gauss: a kernel that solves a system of linear equations $AX = B$ using Gaussian Elimination and back-substitution. For load balance, the rows are distributed among processors cyclically.

SOR: a kernel that uses Red-Black Successive Over-Relaxation for solving partial differential equations. The red and black arrays are divided into roughly equal size bands of rows, with each band assigned to a different processor. Communication occurs across the boundaries between bands.

TSP: a branch-and-bound solution to the traveling salesman problem that uses a work-queue based parallelization strategy. The algorithm is non-deterministic in the sense that the earlier some processor stumbles upon the shortest path, the more quickly other parts of the search space can be pruned. The bound variable is read without any synchronization (i.e., it is not race-free), but the algorithm still works correctly on Cashmere.

4 Results

This section provides an in-depth comparison and analysis of the performance of Shasta and Cashmere. We begin by presenting our overall results for the unmodified applications, followed by a detailed analysis. We next consider minor application and system configuration modifications that improve performance on either Shasta or Cashmere.

4.1 Base Results for Unmodified Applications

Table 1 presents the sequential execution time (without any instrumentation or runtime library overhead), data set size, and memory usage for each application. Wherever possible, we use two dataset sizes — one relatively small, the other larger. This allows us to study performance sensitivity to the size and alignment of data structures

Program	Smaller Problem (Data Set) Size	Time (sec.)	Larger Problem (Data set) Size	Time (sec.)
Barnes-Hut	32K bodies (39M)	15.30	131K bodies (153M)	74.69
LU	1024x1024, block: 16 (8M)	14.21	2048x2048, block: 32 (33M)	74.57
CLU	1024x1024, block: 16 (8M)	6.74	2048x2048, block: 32 (33M)	44.40
Ocean	514x514 (64M)	7.33	1026x1026 (242M)	37.05
Raytrace	balls4 (102M)	44.89	—	—
Volrend	head (23M)	3.81	—	—
Water-nsq	4K mols., 2 steps (3M)	94.20	8K mols., 2 steps (5M)	362.74
Water-sp	4K mols., 2 steps (3M)	10.94	8K mols, 2 steps (5M)	21.12
Em3d	64000 nodes (52M)	47.61	192000 nodes (157M)	158.43
Gauss	1700x1700 (23M)	99.94	2048x2048 (33M)	245.06
Ilink	CLP (15M)	238.05	—	—
Sor	3070x2047 (50M)	21.13	3070x3070 (100M)	28.80
TSP	17 cities (1M)	1580.10	—	—

Table 1: Problem and data set sizes and sequential execution time of applications.

relative to the coherence block size. For Shasta, the instrumentation overhead (not shown) increases the single-processor execution time from 9% to 55%, with an average overhead of 27% across all applications and data set sizes. The relative importance of this checking overhead decreases in parallel executions due to the typical increase in communication and synchronization overheads.

Figures 1 and 2 present the speedups on the two systems using their base configurations for the thirteen applications on 8 and 16 processors (two and four SMP nodes each with four processors), for the smaller and larger data sets, respectively. The base configurations used are a uniform block size of 256 bytes for Shasta, and a block size of 8192 bytes (the underlying page size) for Cashmere. The applications were run without any modifications (except for Barnes, as explained in the next section), and are built with the native C compiler (-O2 optimization level). Application processes are pinned to processors at startup. Execution times are based on the best of three runs, and speedups are calculated with respect to the times of the sequential application without instrumentation or protocol overhead.

Overall, the results in Figures 1 and 2 show that Shasta provides more robust and better performance on average for the eight Splash-2 applications (which have been developed for hardware multiprocessors). Nonetheless, Cashmere provides comparable or better performance on some of the Splash-2 applications, and performs much better than expected on a few applications that are known to exhibit a high degree of fine-grain sharing (e.g., LU, OCEAN). In addition, Cashmere provides superior performance for the other five applications, which have been developed for page-based systems. The next section provides a detailed analysis of these results.

4.2 Detailed Analysis of Base Results

This section presents a detailed comparison and analysis of the performance of the various applications on Shasta and Cashmere. To better understand the reasons for performance differences, we discuss the applications in groups based on their spatial *data access granularity* and temporal *synchronization granularity* (similar to notions used by Zhou et al. [23]). Applications with coarse-grain data access tend to work on contiguous regions at a time, while fine-grain applications are likely to do scattered reads or writes. The temporal synchronization granularity is related to the frequency of synchronization in an application on a given platform. An application has fine-grain synchronization if the average computation time between consecutive synchronization events is

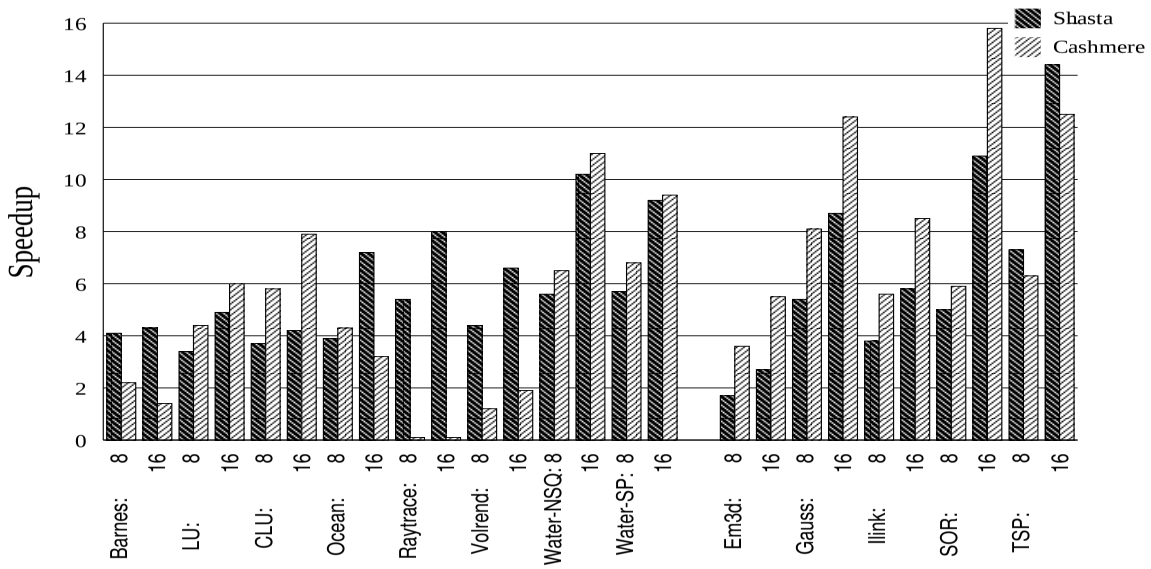

Figure 1: Speedups for the smaller data set at 8 and 16 processors.

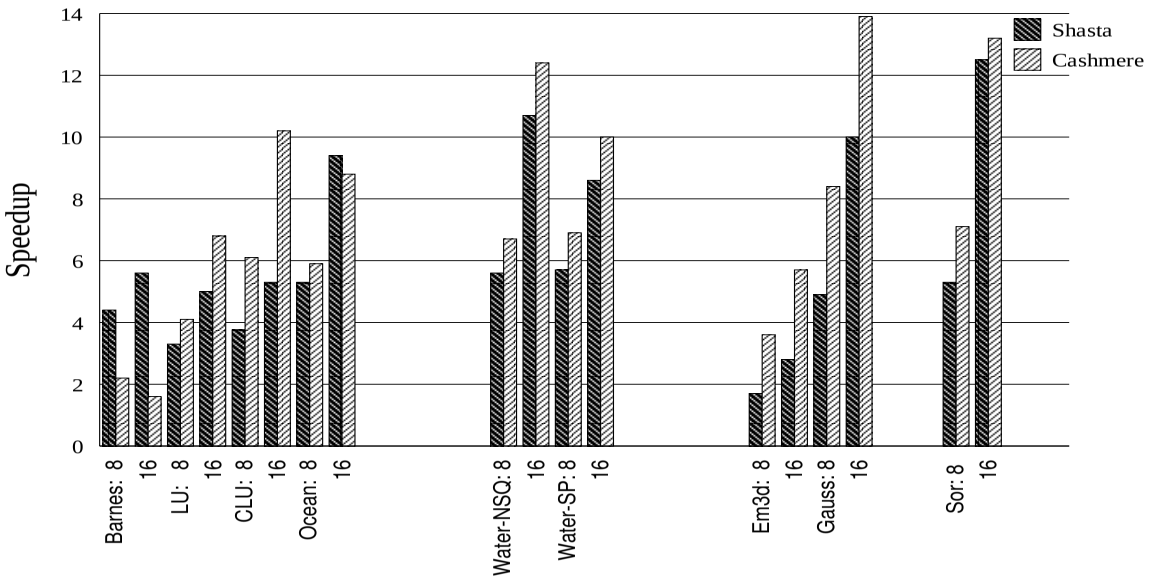

Figure 2: Speedups for the larger data set at 8 and 16 processors (Raytrace, Volrend, TSP, and Ilink not included).

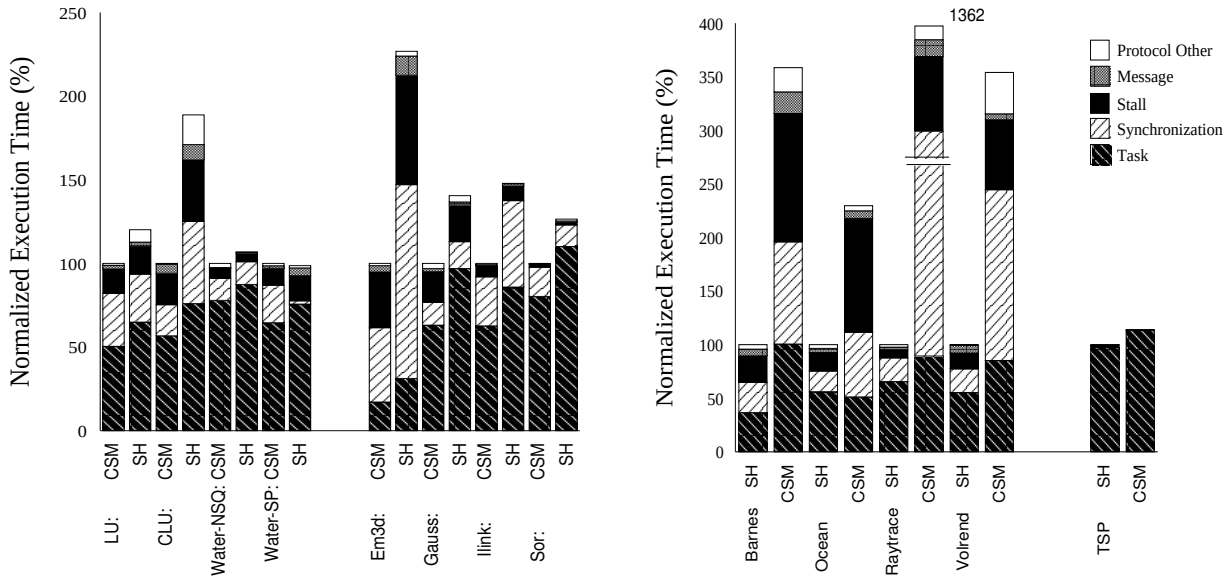

Figure 3: Application execution time breakdown on the smaller data set for 16-processor runs.

not much larger than the cost of the synchronization events themselves.[2]

Throughout this section, we will be referring to Figures 3 and 4. These figures provide a breakdown of the execution time for Shasta and Cashmere (labeled as SH and CSM, respectively) for each of the applications at 16 processors on the base and the larger data set, respectively. Execution time is normalized to that of the fastest system on each application and is split into Task, Synchronization, Data Stall, Messaging, and Protocol time. *Task* time includes the application's compute time, the cost of polling, and the cost of instrumentation in Shasta or page faults in Cashmere. *Synchronization* time is the time spent waiting on locks, flags, or barriers. *Data Stall* time measures the cumulative time spent handling coherence misses. *Messaging* time covers the time spent handling messages when the processor is not already stalled. Finally, *Protocol* time represents the remaining overhead introduced by the protocol.

Table 2 provides more detailed execution statistics on the two systems. The number of lock and flag acquires, the number of barrier operations, the number of messages (including data, synchronization, and protocol) and the amount of the message traffic (including application and protocol) are reported for both protocols. For Shasta, the number of read and write misses and the number of upgrade operations (when block is upgraded from shared to exclusive state) are also included. For Cashmere, the number of read and write page faults and the number of twins are reported.

4.2.1 Coarse-Grain Access and Synchronization

The applications in this group are CLU, Em3d, Gauss, SOR, TSP, and Water-nsquared. Overall, Cashmere is expected to perform better on these applications given the coarse communication and synchronization granularities.

CLU uses a tiled data partitioning strategy with each tile of the matrix allocated as a contiguous chunk. Cashmere performs 1.7 times better on average than Shasta for CLU. Data is propagated more efficiently under Cashmere given its large communication granularity. As shown in Figures 3 and 4, the Cashmere data stall time component is roughly half that of Shasta. The figure also shows a higher "Task" time under Shasta because of

[2]Some applications that are classified as exhibiting coarse-grain synchronization in Zhou et al.'s study [23] exhibit fine-grain behavior in our study due to the faster processors of our platform.

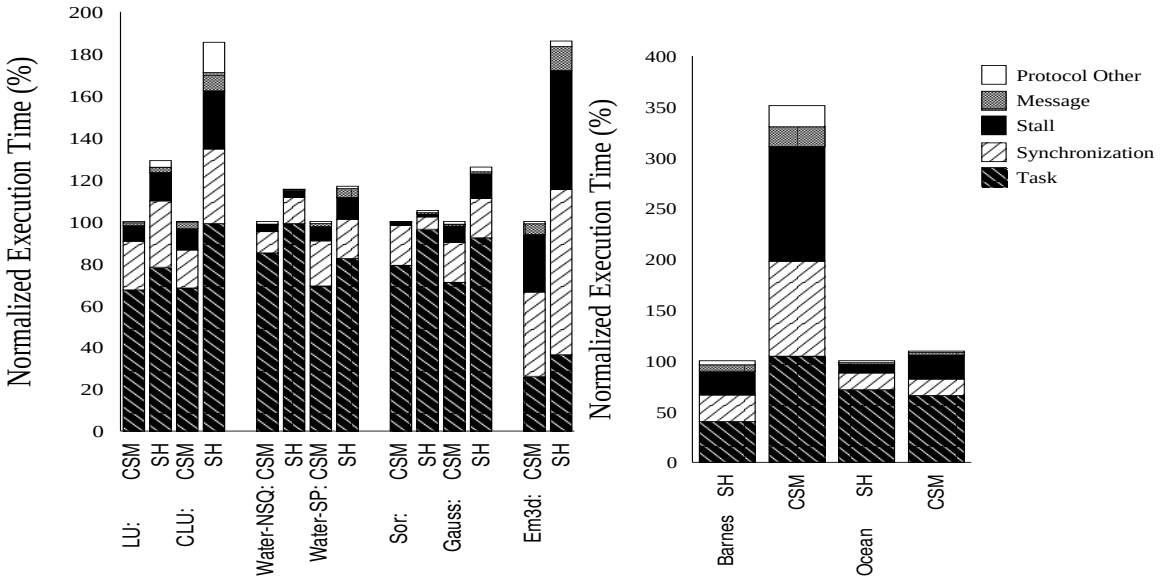

Figure 4: Application execution time breakdown on the larger data set (Raytrace, Volrend, TSP, and Ilink not included) at 16 processors.

Shasta's checking overhead (as high as 50% for a uniprocessor execution). Section 4.3 presents further results under Shasta with variable granularity and different compiler flags to address the communication granularity and checking overhead issues.

Em3d exhibits nearest-neighbor sharing, though the communication is determined at run-time based on indirection arrays. Cashmere performs two times better than Shasta on Em3d. Because of the coarse granularity of communication in this application, Shasta requires ten times more messages to fetch all the data (see Table 2). As we will see in Section 4.3, Shasta's variable granularity feature can be used to close this performance gap.

Gauss uses a cyclic distribution of matrix rows among processors. Cashmere performs 1.7 and 1.4 times better than Shasta for the larger dataset, and 1.5 and 1.4 times better for the smaller dataset, at 8 and 16 processors respectively. Because the matrix is triangularized, fewer elements are modified in each succeeding row, and Cashmere's large granularity causes communication of unnecessary data. For the 1700x1700 dataset, there is also write-write false sharing, since a row is not a multiple of the page size. However, the effects of false sharing are not as large as one might expect due to the SMP-aware protocol. In effect, the distribution of work becomes block-cyclic, with false sharing only on the edges of each block. For the 2048x2048 dataset, the effect of false sharing is eliminated since each row is a multiple of a page size. As can be seen in Table 2, Cashmere sends almost twice as much data as Shasta. However, Shasta's checking overheads ("Task" times in Figures 3 and 4) and larger message count (more than double) lead to the lower relative performance of Shasta.

In SOR, each processor operates on a block of contiguous rows and infrequently communicates with its nearest neighbors. The determining factor is the checking overhead in Shasta (as shown by the higher "Task" times in Figures 3 and 4). Relative to Shasta, Cashmere performs 1.25 better on average.

TSP has a very coarse work granularity, so communication overheads in either system are largely unimportant (see execution breakdowns in Figures 3 and 4). However, the more eager protocol in Shasta can lead to faster propagation of the bound value. This can in turn lead to a more efficient search (given the non-deterministic nature of the branch-and-bound algorithm). Shasta performs approximately 1.15 times better than Cashmere on TSP.

Water-nsquared partitions work such that processors modify contiguous regions of memory. Any false sharing

is only at the boundaries of these regions. In addition, there is considerable node locality in the data access; at least half the lock acquires access data that was last modified within the same SMP node and is therefore fetched via the node's hardware protocol. Hence, the overheads of false sharing are small in Cashmere. Shasta's performance is primarily affected by the extra checking overhead, and Cashmere performs 1.2 times better than Shasta on average.

Overall, the coarse-grain communication and low frequency of synchronization favors Cashmere in these applications. Furthermore, for Gauss and Water-nsquared, the effect of false sharing on the performance of Cashmere is dramatically reduced by the use of SMP-aware protocols, since the false sharing largely occurs among processors on the same node. The relative performance of Shasta is often determined by the checking overhead and in some cases by its smaller data transfer granularity.

4.2.2 Fine-Grain Access with Coarse-Grain Synchronization

The applications in this group are Ilink, LU, Ocean, and Water-spatial. As we will discuss below, the use of SMP nodes with SMP-aware protocols leads to some surprising results for Cashmere.

Ilink computes on sparse arrays of probabilities and uses round-robin work allocation. The sparse data structure causes Cashmere to communicate extra data on pages that have been modified (since whole pages are communicated on a miss). Shasta's performance on Ilink is affected by three factors: the checking overhead, the small communication granularity, and the use of an eager protocol. The instrumentation overhead (as high as 60% on a uniprocessor) is due to the compiler being unable to verify the commonality of certain high-frequency double indirection operations, and Shasta therefore being unable to batch them effectively. Because the work allocation is round-robin on a per-element basis, there is also much false sharing despite the small block size used by Shasta. Shasta eagerly invalidates all copies of a block whenever any processor writes to the block and generates 10 times more protocol messages than Cashmere, which delays invalidations until synchronization points. These effects outweigh the overheads for Cashmere, and Cashmere performs 1.5 times better than Shasta.

LU uses a tiled partitioning strategy. However, unlike CLU, the matrix is allocated as a single object. Hence, each tile consists of small non-contiguous regions of memory on multiple pages. The small read granularity causes a large amount of extra data to be communicated under Cashmere. In addition, the data layout leads to a large amount of false sharing at the page level. However, LU's 2D scatter distribution leads to an assignment of tiles to processors that confines all false sharing to within each 4-processor SMP node, so all false sharing is handled in hardware. (As Table 2 shows, no twin operations are performed despite the known false sharing.) The above effect, along with the checking overheads in Shasta, allows Cashmere to perform better than Shasta by a factor of 1.2 and 1.3 times at 8 and 16 processors respectively.

Ocean also uses tiled data partitioning. For the two datasets, Cashmere performs 1.10 and 1.11 times better than Shasta at 8 processors, while Shasta is 2.25 and 1.07 times better at 16 processors. The communication is nearest-neighbor in both the column and row direction. Hence, while the tiled partitioning reduces true sharing, it increases the amount of unnecessary data communicated when a large coherence unit is used. For example, Cashmere incurs 7 to 8 times more data traffic compared to Shasta at the smaller dataset size (see Table 2). The extra communication generated due to false sharing in Cashmere (incurred on every boundary) increases with the number of processors, which explains Cashmere's lower relative performance at 16 processors. As in LU, the effect of false sharing in Cashmere is greatly reduced because a large portion of the false sharing is confined to individual nodes. Furthermore, both Shasta and Cashmere benefit significantly from the large portion of true sharing communication that is confined to each SMP node [18, 21].

For Water-spatial, Cashmere performs 1.2 times better than Shasta on average, with Shasta's performance being comparable to Cashmere's at the smaller dataset size and 16 processors. In this version of the fluid-flow simulation program (compared to Water-nsquared), a uniform 3-D grid of cells is imposed on the problem domain. Processors own certain cells and only access those cells and their neighbors. Molecules can move

between cells during the simulation, creating a loss of locality, but this effect is small in both Cashmere and Shasta. As with Water-nsquared, the performance difference between the two systems can be attributed to the checking overhead in Shasta (as can be seen by the difference in "Task" time in Figures 3 and 4).

Overall, the performance of Shasta and Cashmere is comparable for the above set of applications. This is surprising for programs such as LU and Ocean, which exhibit frequent false sharing at the page-level. However, much or all of this false sharing turns out to occur between processors on the same SMP node (due to task allocation policy or nearest neighbor communication behavior) and is handled efficiently by the SMP-aware protocol. In addition, the low frequency of synchronization, along with the lazy protocol employed by the page-based system, allows Cashmere to tolerate any false sharing between nodes. At the same time, Shasta's eager protocol causes extra communication in applications (such as Ilink) that exhibit false sharing even at small block sizes.

4.2.3 Fine-Grain Access and Synchronization

The applications in this group are Barnes-Hut, Raytrace, and Volrend. As we will see, the combination of fine-grain data access and synchronization leads to excess communication and false sharing in page-based systems.

The main data structure in Barnes is a tree of nodes, each with a size of 96 bytes, so there is significant false sharing in Cashmere runs. (Note the large number of twins reported in Table 2.) Furthermore, this application relies on processor consistency in the parallel tree-building phase. Hence, while this application can run correctly on Shasta (which can enforce this form of consistency), it must be modified for Cashmere by inserting an extra flag synchronization in the parallel tree-building phase. The performance presented is for the unmodified Barnes program under Shasta, and with the additional flag synchronization under Cashmere. Shasta performs 2 and 3.5 times better than Cashmere at 8 and 16 processors, respectively. The main reason for this difference is the parallel tree building phase. This phase constitutes 2% of the sequential execution time, but slows down by a factor of 24 under Cashmere because of the fine-grain access, excessive false sharing, and extra synchronization. Shasta also suffers a slowdown in this phase, but only by a factor of 2.

Raytrace shows an even more dramatic difference between the performance of Shasta and Cashmere. Shasta performs 7 and 12.5 times better than Cashmere (which actually has a large slowdown) at 8 and 16 processors, respectively. This result is surprising, since there is little communication in the main computational loop that accesses the image plane and ray data structures. The performance difference can be primarily attributed to a single critical section used to increment a global counter in order to identify each ray uniquely. It turns out the ray identifiers are used only for debugging and could easily be eliminated (see Section 4.3.2). Their presence, however, illustrates the sensitivity of Cashmere to synchronization and data access granularity. Although only a single word is modified within the critical section, an entire page must be moved back and forth among the processors. Shasta's performance is insensitive to the synchronization, and is more in line with the behavior of a hardware DSM platform.

Volrend partitions its image plane into small tiles that constitute a unit of work, and relies on task stealing via a central queue to provide load balance. Shasta performs 3.5 times better than Cashmere for this application. Figure 3 shows that data wait and synchronization time account for 60% of the Cashmere execution time, but only about 35% of Shasta's execution. This difference results from the high degree of page-level false sharing present in the application's task queue and image data. As a result of the false sharing, Cashmere communicates over 10MB of data, as opposed to only 2MB in Shasta. The higher amount of data communication in Cashmere leads to more load imbalance among the processes, thereby triggering more task stealing that compounds the communication costs.

Overall, applications in this category exhibit by far the largest performance gap between the two systems, with Cashmere suffering considerably due to the frequent synchronization and communication.

Application		Barnes (sm/lg)		LU (sm/lg)		CLU (sm/lg)		Ocean (sm/lg)	
Shasta	Lock/Flag Acquires (K)	68.7	274.8	0	0	0	0	1.2	0.8
	Barriers	9	9	129	129	129	129	328	248
	Read Misses (K)	128.9	477.0	50.6	199.4	49.8	199.2	28.2	37.9
	Write Misses (K)	51.7	210.0	24.6	98.3	0	0	0	0
	Upgrades (K)	112.1	419.1	0	0	24.6	98.3	26.9	37.3
	Data Fetches (K)	180.3	686.7	75.2	297.7	49.8	199.2)	27.9	37.7
	Messages (K)	985.6	3564.4	217.6	797.0	178.8	699.6	129.4	164.7
	Message Traffic (Mbytes)	77.7	289.8	26.2	101.7	18.5	73.4	11.3	14.9
Cashmere	Lock/Flag Acquires (K)	413.3	1836.3	0	0	0	0	1.2	0.8
	Barriers	9	9	129	129	129	129	328	248
	Read Faults (K)	71.3	253.3	18.5	56.0	3.7	12.1	16.0	14.3
	Write Faults (K)	112.0	462.9	5.6	25.0	1.8	6.9	11.0	10.8
	Twins (K)	10.4	39.1	0	0	0	0	0	0
	Page Transfers (K)	51.0	184.4	6.6	17.6	2.0	6.2	10.2	9.3
	Messages (K)	1185.6	4910.4	54.2	159.1	24.1	56.5	112.0	76.5
	Message Traffic (Mbytes)	425.0	1553	54.6	145.1	16.4	51.2	84.3	76.9

Application		Raytrace	Volrend	Water-NSQ (sm/lg)		Water-SP (sm/lg)	
Shasta	Lock/Flag Acquires (K)	119.9	9.2	73.9	144.2	0.2	0.2
	Barriers	1	3	12	12	12	12
	Read Misses (K)	68.2	4.8	121.2	285.2	36.2	52.9
	Write Misses (K)	51.3	1.1	8.1	14.4	0	0
	Upgrades (K)	0.2	2.0	42.6	98.5	15.9	26.4
	Data Fetches (K)	93.2	5.1	77.7	172.8	34.4	52.0
	Messages (K)	428.8	24.4	556.0	1201.1	165.0	254.8
	Message Traffic (Mbytes)	37.6	2.1	37.7	82.7	14.1	21.5
Cashmere	Lock/Flag Acquires (K)	120.9	9.2	73.9	144.1	0.2	0.3
	Barriers	1	3	12	12	12	12
	Read Faults (K)	134.1	1.7	14.1	26.6	6.0	8.6
	Write Faults (K)	143.0	7.7	50.1	123.6	3.5	4.8
	Twins (K)	5.8	5.5	5.6	2.9	0.3	0.3
	Page Transfers (K)	124.2	1.2	5.9	11.8	1.8	1.8
	Messages (K)	1659.0	46.8	345.0	845.1	21.6	28.7
	Message Traffic (Mbytes)	1027.3	10.4	51.2	102.3	14.6	24.2

Application		Em3d (sm/lg)		Gauss (sm/lg)		Ilink	Sor (sm/lg)		TSP
Shasta	Lock/Flag Acquires (K)	0	0	57.5	69.3	0	0	0	2.5
	Barriers	200	200	6	6	522	48	48	2
	Read Misses (K)	1286.1	3853.6	488.8	526.8	951.8	9.2	13.9	315.7
	Write Misses (K)	0	0	67.5	96.8	17.4	0	0	3.5
	Upgrades (K)	1187.3	3557.1	3.0	5.0	166.8	9.2	13.9	25.4
	Data Fetches (K)	1286.1	3853.6	216.0	307.9	447.9	9.2	13.9	98.5
	Messages (K)	5067.3	15168.1	581.4	811.8	2384.4	38.4	57.3	670.9
	Message Traffic (Mbytes)	491.4	1471.9	73.9	104.8	191.0	3.6	5.4	46.7
Cashmere	Lock/Flag Acquires (K)	0	0	54.1	65.2	0	0	0	2.5
	Barriers	200	200	6	6	512	48	48	2
	Read Faults (K)	45.0	129.8	71.6	78.1	85.6	0.3	0.5	11.2
	Write Faults (K)	41.3	116.8	10.1	13.1	29.5	4.8	6.0	8.7
	Twins (K)	0	0	0	0	4.1	0	0	0
	Page Transfers (K)	42.4	123.3	18.6	21.0	22.6	0.3	0.4	9.7
	Messages (K)	348.2	967.9	174.3	203.8	232.1	10.7	13.5	103.3
	Message Traffic (Mbytes)	348.5	1014.7	153.3	173.4	186.4	2.5	4.0	80.3

Table 2: Detailed statistics for Shasta and Cashmere with 16 processors. Statistics are shown for both smaller and larger data sets (sm/lg) where applicable.

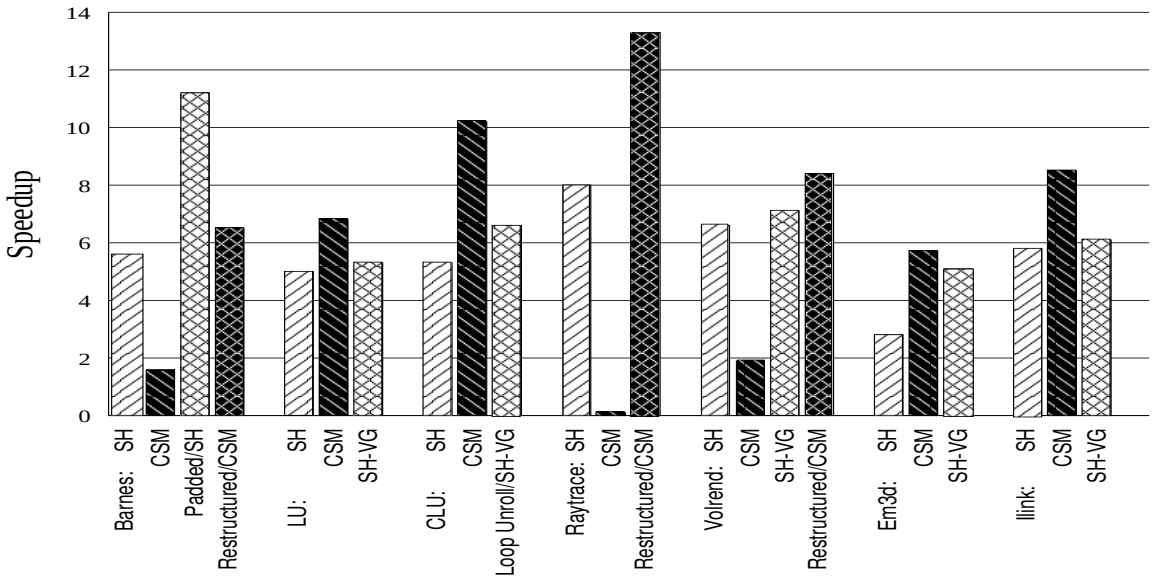

Figure 5: Speedups for the optimized applications at 16 processors. Measurements are taken for the larger data set, except where not possible (Raytrace, Volrend, and Ilink).

4.3 Performance Improvements through Program Modifications

The performance results presented in the previous section were for unmodified programs (except to eliminate a race in Barnes for Cashmere) that were taken from either the hardware or software shared memory domain. In most cases, better performance can be achieved by tailoring the application to the latencies and granularity of the underlying software system. In this section, we present the performance of some of the applications that have been modified for either Shasta or Cashmere.

Figure 5 presents the speedups for the modified applications along with the unmodified results for 16 processor runs. The large dataset size is used where possible. The corresponding execution time breakdowns are shown in Figure 6.

4.3.1 Modifications for Shasta

The modifications we consider for Shasta are guaranteed not to alter program correctness, and can therefore be applied safely without a deep understanding of the application. This is consistent with Shasta's philosophy of transparency and simple portability. The three types of changes we use are variable granularity hints [16], the addition of padding in data structures, and the use of compiler options to reduce instrumentation overhead.

For variable granularity hints, we use a special shared-memory allocator provided by Shasta that allows one to specify the block size for the corresponding region of memory. By allocating certain regions in this manner, the application can cause data to be fetched in large units for important data structures that are accessed in a coarse-grain manner or are mostly-read. Table 3 lists the applications that benefit from using variable granularity, the data structures on which it was used, and the increased block size. As an example of the benefit of variable granularity, the performance of EM3D improves by a factor of 1.8 and CLU by a factor of 1.2 for the large input set on 16 processors (results labeled as "SH-VG" in Figures 5 and 6).

Another change that can sometimes improve performance is padding elements of important data structures. For example, in the Barnes-Hut application, information on each body is stored in a structure which is allocated out of one large array. Since the body structure is 120 bytes, there is some false sharing between different bodies. Shasta's performance improves significantly (by a factor of 1.9 on the large input set for 16 processors)

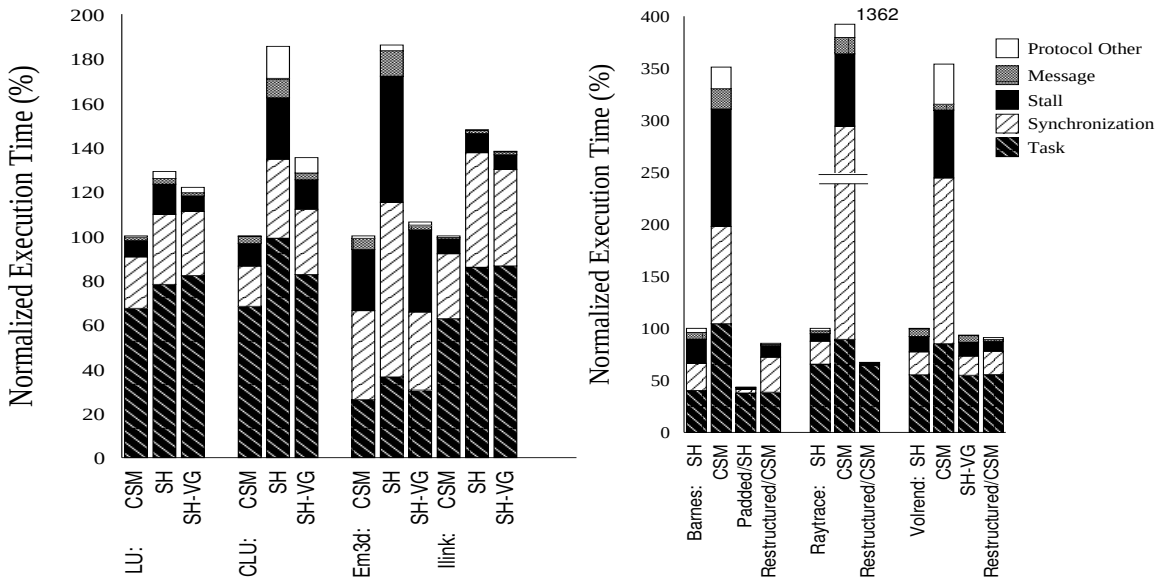

Figure 6: Execution breakdown for the optimized applications at 16 processors. Raytrace, Volrend, and Ilink execute on the smaller data set; all others are measured on the larger data set.

	selected data structure(s)	**block size (bytes)**
LU	matrix array	2048
CLU	matrix block	2048
Volrend	opacity, normal maps	1024
EM3D	node and data array	8192
ILINK	all data	1024

Table 3: Variable block sizes used for Shasta.

by padding the body structure to 128 bytes (labeled as "Padded" in Figures 5 and 6).

A final modification involves using compiler options to reduce instrumentation overhead. Existing compilers typically unroll inner loops to improve instruction scheduling and reduce looping overheads. Batching of checking code is especially effective for unrolled loops, since unrolling increases the number of neighboring loads and stores in the loop bodies. In CLU, instrumentation overhead is still high despite the batching, because the inner loop is scheduled so effectively. The checking overhead is reduced significantly (from 55% to 36% on a uniprocessor with the large input set) by using a compiler option that increases the unrolling of the inner loop from the default four iterations to eight iterations (labelled as "Loop Unroll" in Figure 5 and 6).

There are a large class of other optimizations that would improve application performance under Shasta. However, we have limited our investigation to simple hint optimizations to emphasize the ease of portability of applications from hardware multiprocessors to Shasta.

4.3.2 Modifications for Cashmere

The modifications made to tune the applications for Cashmere aim to reduce the frequency of synchronization or to increase the granularity of sharing. Unlike the changes made for Shasta, the Cashmere modifications are not in the form of hints and require some real understanding of the application in order to maintain correctness. We

have made changes to three applications that exhibit particularly poor performance under Cashmere (results are labeled as "Restructured/CSM" in Figures 5 and 6).

The major source of overhead in Barnes-Hut is in the tree building phase. The application requires processors to position their bodies into a tree data structure of cells, resulting in a large number of scattered accesses to shared memory. In addition, the algorithm requires processors to synchronize in a very fine-grain manner in order to avoid race conditions. The resulting false sharing and fine-grain synchronization cause the tree-building phase to run much slower in parallel under Cashmere than in the sequential execution. While parallel tree-building algorithms suitable for page-based S-DSM [12] exist, we have chosen to use the simple approach of computing the tree sequentially (it constitutes 2% of the total sequential execution time). Building the tree sequentially does, however, have the disadvantage of increasing the memory requirements on the main node and limiting the largest problem size that can be run.

A second source of overhead comes from a parallel reduction in the main computation loop. The reduction modifies two shared variables in a critical section based on per-processor values for these variable. Performance is reduced because of critical section dilation due to page faults. We have modified the code to compute the reduction sequentially on a single processor. These changes result in a factor of three decrease in execution time.

Raytrace is in reality a highly parallel application. There is very little sharing and the only necessary synchronization constructs are per-processor locks on the processor work queues. However, the original version contains some additional locking code that protects a counter used for debugging purposes, as described in Section 4.2.3. Eliminating the locking code and counter update reduces the running time from 71 seconds to 3.7 seconds on 16 processors and the amount of data transferred from 1GByte down to 17MBytes. The magnitude of this improvement illustrates the sensitivity of page-based S-DSM to fine-grain synchronization.

The performance degradation in Volrend comes from false sharing on the task queue data structure as well as the small granularity of work. We have modified the application to change the granularity of tasks as well as to eliminate false sharing in the task queue by padding. Our changes result in a runtime reduction from 2.1 seconds to 0.46 seconds. Similarly, the amount of data transferred drops from 22Mbytes to 5.6Mbytes.

Additional optimizations that would improve the performance of these and other applications in our suite on a page-based system can be implemented [12]. In general, if the size of the coherence block is taken into account in structuring the application, most applications can perform well on page-based systems. Restructuring applications tuned for hardware DSM systems does, however, require knowledge of the underlying computation or data structures.

4.4 Summary of Results

This section provided an in-depth comparison and analysis of the performance of two software DSM systems, Shasta and Cashmere. We summarize our results using the geometric mean of the relative speedups on the two systems. For the eight applications (unmodified Splash-2) that were written and tuned for hardware DSM systems, Shasta exhibits a 1.6 times performance advantage over Cashmere. Most of this difference comes from one application (Raytrace) for which the performance of the two systems differs by a factor of 13. Using the same metric, for the five programs that were written or tuned with page-based DSM in mind, Cashmere exhibits a 1.3 times performance advantage over Shasta. After we allow modifications to the applications, Cashmere performs 1.15 times better than Shasta over all 13 applications. However, it is important to emphasize that the modifications we considered for Shasta were in the form of hints that do not affect application correctness or require detailed application knowledge, while the modifications we considered for Cashmere often required changes in the parallelization strategy.

5 Related Work

There is a large body of literature on S-DSM that has had an impact on the design of the Cashmere and Shasta systems. The focus of this paper is to understand the performance tradeoffs of fine-grain vs. coarse-grain software shared memory rather than to design or study a particular S-DSM system in isolation.

Iftode *et al.* [11] have characterized the performance and sources of overhead of a large number of S-DSM applications, while Jiang *et al.* [12] have provided insights into the restructuring necessary to achieve good performance for a similar S-DSM application suite. Our work builds on their's by providing insight on how a similar class of programs performs under both fine-grain and coarse-grain S-DSM. Also, we use actual systems implemented on a state-of-the-art cluster, allowing us to capture details not present in a simulation environment.

Researchers at Wisconsin and Princeton [23] have also studied the tradeoffs between fine- and coarse-grain S-DSM systems, but our studies have a number of differences. First, we have studied SMP-aware systems running on clusters of SMPs. Second, the Wisconsin/Princeton platform uses custom hardware not available in commodity systems to provide fine-grain access control. Their fine-grain performance results therefore do not include software checking overhead, which limited performance in several of our applications. In addition, the custom hardware delivers an access control fault in only 5 μs, which is fourteen times faster than the delivery of a page fault on our platform. Third, the processors in our cluster are an order of magnitude faster than those in the Wisconsin/Princeton cluster (400MHz vs 66MHz), while our network is only 3-4 times better in latency and bandwidth, thus increasing the relative cost of communication. All of these differences have manifested themselves in a number of ways in our performance results for both the fine-grain and coarse-grain systems.

Some of the results of our study mirror those of the Wisconsin/Princeton study, but others offer new insight into the granularity issue. For example, both Raytrace and Volrend perform well on the coarse-grain protocol in the Wisconsin/Princeton study, but perform very poorly on Cashmere in our study. The performance gap can be attributed to our fast hardware platform, which causes accelerated synchronization and in turn magnifies the effect of unnecessary data transferred in a coarse-grain protocol. More favorably for coarse-grain protocols, we also found that an SMP-aware implementation can greatly mitigate the effects of false sharing.

We believe that Cashmere and Shasta are among the most efficient S-DSMs in their class. There are still relatively few S-DSMs that are SMP-aware and capable of executing on commodity hardware. The Sirocco system [20] is a fine-grain S-DSM that uses an SMP-aware protocol, but its instrumentation overheads are much higher than Shasta's. SoftFlash [7] was one of the first page-based implementations designed for SMP clusters. The SoftFlash results showed that intra-node synchronization could be excessive. Cashmere-2L [21], however, combines existing techniques with a novel incoming diff operation to eliminate most intra-node synchronization. HLRC-SMP [15] is a more recent protocol that shares several similarities with Cashmere-2L. The Cashmere-2L protocol, however, has been optimized to take advantage of the Memory Channel network and allows home nodes to migrate to active writers, thereby potentially reducing twin/diff overhead.

6 Conclusions

In this paper, we have examined the performance tradeoffs between fine-grain and coarse-grain S-DSM in the context of two state-of-the-art systems: Shasta and Cashmere. In general, we found that the fine-grain, instrumentation-based approach to S-DSM offers a higher degree of robustness and superior performance in the presence of fine-grain synchronization, while the coarse-grain, VM-based approach offers higher performance when coarse-grain synchronization is used.

The performance of applications running under Shasta is most affected by the instrumentation overhead and by the smaller default block size when accessing data at a coarse granularity. Conversely, for Cashmere, the main sources of overhead are critical section dilation in the presence of fine-grain synchronization, and the communication of unneeded data in computations with fine-grain data access. Standard programming idioms such as work-queues, parallel reductions, and atomic counters can cause excessive communication overhead if they are not tuned for coarse-grain systems. However, a number of applications with false sharing at a page-level performed better than expected on Cashmere because its SMP-aware protocol enabled most or all of the false sharing effects to be handled in hardware. Finally, we found that most of the remaining performance differences between Shasta and Cashmere could be eliminated by program modifications that take the coherence granularity into account.

References

[1] S. Adve and K. Gharachorloo. Shared Memory Consistency Models: A Tutorial. In *Computer*, 29(12):66–76, December 1996.

[2] C. Amza, A. L. Cox, S. Dwarkadas, P. Keleher, H. Lu, R. Rajamony, W. Yu, and W. Zwaenepoel. TreadMarks: Shared Memory Computing on Networks of Workstations. In *Computer*, 29(2):18-28, February 1996.

[3] J. B. Carter, J. K. Bennett, and W. Zwaenepoel. Implementation and Performance of Munin. In *Proceedings of the Thirteenth ACM Symposium on Operating Systems Principles*, pages 152–164, October 1991.

[4] D. Culler, A. Dusseau, S. Goldstein, A. Krishnamurthy, S. Lumetta, T. von Eicken, and K. Yelick. Parallel Programming in Split-C. In *Proceedings of Supercomputing '93*, pp. 262–273, Nov. 1993.

[5] S. Dwarkadas, R. W. Cottingham, A. L. Cox, P. Keleher, A. A. Scaffer, and W. Zwaenepoel. Parallelization of General Linkage Analysis Problems. In *Human Heredity*, 44:127–141, July 1994.

[6] S. Dwarkadas, K. Gharachorloo, L. Kontothanassis, D. J. Scales, M. L. Scott, and R. Stets. Comparative Evaluation of Fine- and Coarse-Grain Approaches for Software Distributed Shared Memory. University of Rochester CS TR 699, October 1998. Also available as Western Research Lab TR 98/7.

[7] A. Erlichson, N. Nuckolls, G. Chesson, and J. Hennessy. SoftFLASH: Analyzing the Performance of Clustered Distributed Virtual Shared Memory. In *Proceedings of the Seventh International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 210–220, October 1996.

[8] K. Gharachorloo, D. Lenoski, J. Laudon, P. Gibbons, A. Gupta, and J. L. Hennessy. Memory Consistency and Event Ordering in Scalable Shared-Memory Multiprocessors. In *Proceedings of the Seventeenth Annual International Symposium on Computer Architecture*, pages 15–26, May 1990.

[9] R. Gillett. Memory Channel: An Optimized Cluster Interconnect. *IEEE Micro*, 16(2), February 1996.

[10] L. Iftode, C. Dubnicki, E. W. Felten, and K. Li. Improving Release-Consistent Shared Virtual Memory Using Automatic Update. In *Proceedings of the Second Conference on High Performance Computer Architecture*, February 1996.

[11] L. Iftode, J. P. Singh, and K. Li. Understanding Application Performance on Shared Virtual Memory. In *Proceedings of the Twenty-Third Annual International Symposium on Computer Architecture*, May 1996.

[12] D. Jiang, H. Shan, and J. P. Singh. Application Restructuring and Performance Portability on Shared Virtual Memory and Hardware-Coherence Multiprocessors. In *Proceedings of the Sixth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, June 1997.

[13] L. Kontothanassis, G. Hunt, R. Stets, N. Hardavellas, M. Cierniak, S. Parthasarathy, W. M. Jr., S. Dwarkadas, and M. L. Scott. VM-Based Shared Memory on Low-Latency, Remote-Memory-Access Networks. In *Proceedings of the Twenty-Fourth Annual International Symposium on Computer Architecture*, June 1997.

[14] K. Li and P. Hudak. Memory Coherence in Shared Virtual Memory Systems. *ACM Transactions on Computer Systems*, 7(4):321-359, November 1989.

[15] R. Samanta, A. Bilas, L. Iftode, and J. Singh. Home-Based SVM Protocols for SMP Clusters: Design and Performance. In *Proceedings of the Fourth Conference on High Performance Computer Architecture*, pages 113–124, February 1998.

[16] D. Scales, K. Gharachorloo, and C. Thekkath. Shasta: A Low-Overhead Software-Only Approach to Fine-Grain Shared Memory. In *Proceedings of the Seventh International Conference on Architectural Support for Programming Languages and Operating Systems*, October 1996.

[17] D. Scales and K. Gharachorloo. Toward Transparent and Efficient Software Distributed Shared Memory. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*, October 1997.

[18] D. J. Scales, K. Gharachorloo, and A. Aggarwal. Fine-Grain Software Distributed Shared Memory on SMP Clusters. In *Proceedings of the Fourth Conference on High Performance Computer Architecture*, February 1998.

[19] I. Schoinas, B. Falsafi, A. R. Lebeck, S. K. Reinhardt, J. R. Larus, and D. A. Wood. Fine-grain Access Control for Distributed Shared Memory. In *Proceedings of the Sixth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 297–306, October 1994.

[20] I. Schoinas, B. Falsafi, M. Hill, J. Larus, and D. Wood. Sirocco: Cost-Effective Fine-Grain Distributed Shared Memory. In *Proceedings of PACT '98*, October 1998.

[21] R. Stets, S. Dwarkadas, N. Hardavellas, G. Hunt, L. Kontothanassis, S. Parthasarathy, and M. L. Scott. CSM-2L: Software Coherent Shared Memory on a Clustered Remote-Write Network. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*, October 1997.

[22] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta. Methodological Considerations and Characterization of the SPLASH-2 Parallel Application Suite. In *Proceedings of the Twenty-Second Annual International Symposium on Computer Architecture*, June 1995.

[23] Y. Zhou, L. Iftode, J. P. Singh, K. Li, B. R. Toonen, I. Schoinas, M. D. Hill, and D. A. Wood. Relaxed Consistency and Coherence Granularity in DSM Systems: A Performance Evaluation. In *Proceedings of the Sixth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, June 1997.

WRL Research Reports

‘‘Titan System Manual.’’ **Michael J. K. Nielsen.** WRL Research Report 86/1, September 1986.

‘‘Global Register Allocation at Link Time.’’ **David W. Wall.** WRL Research Report 86/3, October 1986.

‘‘Optimal Finned Heat Sinks.’’ **William R. Hamburgen.** WRL Research Report 86/4, October 1986.

‘‘The Mahler Experience: Using an Intermediate Language as the Machine Description.’’ **David W. Wall and Michael L. Powell.** WRL Research Report 87/1, August 1987.

‘‘The Packet Filter: An Efficient Mechanism for User-level Network Code.’’ **Jeffrey C. Mogul, Richard F. Rashid, Michael J. Accetta.** WRL Research Report 87/2, November 1987.

‘‘Fragmentation Considered Harmful.’’ **Christopher A. Kent, Jeffrey C. Mogul.** WRL Research Report 87/3, December 1987.

‘‘Cache Coherence in Distributed Systems.’’ **Christopher A. Kent.** WRL Research Report 87/4, December 1987.

‘‘Register Windows vs. Register Allocation.’’ **David W. Wall.** WRL Research Report 87/5, December 1987.

‘‘Editing Graphical Objects Using Procedural Representations.’’ **Paul J. Asente.** WRL Research Report 87/6, November 1987.

‘‘The USENET Cookbook: an Experiment in Electronic Publication.’’ **Brian K. Reid.** WRL Research Report 87/7, December 1987.

‘‘MultiTitan: Four Architecture Papers.’’ **Norman P. Jouppi, Jeremy Dion, David Boggs, Michael J. K. Nielsen.** WRL Research Report 87/8, April 1988.

‘‘Fast Printed Circuit Board Routing.’’ **Jeremy Dion.** WRL Research Report 88/1, March 1988.

‘‘Compacting Garbage Collection with Ambiguous Roots.’’ **Joel F. Bartlett.** WRL Research Report 88/2, February 1988.

‘‘The Experimental Literature of The Internet: An Annotated Bibliography.’’ **Jeffrey C. Mogul.** WRL Research Report 88/3, August 1988.

‘‘Measured Capacity of an Ethernet: Myths and Reality.’’ **David R. Boggs, Jeffrey C. Mogul, Christopher A. Kent.** WRL Research Report 88/4, September 1988.

‘‘Visa Protocols for Controlling Inter-Organizational Datagram Flow: Extended Description.’’ **Deborah Estrin, Jeffrey C. Mogul, Gene Tsudik, Kamaljit Anand.** WRL Research Report 88/5, December 1988.

‘‘SCHEME->C A Portable Scheme-to-C Compiler.’’ **Joel F. Bartlett.** WRL Research Report 89/1, January 1989.

‘‘Optimal Group Distribution in Carry-Skip Adders.’’ **Silvio Turrini.** WRL Research Report 89/2, February 1989.

‘‘Precise Robotic Paste Dot Dispensing.’’ **William R. Hamburgen.** WRL Research Report 89/3, February 1989.

‘‘Simple and Flexible Datagram Access Controls for Unix-based Gateways.’’ **Jeffrey C. Mogul.** WRL Research Report 89/4, March 1989.

‘‘Spritely NFS: Implementation and Performance of Cache-Consistency Protocols.’’ **V. Srinivasan and Jeffrey C. Mogul.** WRL Research Report 89/5, May 1989.

‘‘Available Instruction-Level Parallelism for Superscalar and Superpipelined Machines.’’ **Norman P. Jouppi and David W. Wall.** WRL Research Report 89/7, July 1989.

‘‘A Unified Vector/Scalar Floating-Point Architecture.’’ **Norman P. Jouppi, Jonathan Bertoni, and David W. Wall.** WRL Research Report 89/8, July 1989.

''Architectural and Organizational Tradeoffs in the Design of the MultiTitan CPU.'' **Norman P. Jouppi.** WRL Research Report 89/9, July 1989.

''Integration and Packaging Plateaus of Processor Performance.'' **Norman P. Jouppi.** WRL Research Report 89/10, July 1989.

''A 20-MIPS Sustained 32-bit CMOS Microprocessor with High Ratio of Sustained to Peak Performance.'' **Norman P. Jouppi and Jeffrey Y. F. Tang.** WRL Research Report 89/11, July 1989.

''The Distribution of Instruction-Level and Machine Parallelism and Its Effect on Performance.'' **Norman P. Jouppi.** WRL Research Report 89/13, July 1989.

''Long Address Traces from RISC Machines: Generation and Analysis.'' **Anita Borg, R.E.Kessler, Georgia Lazana, and David W. Wall.** WRL Research Report 89/14, September 1989.

''Link-Time Code Modification.'' **David W. Wall.** WRL Research Report 89/17, September 1989.

''Noise Issues in the ECL Circuit Family.'' **Jeffrey Y.F. Tang and J. Leon Yang.** WRL Research Report 90/1, January 1990.

''Efficient Generation of Test Patterns Using Boolean Satisfiablilty.'' **Tracy Larrabee.** WRL Research Report 90/2, February 1990.

''Two Papers on Test Pattern Generation.'' **Tracy Larrabee.** WRL Research Report 90/3, March 1990.

''Virtual Memory vs. The File System.'' **Michael N. Nelson.** WRL Research Report 90/4, March 1990.

''Efficient Use of Workstations for Passive Monitoring of Local Area Networks.'' **Jeffrey C. Mogul.** WRL Research Report 90/5, July 1990.

''A One-Dimensional Thermal Model for the VAX 9000 Multi Chip Units.'' **John S. Fitch.** WRL Research Report 90/6, July 1990.

''1990 DECWRL/Livermore Magic Release.'' **Robert N. Mayo, Michael H. Arnold, Walter S. Scott, Don Stark, Gordon T. Hamachi.** WRL Research Report 90/7, September 1990.

''Pool Boiling Enhancement Techniques for Water at Low Pressure.'' **Wade R. McGillis, John S. Fitch, William R. Hamburgen, Van P. Carey.** WRL Research Report 90/9, December 1990.

''Writing Fast X Servers for Dumb Color Frame Buffers.'' **Joel McCormack.** WRL Research Report 91/1, February 1991.

''A Simulation Based Study of TLB Performance.'' **J. Bradley Chen, Anita Borg, Norman P. Jouppi.** WRL Research Report 91/2, November 1991.

''Analysis of Power Supply Networks in VLSI Circuits.'' **Don Stark.** WRL Research Report 91/3, April 1991.

''TurboChannel T1 Adapter.'' **David Boggs.** WRL Research Report 91/4, April 1991.

''Procedure Merging with Instruction Caches.'' **Scott McFarling.** WRL Research Report 91/5, March 1991.

''Don't Fidget with Widgets, Draw!.'' **Joel Bartlett.** WRL Research Report 91/6, May 1991.

''Pool Boiling on Small Heat Dissipating Elements in Water at Subatmospheric Pressure.'' **Wade R. McGillis, John S. Fitch, William R. Hamburgen, Van P. Carey.** WRL Research Report 91/7, June 1991.

''Incremental, Generational Mostly-Copying Garbage Collection in Uncooperative Environments.'' **G. May Yip.** WRL Research Report 91/8, June 1991.

''Interleaved Fin Thermal Connectors for Multichip Modules.'' **William R. Hamburgen.** WRL Research Report 91/9, August 1991.

''Experience with a Software-defined Machine Architecture.'' **David W. Wall.** WRL Research Report 91/10, August 1991.

‘‘Network Locality at the Scale of Processes.’’ **Jeffrey C. Mogul.** WRL Research Report 91/11, November 1991.

‘‘Cache Write Policies and Performance.’’ **Norman P. Jouppi.** WRL Research Report 91/12, December 1991.

‘‘Packaging a 150 W Bipolar ECL Microprocessor.’’ **William R. Hamburgen, John S. Fitch.** WRL Research Report 92/1, March 1992.

‘‘Observing TCP Dynamics in Real Networks.’’ **Jeffrey C. Mogul.** WRL Research Report 92/2, April 1992.

‘‘Systems for Late Code Modification.’’ **David W. Wall.** WRL Research Report 92/3, May 1992.

‘‘Piecewise Linear Models for Switch-Level Simulation.’’ **Russell Kao.** WRL Research Report 92/5, September 1992.

‘‘A Practical System for Intermodule Code Optimization at Link-Time.’’ **Amitabh Srivastava and David W. Wall.** WRL Research Report 92/6, December 1992.

‘‘A Smart Frame Buffer.’’ **Joel McCormack & Bob McNamara.** WRL Research Report 93/1, January 1993.

‘‘Recovery in Spritely NFS.’’ **Jeffrey C. Mogul.** WRL Research Report 93/2, June 1993.

‘‘Tradeoffs in Two-Level On-Chip Caching.’’ **Norman P. Jouppi & Steven J.E. Wilton.** WRL Research Report 93/3, October 1993.

‘‘Unreachable Procedures in Object-oriented Programing.’’ **Amitabh Srivastava.** WRL Research Report 93/4, August 1993.

‘‘An Enhanced Access and Cycle Time Model for On-Chip Caches.’’ **Steven J.E. Wilton and Norman P. Jouppi.** WRL Research Report 93/5, July 1994.

‘‘Limits of Instruction-Level Parallelism.’’ **David W. Wall.** WRL Research Report 93/6, November 1993.

‘‘Fluoroelastomer Pressure Pad Design for Microelectronic Applications.’’ **Alberto Makino, William R. Hamburgen, John S. Fitch.** WRL Research Report 93/7, November 1993.

‘‘A 300MHz 115W 32b Bipolar ECL Microprocessor.’’ **Norman P. Jouppi, Patrick Boyle, Jeremy Dion, Mary Jo Doherty, Alan Eustace, Ramsey Haddad, Robert Mayo, Suresh Menon, Louis Monier, Don Stark, Silvio Turrini, Leon Yang, John Fitch, William Hamburgen, Russell Kao, and Richard Swan.** WRL Research Report 93/8, December 1993.

‘‘Link-Time Optimization of Address Calculation on a 64-bit Architecture.’’ **Amitabh Srivastava, David W. Wall.** WRL Research Report 94/1, February 1994.

‘‘ATOM: A System for Building Customized Program Analysis Tools.’’ **Amitabh Srivastava, Alan Eustace.** WRL Research Report 94/2, March 1994.

‘‘Complexity/Performance Tradeoffs with Non-Blocking Loads.’’ **Keith I. Farkas, Norman P. Jouppi.** WRL Research Report 94/3, March 1994.

‘‘A Better Update Policy.’’ **Jeffrey C. Mogul.** WRL Research Report 94/4, April 1994.

‘‘Boolean Matching for Full-Custom ECL Gates.’’ **Robert N. Mayo, Herve Touati.** WRL Research Report 94/5, April 1994.

‘‘Software Methods for System Address Tracing: Implementation and Validation.’’ **J. Bradley Chen, David W. Wall, and Anita Borg.** WRL Research Report 94/6, September 1994.

‘‘Performance Implications of Multiple Pointer Sizes.’’ **Jeffrey C. Mogul, Joel F. Bartlett, Robert N. Mayo, and Amitabh Srivastava.** WRL Research Report 94/7, December 1994.

‘‘How Useful Are Non-blocking Loads, Stream Buffers, and Speculative Execution in Multiple Issue Processors?.’’ **Keith I. Farkas, Norman P. Jouppi, and Paul Chow.** WRL Research Report 94/8, December 1994.

‘‘Drip: A Schematic Drawing Interpreter.’’ **Ramsey W. Haddad.** WRL Research Report 95/1, March 1995.

‘‘Recursive Layout Generation.’’ **Louis M. Monier, Jeremy Dion.** WRL Research Report 95/2, March 1995.

‘‘Contour: A Tile-based Gridless Router.’’ **Jeremy Dion, Louis M. Monier.** WRL Research Report 95/3, March 1995.

‘‘The Case for Persistent-Connection HTTP.’’ **Jeffrey C. Mogul.** WRL Research Report 95/4, May 1995.

‘‘Network Behavior of a Busy Web Server and its Clients.’’ **Jeffrey C. Mogul.** WRL Research Report 95/5, October 1995.

‘‘The Predictability of Branches in Libraries.’’ **Brad Calder, Dirk Grunwald, and Amitabh** Srivastava. WRL Research Report 95/6, October 1995.

‘‘Shared Memory Consistency Models: A Tutorial.’’ **Sarita V. Adve, Kourosh Gharachorloo.** WRL Research Report 95/7, September 1995.

‘‘Eliminating Receive Livelock in an Interrupt-driven Kernel.’’ **Jeffrey C. Mogul and K. K. Ramakrishnan.** WRL Research Report 95/8, December 1995.

‘‘Memory Consistency Models for Shared-Memory Multiprocessors.’’ **Kourosh Gharachorloo.** WRL Research Report 95/9, December 1995.

‘‘Register File Design Considerations in Dynamically Scheduled Processors.’’ **Keith I. Farkas, Norman P. Jouppi, Paul Chow.** WRL Research Report 95/10, November 1995.

‘‘Optimization in Permutation Spaces.’’ **Silvio** Turrini. WRL Research Report 96/1, November 1996.

‘‘Shasta: A Low Overhead, Software-Only Approach for Supporting Fine-Grain Shared Memory.’’ **Daniel J. Scales, Kourosh Gharachorloo, and Chandramohan A. Thekkath.** WRL Research Report 96/2, November 1996.

‘‘Efficient Procedure Mapping using Cache Line Coloring.’’ **Amir H. Hashemi, David R. Kaeli, and Brad Calder.** WRL Research Report 96/3, October 1996.

‘‘Optimizations and Placement with the Genetic Workbench.’’ **Silvio Turrini.** WRL Research Report 96/4, November 1996.

‘‘Memory-system Design Considerations for Dynamically-scheduled Processors.’’ **Keith I. Farkas, Paul Chow, Norman P. Jouppi, and Zvonko Vranesic.** WRL Research Report 97/1, February 1997.

‘‘Performance of the Shasta Distributed Shared Memory Protocol.’’ **Daniel J. Scales and Kourosh Gharachorloo.** WRL Research Report 97/2, February 1997.

‘‘Fine-Grain Software Distributed Shared Memory on SMP Clusters.’’ **Daniel J. Scales, Kourosh Gharachorloo, and Anshu Aggarwal.** WRL Research Report 97/3, February 1997.

‘‘Potential benefits of delta encoding and data compression for HTTP.’’ **Jeffrey C. Mogul, Fred Douglis, Anja Feldmann, and Balachander Krishnamurthy.** WRL Research Report 97/4, July 1997.

‘‘Neon: A (Big) (Fast) Single-Chip 3D Workstation Graphics Accelerator.’’ **Joel McCormack, Robert McNamara, Christopher Gianos, Larry Seiler, Norman P. Jouppi, and Ken Correll.** WRL Research Report 98/1, August 1998.

‘‘Efficient Dynamic Procedure Placement.’’ **Daniel J. Scales.** WRL Research Report 98/5, August 1998.

‘‘Scalable kernel performance for Internet servers under realistic loads.’’ **Gaurav Banga and Jeffrey C. Mogul.** WRL Research Report 98/6, October 1998.

WRL Technical Notes

‘‘TCP/IP PrintServer: Print Server Protocol.’’ **Brian K. Reid and Christopher A. Kent.** WRL Technical Note TN-4, September 1988.

‘‘TCP/IP PrintServer: Server Architecture and Implementation.’’ **Christopher A. Kent.** WRL Technical Note TN-7, November 1988.

‘‘Smart Code, Stupid Memory: A Fast X Server for a Dumb Color Frame Buffer.’’ **Joel McCormack.** WRL Technical Note TN-9, September 1989.

‘‘Why Aren’t Operating Systems Getting Faster As Fast As Hardware?.’’ **John Ousterhout.** WRL Technical Note TN-11, October 1989.

‘‘Mostly-Copying Garbage Collection Picks Up Generations and C++.’’ **Joel F. Bartlett.** WRL Technical Note TN-12, October 1989.

‘‘Characterization of Organic Illumination Systems.’’ **Bill Hamburgen, Jeff Mogul, Brian Reid, Alan Eustace, Richard Swan, Mary Jo Doherty, and Joel Bartlett.** WRL Technical Note TN-13, April 1989.

‘‘Improving Direct-Mapped Cache Performance by the Addition of a Small Fully-Associative Cache and Prefetch Buffers.’’ **Norman P. Jouppi.** WRL Technical Note TN-14, March 1990.

‘‘Limits of Instruction-Level Parallelism.’’ **David W. Wall.** WRL Technical Note TN-15, December 1990.

‘‘The Effect of Context Switches on Cache Performance.’’ **Jeffrey C. Mogul and Anita Borg.** WRL Technical Note TN-16, December 1990.

‘‘MTOOL: A Method For Detecting Memory Bottlenecks.’’ **Aaron Goldberg and John** Hennessy. WRL Technical Note TN-17, December 1990.

‘‘Predicting Program Behavior Using Real or Estimated Profiles.’’ **David W. Wall.** WRL Technical Note TN-18, December 1990.

‘‘Cache Replacement with Dynamic Exclusion.’’ **Scott McFarling.** WRL Technical Note TN-22, November 1991.

‘‘Boiling Binary Mixtures at Subatmospheric Pressures.’’ **Wade R. McGillis, John S. Fitch, William R. Hamburgen, Van P. Carey.** WRL Technical Note TN-23, January 1992.

‘‘A Comparison of Acoustic and Infrared Inspection Techniques for Die Attach.’’ **John S. Fitch.** WRL Technical Note TN-24, January 1992.

‘‘TurboChannel Versatec Adapter.’’ **David Boggs.** WRL Technical Note TN-26, January 1992.

‘‘A Recovery Protocol For Spritely NFS.’’ **Jeffrey C. Mogul.** WRL Technical Note TN-27, April 1992.

‘‘Electrical Evaluation Of The BIPS-0 Package.’’ **Patrick D. Boyle.** WRL Technical Note TN-29, July 1992.

‘‘Transparent Controls for Interactive Graphics.’’ **Joel F. Bartlett.** WRL Technical Note TN-30, July 1992.

‘‘Design Tools for BIPS-0.’’ **Jeremy Dion & Louis Monier.** WRL Technical Note TN-32, December 1992.

‘‘Link-Time Optimization of Address Calculation on a 64-Bit Architecture.’’ **Amitabh Srivastava and David W. Wall.** WRL Technical Note TN-35, June 1993.

‘‘Combining Branch Predictors.’’ **Scott McFarling.** WRL Technical Note TN-36, June 1993.

‘‘Boolean Matching for Full-Custom ECL Gates.’’ **Robert N. Mayo and Herve Touati.** WRL Technical Note TN-37, June 1993.

‘‘Piecewise Linear Models for Rsim.’’ **Russell Kao, Mark Horowitz.** WRL Technical Note TN-40, December 1993.

‘‘Speculative Execution and Instruction-Level Parallelism.’’ **David W. Wall.** WRL Technical Note TN-42, March 1994.

‘‘Ramonamap - An Example of Graphical Groupware.’’ **Joel F. Bartlett.** WRL Technical Note TN-43, December 1994.

‘‘ATOM: A Flexible Interface for Building High Performance Program Analysis Tools.’’ **Alan Eustace and Amitabh Srivastava.** WRL Technical Note TN-44, July 1994.

‘‘Circuit and Process Directions for Low-Voltage Swing Submicron BiCMOS.’’ **Norman P. Jouppi, Suresh Menon, and Stefanos Sidiropoulos.** WRL Technical Note TN-45, March 1994.

‘‘Experience with a Wireless World Wide Web Client.’’ **Joel F. Bartlett.** WRL Technical Note TN-46, March 1995.

‘‘I/O Component Characterization for I/O Cache Designs.’’ **Kathy J. Richardson.** WRL Technical Note TN-47, April 1995.

‘‘Attribute caches.’’ **Kathy J. Richardson, Michael J. Flynn.** WRL Technical Note TN-48, April 1995.

‘‘Operating Systems Support for Busy Internet Servers.’’ **Jeffrey C. Mogul.** WRL Technical Note TN-49, May 1995.

‘‘The Predictability of Libraries.’’ **Brad Calder, Dirk Grunwald, Amitabh Srivastava.** WRL Technical Note TN-50, July 1995.

‘‘Simultaneous Multithreading: A Platform for Next-generation Processors.’’ **Susan J. Eggers, Joel Emer, Henry M. Levy, Jack L. Lo, Rebecca Stamm and Dean M. Tullsen.** WRL Technical Note TN-52, March 1997.

‘‘Reducing Compulsory and Capacity Misses.’’ **Norman P. Jouppi.** WRL Technical Note TN-53, August 1990.

‘‘The Itsy Pocket Computer Version 1.5: User’s Manual.’’ **Marc A. Viredaz.** WRL Technical Note TN-54, July 1998.

‘‘The Memory Daughter-Card Version 1.5: User’s Manual.’’ **Marc A. Viredaz.** WRL Technical Note TN-55, July 1998.

WRL Research Reports and Technical Notes are available on the World Wide Web, from `http://www.research.digital.com/wrl/techreports/index.html`.