

An Evaluation of an OS-Based Coherence Scheme for Tiled CMPs

Christian Fensch · Marcelo Cintra

Received: 22 January 2009 / Accepted: 9 December 2010 / Published online: 29 December 2010
© Springer Science+Business Media, LLC 2010

Abstract The interconnect mechanisms (shared bus or crossbar) used in current chip-multiprocessors (CMPs) are expected to become a bottleneck that prevents these architectures from scaling to a larger number of cores. Tiled CMPs offer better scalability by integrating relatively simple cores with a lightweight point-to-point interconnect. However, such interconnects make snooping impractical and, thus, require alternative solutions to cache coherence. In this article, we investigate a novel, cost-effective mechanism to support shared-memory parallel applications that forgoes hardware maintained cache coherence. This mechanism is based on the key ideas that mapping of lines to physical caches is done at the page level with OS support and that hardware supports remote cache accesses. We extend our previous work by investigating in detail the impact of system design parameters and extending the system to support multi-level cache hierarchies. Results show that the choice of implementation of multi-level cache hierarchies can have a significant impact on performance.

Keywords Many-core architectures · Cache coherence · Shared memory

1 Introduction

Chip-multiprocessors (CMPs) have now replaced very wide-issue out-of-order superscalar processors as they provide higher aggregate computational power, multiple clock domains, better power efficiency and simpler design through replicated building blocks. Current CMPs are commonly built around a relatively small number

C. Fensch (✉) · M. Cintra
School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, UK
e-mail: c.fensch@ed.ac.uk

M. Cintra
e-mail: mc@staffmail.ed.ac.uk

of cores (2 to 8), each with its own L1, and possibly L2, cache, connected through an on-chip interconnect that is either a shared bus or crossbar. Supporting shared-memory, parallel applications requires cache coherence, which is greatly facilitated by the use of buses and crossbars in current CMPs. Such interconnects allow for straightforward hardware cache coherence mechanisms based on snooping [31] and directories [19,22].

However, such types of interconnect are expected to become a bottleneck as the number of cores increases [25]. Either access latencies have to be significantly stretched or the area required by the interconnects has to be increased to the point of becoming impractical. Tiled CMPs [6,8,24,33,34] have been advocated as a possible alternative. Such systems are built from a relatively large number (≥ 32) of relatively simple cores plus a tightly integrated and lightweight point-to-point interconnect. Unfortunately, such scalable interconnects complicate the implementation of snooping and directory protocols. In fact, the existing hardware solution to cache coherence on such interconnects is to use fully distributed directory coherence protocols [27], which are notoriously hard to implement and verify (e.g., [1]).

We proposed an alternative, cost-effective software/hardware mechanism to support shared-memory parallel applications that forgoes hardware maintained cache coherence [13]. The mechanism is based on the key ideas that mapping of lines to physical caches is done at the page level with OS support and that the hardware efficiently supports remote cache accesses. An extension of the basic scheme only allows some *controlled* migration and replication of data. Data is migrated by refreshing the page mappings at barriers. Read-only sharing is done with the help of the existing write-protection mechanism in the TLB/OS. Overall, the mechanisms allow a sufficient degree of flexibility in the mapping and sharing.

In this article, we build on our previous work in [13]. We investigate how the page size affects the system performance of our scheme. Furthermore, we extend our system to support a multi-level cache hierarchy and consider several configuration options for such a hierarchy.

We evaluate our tiled CMP architecture on benchmarks from two very different domains—the Splash-2 scientific benchmarks and the ALPBench multimedia benchmarks. We compare the system against one with an SGI Origin like distributed directory cache coherence mechanism. Experimental results show that our scheme performs very close to this system with a performance gap as close as 0% (no gap), and 16% on average, across all benchmarks for 16 and 32 processors. Experiments with two possible configurations with L2 caches show that the best gap between the two configurations can be as much as 18%, and 6.4% on average for 32 processors.

The rest of the paper is organized as follows: Sect. 2 describes the tiled CMP architecture that we assume; Sect. 3 describes our scheme to support shared-memory parallel programs; Sect. 4 describes consideration when extending our scheme for a multi-level cache hierarchy; Sect. 5 describes our simulation infrastructure and our evaluation methodology; Sect. 6 presents the experimental results; Sect. 7 discusses related work; and Sect. 8 concludes the article.

2 Tiled Architectures

2.1 Current CMPs and Coherence Mechanisms

Current chip-multiprocessors are currently commonly built around a relatively small number of cores (2 to 8), each with its own L1, and possibly L2, cache, and are connected through an on-chip interconnect to a lower level shared cache. So far, the choice of on-chip interconnect has followed those of multi-chip symmetric multiprocessor (SMP) systems: shared bus fabrics and crossbars. The main reason for this choice is that such interconnects allow a straightforward implementation of coherence via snooping (bus) or directory at the shared cache level (crossbar). Unfortunately, as pointed out in [25], future technology scaling will lead to on-chip interconnects having different sets of tradeoffs and design issues than traditional off-chip interconnects. In particular, wire widths and the area required by connectors do not scale down at the same rate as other features shrink, which means that either the delay or the area overheads, or both, of buses and crossbars increase as process scales. In fact, the detailed study in [25] clearly shows that the area and delay overheads of buses and crossbars will become prohibitively high in CMPs with more than 16 cores in 65 nm and smaller processes.

In order to scale the number of cores in a CMP above this barrier, and into the numbers of cores proposed for tiled architectures [6, 8, 24, 33, 34], it is necessary to resort to scalable (i.e., point-to-point) interconnect types. Such interconnects are suitable not only because their peak bandwidth naturally scales with the number of cores, but also because, due to the short-length wires and low radix, their area overhead is a fixed, independent fraction of the number of cores. However, they do not lend themselves well to the implementation of snooping cache coherence protocols (although recent research attempts to address this limitation [30]). The alternative to continue enforcing cache coherence in such systems is to employ distributed directory schemes, which have been used in multi-chip multiprocessors in the past (e.g., [1, 27]). These have proven fairly scalable, reaching up to hundreds of processors. Snooping protocols are already somewhat difficult to completely debug and verify due to subtle corner cases and state transitions [15],¹ and distributed directories, with even more states, races, and corner cases, are notoriously even harder to debug and verify (e.g., [1]). Most of this complexity stems from the fact that requests cannot always be resolved at the home directory, but must in some cases generate further requests, such as forwarding and invalidation requests, which lead to complex protocols with subtle race conditions and several pending states. All this complexity is of serious concern and designing and verifying the directory coherence protocol for each new generation of the CMP architecture may likely become an expensive bottleneck.

An alternative to enforce coherence in a distributed memory system is to use the OS' virtual memory (VM) system to handle the copies of virtual pages, as was done on software DSM systems (e.g., [7, 20, 23, 28]). In this scheme, all caches are private and it is the responsibility of software to maintain coherence. As with distributed

¹ Further suggestion to the difficulty of complete verification is the recent Core 2 Duo Errata AI39: "Cache Data Access Request from One Core Hitting a Modified Line in the L1 Data Cache of the Other Core May Cause Unpredictable System Behavior" [18].

directories, such schemes have only been tested in multi-chip systems and must be adapted to operate on a CMP. A major drawback of directly porting software DSM schemes to the CMP environment is that such schemes require moving, comparing (“diff”), and copying data in physical memory pages to enforce coherence. This is because creating multiple physical copies of the same virtual page is the only way to cope with false sharing and the inability of the hardware to identify which parts of a cache line have been modified. In this way, at communication points, such as lock transfers and barriers, the individual copies must be compared against the previous stable copy of the page and the modifications must be merged into a single new stable copy of the page. These operations are likely to be extremely costly in a CMP, will consume precious off-chip memory bandwidth, and generate much pollution in the relatively small on-chip caches.

Overall, the potentially complex hardware solution of distributed directories and the potentially high-overhead software-only solution of a VM-based scheme are two extremes in the spectrum of solutions for the cache coherence problem in tiled CMPs. In Sect. 3 we describe our alternative to such cache coherence mechanisms, after we define the baseline tiled architecture in the next section.

2.2 A Baseline Architecture

In this paper, we are concerned with tiled CMPs consisting of 32 or more processors. Such systems are built by replicating regular building blocks, which are usually simple and small enough that the maximum intra-tile wire delay is small (1 to 2 cycles). As discussed in the previous section, snooping cache coherence approaches are unlikely to be suitable at such a scale due to the area overheads of the broadcast interconnects they require, and the only currently established alternative, namely distributed directory coherence, could prove to be prohibitively complex.

Before we describe our design, we first present the baseline tiled CMP that we assume. We assume a fairly generic tile that consists of a compute processor (PE) that is a simple single-issue RISC processor with separate and private instruction and data caches. These first level caches are virtually indexed and physically tagged.

The on-chip interconnect fabric consists of a point-to-point network with a mesh topology where each tile is connected to its four neighbors. Each tile contains a very simple network controller (NC) that performs simple dimension-ordered routing. The number of message buffers in the NC is enough to guarantee maximum throughput, which corresponds to four non-conflicting transfers per cycle. Figure 1 gives a high-level overview of the architecture (the shaded gray components are explained in Sect. 3).

3 A Hardware/OS Scheme to Avoid Cache Incoherence

As described in Sect. 2.2, the baseline architecture does not support shared-memory parallel applications because it suffers from the cache coherence problem. One option is to enforce cache coherence in hardware with a distributed directory protocol. For this purpose, one would add directory tags and a directory controller to each node next

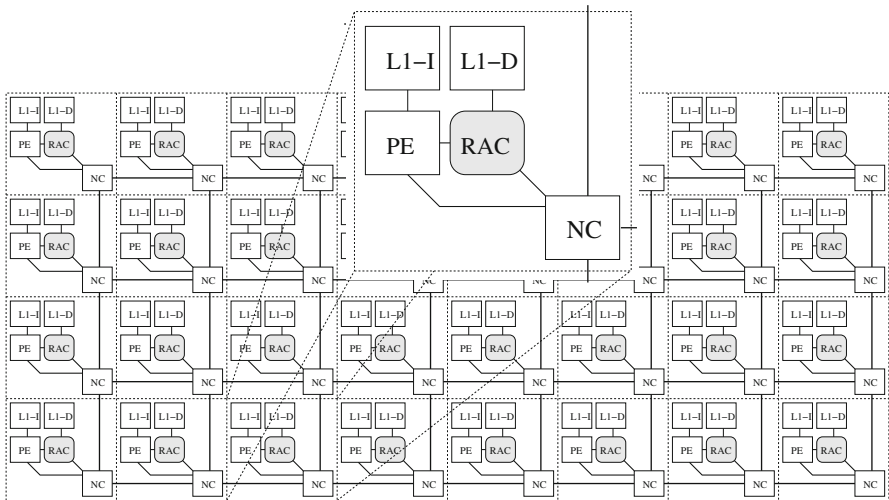


Fig. 1 Proposed tiled CMP overview

to the L1 data cache in Fig. 1. In this section, we present an alternative solution. The scheme divides the work between the hardware and the OS and, in reality, does not enforce coherence across cached copies of data, but rather avoids the possibility of incoherence by not allowing multiple modifiable shared copies of data. The key ideas are to map data to tiles at the granularity of pages under OS control and to support remote cache accesses in hardware.

3.1 Caches and Coherence

This section describes the mechanism we designed to avoid incoherence. The basic idea is to treat all L1s as a single logical cache and, thus, avoid replication of data, which can lead to data incoherence. This initial architecture is extended later in Sects. 3.2 and 3.3 to allow some *controlled* migration and replication.

3.1.1 Data Placement and Remote Cache Accesses

Instead of trying to keep the L1 caches coherent, our scheme avoids duplicate copies of a single cache line. To achieve this, every memory line can only reside in one L1 cache (the home cache or tile) and processors in other tiles must perform remote cache reads and writes to access the data. Thus, instead of a directory controller we add a remote cache access controller (RAC) to each tile, as shown in Fig. 1. To receive and service remote data requests the RAC is given access to the network and it uses the dedicated ports to the cache’s data and tag arrays that would be otherwise used by the snooping or directory controller.

The simplest way to place and locate data in the L1 caches whilst enforcing a single copy of each line would be to statically map lines to L1 caches based on the address.

This, however, is too restrictive and takes no account of the data access patterns. At the other end of the spectrum, each line would be dynamically mapped to any one L1 cache and it would be located through broadcasts, centralized tag stores, or redundant tag stores, as has been previously proposed [4, 11, 21, 40]. What we propose is to map whole memory pages to L1 caches through extensions to the OS page table and the hardware TLB mechanisms.

We expose the internal chip structure to the OS and extend the traditional page table with a new table that maps virtual pages to architectural tiles. This is matched with a new TLB-like hardware table that caches these translations and allows for fast identification of the home L1 cache where data in the page can be found. Each tile is given one of such hardware structures, which we call a MAP. The default policy for the OS to map virtual pages to tiles is first-touch. Note that our mechanism is different from simply mapping memory pages to L1 caches based on the physical address and using the virtual-to-physical page translation mechanism to provide the run-time mapping. The problem with the latter is that physical addresses are bound to specific L1 caches, which limits the OS flexibility in allocating physical memory and may lead to fragmentation and inefficient use of physical memory. Additionally, it makes any changes to the mappings much more involved, as the physical pages have to be moved in memory. It is for these reasons that we decided to add this extra level of indirection.

One important design decision at this point is where to provide virtual-to-physical address translation. Traditional CMPs keep all the translations of the local processor in the local TLB and ship only physical addresses to access lower level caches. A problem with using physical addresses for the remote cache accesses in our architecture appears when virtually indexed L1 caches are used, which is often the case in order to speed up accesses from the local processor. Thus, performing the virtual-to-physical address translations locally in the case of remote L1 accesses would require some (impractical) inverse translation at the remote tile. Our solution is to keep the virtual-to-physical address translations only in the TLB next to the home L1 cache and to ship virtual addresses over the network for remote cache accesses. Note that this is not intrinsic to our scheme, but a solution in case one wants to use virtually indexed L1 caches; with physically indexed L1 caches, our scheme would work as usual with physical addresses on the network.

In this scheme a processor request proceeds as follows (Fig. 2). Firstly, the virtual address is simultaneously used to index the local L1 cache, to perform a local TLB lookup to obtain the physical address, and to perform a local MAP lookup to obtain the identity of the home tile. If the result of the MAP translation points to a remote L1 cache, the local cache access is aborted. In this case, the result of the local TLB lookup is also ignored, including a possible TLB miss. The virtual address is then shipped to the RAC in the remote tile over the network. At the remote tile, the virtual address is simultaneously used to index the L1 cache and to perform a TLB lookup. To avoid delaying local cache requests due to remote cache requests we provide an extra port to the TLBs. Since our L1 caches are virtually indexed, the cache lookup can proceed in parallel with the TLB lookup and the extra TLB latency due to the extra port is unlikely to have any impact on the overall L1 access latency. If the TLB lookup succeeds then a tag comparison follows, using the physical tag. A cache or

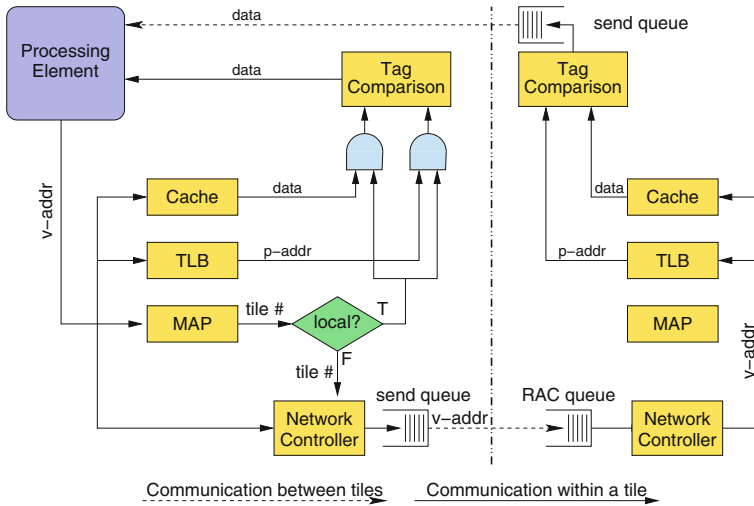


Fig. 2 Remote cache access mechanism

TLB miss is handled as usual. If the result of the MAP translation points to the local L1 cache then the local cache access proceeds as usual.

The above discussion only applies to data caches. Each tile has its own read-only instruction cache.

3.2 Data Migration

The first-touch data mapping strategy, combined with the fact that mapping is done at the granularity of pages, may lead to poor performance when data migrates across threads. Mechanisms have been proposed to allow migration and replication of memory pages in CC-NUMA machines [36]. These are tailored to much larger systems with larger latencies, and, thus, we borrow some of their ideas but adapt the mechanisms and policies.

We propose a simple mechanism that allows for some degree of migration by invalidating the mappings of virtual memory pages to L1 caches. This is done by invalidating the MAP table in all tiles. After an invalidation, a first-touch policy is again used for the new mappings. Thus, invalidating the mappings does not in itself migrate pages, but it creates an opportunity for this to happen. The invalidation is more easily implemented at a quiescent state where there are no pending memory requests on chip. A natural point to perform such invalidation is at barriers. In many well-designed applications barriers are used to signal change in the data access pattern and communication across threads. Thus, barriers are also naturally good points for re-mapping and migration. Finally, to effect the migration of the data, all dirty lines in the L1 caches must be written back at a mapping invalidation such that the modified data may be reachable after the re-mapping.

The actual invalidation is done in two phases. First, each processor invalidates the local MAP table just before joining the barrier. This is done with a new instruction that is very similar to the existing `tlbia` instruction in the PowerPC IS. At this point, the local cache controller starts writing back dirty cache lines to main memory with the goal of hiding the write-back overhead with the idle synchronization time.

In the second phase, before releasing the barrier, one processor invokes a special system call to invalidate the OS' internal MAP table. Also at this point (before barrier release) all tiles write back all remaining dirty cache lines. When the writebacks are completed, the contents of the caches are invalidated and the barrier is released.

3.3 Read-Only Data Sharing

The baseline scheme, coupled with the extension to refresh mappings to allow migration, is likely to work well as long as there is not much sharing of data at the granularity of pages. Whilst full-blown sharing requires line-based hardware coherence, some degree of sharing can be easily enforced by the OS with minimal hardware support. What we propose is a simple mechanism that allows sharing of pages across multiple readers and a single writer at any given time. The mechanism works as follows. The first processor to touch a given page for reading obtains a local mapping for it whilst the OS marks the page as read-only in the page table and in the processor's TLB and MAP. Other processors touching the same page for reading are allowed to create local mappings for it, also in read-only mode. At this point, the OS *does not* need to keep track of which processors are sharing the page. The first write by a processor to a page is intercepted by the OS, which then marks the page as modified and makes this processor the owner of the page. Subsequent reads by other processors with existing local mappings can continue to use these mappings and access local data. However, subsequent writes by other processors when intercepted by the OS will not be allowed to proceed locally, but will generate a MAP entry (or change it if one already exists) that points to the owner node. Similarly, reads by processors without a local mapping for the page will generate an entry pointing to the owner node. Figure 3 shows a state diagram for the complete protocol. Note that most of the state transitions occur only at the OS level and the hardware state machine (corresponding to the MAP state in the figure) is fairly simple.

The mechanism just described allows processors to continue using local mappings and locally cached data even after other processors write to data in the page. To prevent stale data from being used we assume a release consistency memory model and invalidate the MAP entries for shared pages on lock acquire operations. This is done by a special instruction that clears the valid bit of a MAP entry if the shared bit is set. By doing so we guarantee that all accesses to data modified by other processors will use a new remote mapping and will become remote. Entries that point to non-shared data don't have to be invalidated, since no migration happens at lock acquires and, thus, they do not change. It is also necessary to extend the barrier actions used with the migration mechanism of Sect. 3.2 to include a full cache flush in addition to the writebacks and the refresh of the mapping. No special action is required on lock releases.

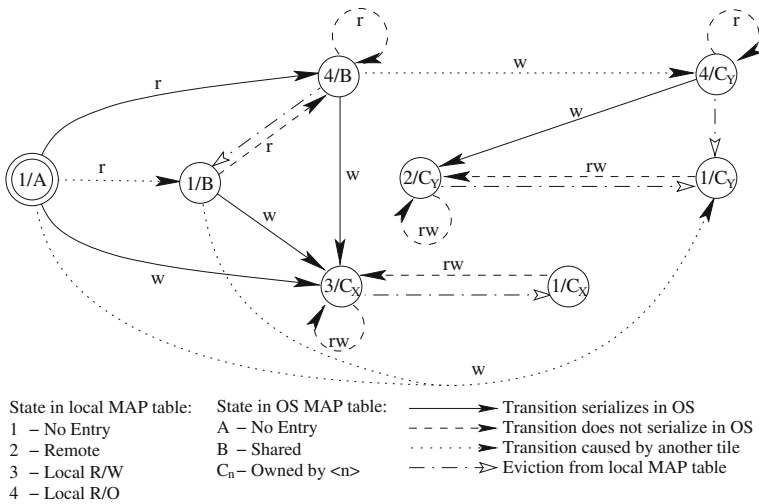


Fig. 3 Sharing protocol for node X. The state of a page depends on its local (*numbers*) and OS (*letters*) state

While this mechanism may seem very similar to previous software cache coherence mechanisms (e.g., [23]), it differs from these in one crucial way. Namely, it does not allow multiple writers, reverts to a single up-to-date copy of every page upon a write, and enforces remote cache accesses in such cases. The key benefit of this is that in our scheme, no multiple modified copies of physical pages exist at any time and, thus, there is no need to perform expensive diff operations and copy data in memory.

3.4 Synchronization

Memory locks have been implemented in the past using either *compare&swap*-style atomic instructions or *load-link store-conditional* (LL-SC) pairs. The latter approach has been favored recently because it is easier to implement in hardware with cache coherence. Our architecture can implement both approaches by using the RCA on the remote tile as a synchronization point. *Compare&swap*-style primitives can be more easily implemented than in current multiprocessors. This is because there is no replication of the lock variable in multiple caches and it is, thus, easier to enforce the atomicity of the primitive. Implementing this primitive then only requires adding the compare logic to the cache controller and blocking subsequent requests from other processors until the swap is performed. On the other hand, implementing *load-link store-conditional* pairs in our architecture is more difficult than in current multiprocessors with cache coherence [13].

3.5 Cost Comparison with Directory Coherence

Since we are proposing to replace a directory controller and its protocol with our RAC+MAP and a combined hardware/OS protocol, it is relevant to compare both

schemes' area and complexity overheads. In particular, our main goal is to provide a less complex alternative. A comprehensive comparison between the two competing approaches would require the full design of the controllers and their circuit implementation. This is a highly involved task and, instead, we attempt to provide some intuition into why we believe our scheme is less complex.

Like a directory controller, the RAC has to handle remote read and write requests. Unlike a directory controller, it does not have to deal with forwarded transactions and multiple invalidations, which lead to complex protocols with subtle race conditions and several pending states. The RAC can directly handle requests and generate responses for all transactions in our protocol. Thus, the RAC has fewer states and a much simpler finite state machine, which means that it has simpler logic than a directory controller does. This means that the resulting protocol is simpler to verify and validate.

As far as state storage is concerned, there is probably no significant difference. For instance, for a 32 tile system a MAP table with 128 entries, each with 22 bits (15 bits for the virtual address tag, 5 bits for the tile ID, 1 shared bit, and 1 valid bit) would have a total of 352 bytes. A directory for 32 Kbytes L1 caches and 32 bytes lines would have 34 bits per entry (32 bits for the sharing vector and 2 bits for line state), for a total of about 4 Kbytes.

On the negative side, our system requires an additional port to the 4-way associative TLB to handle remote accesses independently from the CPU. As we mentioned earlier this is unlikely to impact the overall L1 access latency with our virtually indexed caches.

3.6 Software Implications

Our scheme has implications for the operating system and for the applications. The operating system requires changes that make it aware of the presence of the MAP table. Such changes include system calls to manipulate the MAP table, traps that deal with MAP table misses and saving the state of the MAP table in case of a context switch. These operations are all very similar to the already existing support for TLBs and as such can easily be implemented.

On the application side, several modifications may be required. The number of modifications depends on if the data migration or read-only sharing extensions are being used. If neither extension is being used then no modifications to application are necessary. If the migration extension is used, then it becomes necessary to flush all caches and reset the MAP table (as described in Sect. 3.2). While such an activities could be performed by the application, it is more reasonable to assume that they are encapsulated in a system library. If no extension or only the migration extension is used, our system provides sequential consistency.

Using the read-only sharing extension has two implications: first, the actions described in Sect. 3.3 have to be performed whenever a lock is acquired or released. This can be ensured by implementing these actions within a system library. Second, the system now only provides release consistency, potentially affecting every memory access in the application.

Relaxed memory consistency models were originally introduced mainly because of better performance with respect to sequential consistency and because they are arguably easier to implement in systems with complex memory hierarchies [2]. Later it was noted that aggressive dynamic instruction scheduling in processors could close a large fraction of this performance gap while still providing a tight memory model to the programmer, which led to the suggestion that this should be the case [16]. This approach, however, has only been implemented in one processor, the MIPS R10K. The suggestion was based on the assumption that sequential consistency is a better memory consistency model from a software engineering perspective. However, many believe that most programmers have neither the skills nor the tools to safely use handcrafted communication and synchronization that relies on the strict ordering of memory operations imposed by sequential consistency, thus negating its main purported benefit.

While the debate is still open on what memory consistency models and the architecture should provide and what models programmers should see, in the end what matters is the models used in de-facto standards for parallel programming environments. For scientific HPC the leader is OpenMP, and for embedded, desktop and server applications the dominant models are Java, and POSIX threads. The consistency model of OpenMP is only slightly less relaxed than release consistency in that it does not allow loads to cross below a release or stores to cross below an acquire. As such, our scheme can be easily extended to work with it by the addition of memory fences to acquires and releases. The memory model for POSIX thread libraries is deliberately informal, but essentially conforms to release consistency in that all communication must be protected by calls to synchronization routines. Thus, our scheme works “out of the box” with POSIX threads, assuming a POSIX thread library that is adapted to our system. In fact, no changes were required to port Splash-2 to our system when using a PARMACS macro package for POSIX threads. While the Java memory model is much more involved than those for C, in essence it also requires that programs use explicit synchronization. Finally, a trend in recent programming languages is to use “atomic” blocks (e.g., [3, 12]), which also fit straightforwardly with our proposed scheme.

In summary, almost all software complexity is hidden from the programmer. Our only requirement is that application conforms to release consistency, which as discussed is not a strong limitation. All other complexity is hidden from the programmer, assuming the application uses system provided libraries to implement locks, barriers and conditional variables.

In addition, our system offers the programmer to control data placement manually by preloading a page-to-tile mapping table using a system call. However, an evaluation of this functionality is beyond the scope of this article.

4 Multi-Level On-Chip Cache Hierarchies

The design described so far assumes only a single level of cache per tile and no other level of cache on chip. In some cases, a higher storage capacity per tile may be required. Our architecture can be extended to work with L2 caches in each tile and our key ideas can still be applied. In this case, the L2s take the roles of the L1s in the architecture




Fig. 4 Possible options to design a tile with multiple levels of cache

described so far and constitute a single logical shared cache. Again, mapping of memory lines to L2 caches is done at the granularity of pages with both OS and hardware support. The migration and shared-only replication mechanisms can still be applied. The only requirement is that the L1 caches must only be allowed to cache lines that are mapped to the local L2 cache.

However, the introduction of an L2 cache offers several options as to which cache the RAC should be connected. Two design options are shown in Fig. 4. The probably most intuitive version is shown in Fig. 4a, where the second level cache is simply attached as a backing store to first level cache. The coherence between L1 and L2 can be easily arranged by a write-through L1.

However, other arrangements are possible. By connecting the remote cache access controller (RAC) directly to the second level cache, one could hope that the first level cache gets less polluted by remotely requested data and, thus, offering a higher hit rate for local CPU requests. On the downside, serving remote accesses would now take longer, since they can only be served at the hit latency of the L2. Furthermore, the design of the L1 becomes more complicated, since it also has to respond to invalidation messages from the L2, in case a cache line has been modified remotely. On the other side, the L1 has no longer requires additional ports to service the requests from the RAC. Such an arrangement is shown in Fig. 4b.

5 Evaluation Setup

5.1 Applications

For our performance analysis, we use the Splash-2 benchmarks [37] and three ALP-Bench benchmarks [29]. The Splash-2 benchmarks are representative of scientific and engineering workloads and the ALPBench benchmarks are representative of multimedia workloads. Both benchmark suites use explicit locks and barriers, assume the release consistency memory model, and rely on hardware maintained cache coherence if caches are used. We use the reference inputs for the Splash-2 benchmarks except *radiosity* for which we use a reduced input set to keep simulation time manageable. Similarly, we reduced the input for *mpegdec* to only 20 frames. Because the input

Table 1 Characteristics of the applications used

Benchmark	Input size	Instructions	Locks	Barriers
Splash-2 Kernel				
cholesky	tk29.O	1,234 M	72,075	3
fft	65,536 points	58 M	32	7
lu	512 × 512 matrix, 16 × 16 block	389 M	32	67
radix	262,144 keys	54 M	406	12
Splash-2 Application				
barnes	16,384 particles	4,361 M	69,360	18
fmm	16,384 particles	2,903 M	47,074	34
ocean	258 × 258 grid	412 M	6,656	900
radiosity	demo	646 M	281,217	19
raytrace	car	2,006 M	95,528	2
volrend	head	1,344 M	38,604	20
water-nsq	512 molecules	652 M	35,360	19
water-spa	512 molecules	664 M	609	19
ALP Bench				
facerec	ALP Training	2,826 M	30	3
mpegdec	525_tens_040.m2v	1,049 M	29	41
mpegenc	Output of mpegdec	9,477 M	29	40

The number of instructions refers to the total number for a sequential execution of the benchmark. The number of locks refers to those encountered by all 32/16 tiles (Splash-2/ALPBench) within the application (not library) code

sets for the ALPBench benchmarks were not intended to be used with more than 16 processors, we do not simulate larger systems for these benchmarks.

Speedups are reported with respect to the execution time of the sequential programs on a single processor after initialization. Table 1 lists the benchmarks we used.

The benchmarks were compiled with gcc 3.4.4 and glibc 2.3.5 for PowerPC. Compiler and library were modified such that they use synchronization primitives that have been adapted to our architecture.

5.2 Simulation Environment

We used the Liberty Simulation Environment (LSE) [35] to implement our simulator. A tile consists of a PowerPC core, a network controller, a data cache module, and a private instruction cache. The single-issue CPU is implemented as an 8-stage pipeline running at 2 GHz and is simulated in detail. The cache has been implemented with the cache module from SimpleScalar [5]. The details for the memory system are shown in Table 2. We also implemented a detailed wormhole routed interconnect, where contention is accurately simulated at the network end points as well intermediate nodes. System calls and interrupts to the OS are assumed to take 2000 cycles.

Table 2 Memory system configuration

Remote cache access latency without any congestion: $2 * (h + w) + t + 1$, where h is the number of hops, w is the number of words in the message (2 or 3), and t is the access time at the remote cache	L1 D-cache configuration	
	size	32 K
	hit latency	3 cycles
	miss latency	200 + 16 cycles
	block size	32 bytes
	associativity	4-way
	TLB/MAP configuration	
	entries	128
	page size	4 K
	associativity	4-way
	hit latency	1 cycle
	RAC input queue configuration	
	entries	32

5.3 Systems Evaluated

We compare our architecture against a similarly configured one where the L1 caches are kept coherent on a cache line basis through an SGI-Origin-like distributed directory protocol [27]. We note that developing this protocol was greatly simplified by the use of common simulation artifices and the complexity we found is very different from the complexity we expect from a real implementation. For fairness of comparison, we augment the directory scheme with migration of pages at barriers, which minimizes any negative effects from the first-touch home-allocation policy in the initialization phase. The cost of migration is the same as in our system: 2000 cycles plus the cost of flushing the caches. For the directory controller, we assume an aggressive hardware implementation that requires only 5 cycles to process each request. We also compare our architecture against one that maintains cache coherence through a TreadMarks-like software DSM protocol [20]. Our implementation is much simplified in that it only takes into consideration the overhead of creating diffs and cache pollution by twin and diff creation. To estimate the cost of the diff we wrote a highly optimized kernel that compares the contents of two physical pages in memory and writes back one of the values if they differ. The cost was measured to be about 50 K cycles.

We refer to the systems as *NUCA-Dist* for a system that implements our architecture with both re-mapping and read-only sharing of pages, *Dir-Coh* for the system with directory coherence, and *SW DSM* for the system with software DSM coherence.

6 Experimental Results

6.1 Overall Performance

We start by comparing the overall performance of our architecture against the hardware distributed directory system. Figure 5 presents the speedups for 32 (Splash-2)

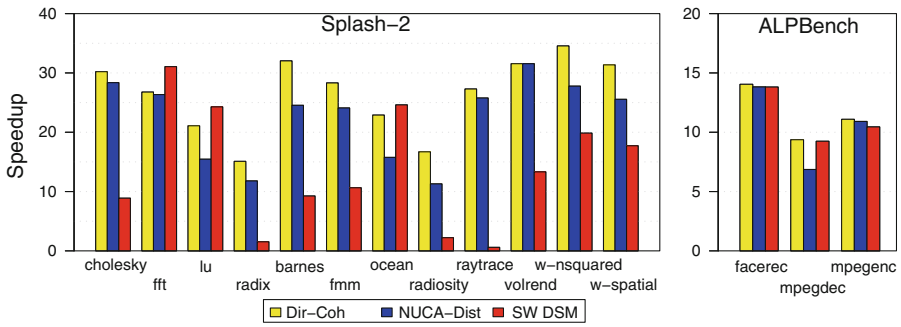


Fig. 5 Speedups for 32 (Splash-2) and 16 (ALPBench) tiles compared to the execution time of a single tile

and 16 (ALPBench) tile systems of *Dir-Coh*, *NUCA-Dist*, and *SW DSM*. We can see that *Dir-Coh* scales well for most benchmarks, with an efficiency (speedup divided by number of processors) of 81% on average. These results are somewhat better than those in [27] mainly due to the lower communication latencies observed in a single chip multiprocessor.

Looking at the performance of our scheme (*NUCA-Dist*) we can see that it performs fairly close to the hardware directory coherence system, with a performance gap for 32 processors ranging from 0% (no gap) to 32% (for *radiosity*), and 16% on average. Moreover, the performance gap is less than 10% for 6 out of 15 benchmarks, which is an impressive result considering that the directory coherence system uses a very aggressive hardware implementation and that our architecture requires only simple hardware support.

Finally, *SW DSM* performs, with few exceptions, significantly worse than the other systems. While the system performs very well on benchmarks that mainly use barriers for synchronization (the good results are possibly due to our simplifications), the results show that it is not able to provide sufficient scalability for most applications. The gap to our system is on average 27%, ranging from -57% to 98%. These results are in line with those reported in [17].

We furthermore investigated our scheme with a varying numbers of tiles. By doing so, we hope to get some insight into the overall scalability of our scheme. Figure 6 shows the speedups observed for 2 to 32 tiles. In general, we notice that for most benchmark the observed speedup doubles if the number of tiles is doubled. A notable exception from this trend is when the number of tiles is increased from 8 to 16. For several benchmarks (*fft*, *lu*, *barnes*, *fmm*, *raytrace*, *volrend*, *water-nsq* and *water-spa*) the observed speedup more than doubles. The reason is that by doubling the number of tiles, we also double the number the amount of on-chip storage. At this point, most of the working set can be stored on-chip reducing off-chip memory accesses and improving performance. Looking at the scalability of *radix* and *radiosity* when increasing the number of tiles from 16 to 32 reveals a drop in scalability. This behavior is in line with the results previously reported [37].




Fig. 6 Speedups obtained with *NUCA-Dist* for the Splash-2 benchmarks when compared to the execution time of a single tile

6.2 Memory Access Breakdown

To better understand the behavior of our architecture, we track the outcome of each processor memory request. Figure 7 shows the breakdown of memory requests for each benchmark and for configurations with 32 (Splash-2) and 16 (ALPBench) processors. For each benchmark and configuration, the bar is normalized to the total number of processor memory requests, which does not vary noticeably across the different systems. The bars are broken down into the following components: accesses that hit in the local L1 cache (*local hits*); accesses that hit in a remote L1 cache (*remote hits*); accesses that go off-chip following a miss in the local cache (*local miss*); and accesses that go off-chip following a miss in a remote cache (*remote miss*).

The figure shows that the fraction of off-chip accesses is fairly small in most cases, with the exception being *ocean*, where the off-chip accesses for all systems account for about 12% of all requests. Another exception is *facerec*, where sequential execution

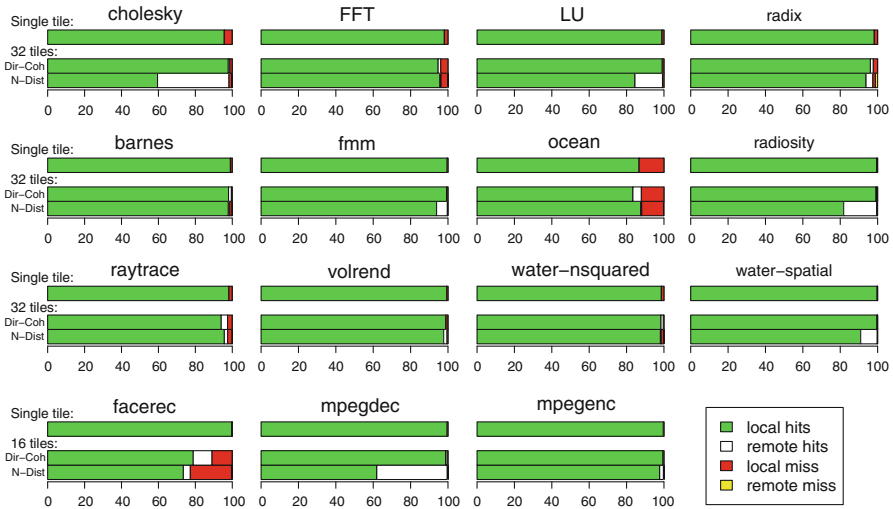


Fig. 7 Distribution of memory accesses into local and remote, further divided into cache hits and misses




Fig. 8 Average latencies for local, remote, and all loads for *NUCA-Dist*. The average latencies for *Dir-Coh* are shown as a comparison

shows only a small number of off-chip accesses, but both parallel systems show a large fraction of off-chip accesses. Given this generally small number of off-chip accesses, we expect the main differentiating factor to be the ratio of local to remote cache accesses.

The results for *NUCA-Dist* show that the fraction of remote cache accesses is fairly small for most benchmarks, except *cholesky* and *mpegdec*, and, to a lesser extent, *lu* and *radiosity*. Such a relatively small number of remote cache accesses partially explains the good performance of our architecture for many benchmarks. An interesting case is *cholesky* where the fraction of remote cache accesses is high compared to most benchmarks, but its performance with *NUCA-Dist* is good. On the other hand, some benchmarks, such as *ocean* and *barnes*, show a small fraction of remote cache accesses, but their performance with *NUCA-Dist* is not as good as some of the other benchmarks. The results for *Dir-Coh*, on the other hand, show that it incurs very few remote accesses (i.e., cache-to-cache transfers), which mainly explains its very good performance.

The impact of local, remote, and off-chip accesses can be further seen in Fig. 8, which shows the average load latencies, in cycles, for the different types of loads for *NUCA-Dist* and for the average load for *Dir-Coh*. While the latencies for remote loads in *NUCA-Dist* are significantly larger than those of local loads, the average latencies are fairly close to the local ones and, thus, very close to those in *Dir-Coh*.

Migration and replication not only improve the average load latency by converting remote accesses to local ones, but also reduce the average load latency of the remote loads themselves (results not shown). This is because reducing remote accesses reduces the contention that occurs when multiple requests target the same tile.

6.3 Impact of the Granularity of the MAP Table

To try to further reduce the amount of remote accesses in *NUCA-Dist* we experimented with 1 KByte pages. Table 3 shows the results of this analysis. As expected, the number of misses in the MAP table increases and, on average the miss rate is now 3.6 times higher than in a system with 4 K MAP pages. However, the impact on the overall performance of the system is less obvious, some benchmarks (like *lu* or

Table 3 Difference in speedup and MAP table miss rate between a system that uses 4K and 1K MAP table pages

	4 k Pages		1 k Pages		Change in %	
	speedup	miss rate (%)	speedup	miss rate (%)	speedup	miss rate
cholesky	28.36	0.11	26.65	0.20	-6.42	82
fft	26.35	0.35	26.35	0.42	0.00	20
lu	15.45	0.01	17.78	0.04	13.10	300
radix	11.79	1.21	11.85	1.84	0.51	52
barnes	24.54	0.08	23.81	0.32	-3.07	300
fmm	24.09	0.10	21.08	0.21	-14.28	110
ocean	15.76	0.36	9.03	0.79	-74.53	119
radiosity	11.29	0.31	12.16	1.09	7.15	251
raytrace	25.77	0.90	24.26	1.36	-6.22	51
volrend	31.56	0.24	30.35	0.42	-3.99	75
water-nsq	27.77	0.03	25.89	0.28	-7.26	833
water-spa	25.55	0.01	25.65	0.03	0.39	200
facerec	13.82	0.12	13.81	0.73	-0.07	508
mpegdec	6.86	0.01	8.93	0.02	23.18	100
mpegenc	10.91	0.001	9.58	0.01	-13.88	900
Average					-5.69	260

As before, Splash-2 benchmarks were executed on a 32 tile system, while ALPBench benchmarks were executed on a 16 tile system




Fig. 9 Distribution of memory accesses into local and remote accesses for *NUCA-Dist+MS* with a MAP table with either 4k or 1k granularity. The accesses are then further divided into cache hits and misses

mpegdec) greatly benefit from smaller MAP pages, while other benchmarks do not benefit at all and actually lose performance (like *fmm*, *ocean* or *mpegenc*). This behavior can be explained by looking at the cache access distribution of the system with the finer granularity (shown in Fig. 9). For benchmarks that profit from the smaller granularity the number of remote accesses is reduced (in the case *lu* they are almost




Fig. 10 Number of messages in the system during the parallel part of the execution on 32 tiles. The *x*-axis represents time, while the *y*-axis shows the number of messages

completely eliminated). For benchmarks that do not profit, the local to remote ratio almost remains unchanged. These benchmarks just have to pay the penalty of missing more often in the local MAP table. In total only 4 benchmarks benefit from the smaller MAP table pages; all other benchmarks (with the exception of *fft* whose performance is not affected) lose performance. The average loss is 5.7%. One interesting thing to notice is that there is no correlation between the increase of the MAP miss rate and the impact on the performance. For *fft*, *ocean* and *mpegdec* the MAP miss rate approximately doubles, however the impact on performance varies between a critical slowdown to a significant speedup.

6.4 Network and Contention Effects

One important effect of our mechanism is a potential increase in the number of messages in the network, due to the remote accesses used in the scheme. To properly account for this effect, we modeled the network in detail including congestion both at intermediate nodes and at the end points. Congestion at the end points does lead to some performance degradation and is one of the main reasons for the relatively large remote cache access latency shown in Fig. 8 (note that for a 32 tile system the uncontended remote cache access latency should be around 18 processor cycles). Figure 10 shows the total number of messages across time for the three benchmarks with the largest fraction of remote accesses (Fig. 7) (the other benchmarks have much fewer messages and we do not show their results for lack of space). As can be seen from the figure, the number of messages in the system is usually low and often less than the number of tiles in the system (except for the latter part of the execution of *radiosity*, which has several peaks close to 60 messages). This relatively small number of messages in the system is the reason for the small contention inside the network, and hot spot effects explain the noticeable contention at the end nodes.

6.5 Impact of Flushing and Invalidations

Our read-only sharing scheme (Sect. 3.3) involves the potentially very expensive operations of flushing caches on barriers and invalidating the MAP table on lock acquires. To assess the actual impact of these operations' overheads on our benchmarks, we run a modified version of our scheme that does not suffer from these overheads.

Table 4 Overhead for *NUCA-Dist* with 32/16 processors in % caused by flushing the cache at barriers and invalidating the MAP table on a lock acquire.

	Barriers	Locks		Barriers	Locks
cholesky	<0.01	0.77	raytrace	<0.01	5.57
fft	0.23	0.04	volrend	0.00	0.00
lu	8.01	<0.01	water-nsq	0.18	0.43
radix	8.75	0.17	water-spa	0.47	0.04
barnes	<0.01	0.61	facerec	<0.01	0.07
fmm	1.67	0.21	mpegdec	<0.01	<0.01
ocean	10.56	<0.01	mpegenc	<0.01	<0.01
radiosity	3.74	<0.01			

The results of this analysis for a 32 core system are shown in Table 4 for the Splash-2 benchmarks. For most benchmarks, the overhead stays well below 1%. Exceptions are *ocean*, which experiences close to 11% overhead at barriers, and *raytrace*, which experiences 6% overhead at locks. The overhead for *ocean* was expected considering that this benchmark has 900 barriers. Similarly, since *raytrace* has a high number of locks, it is not much of a surprise that it suffers from invalidating the MAP table. Still, other benchmarks have similar numbers of locks and do not suffer as much. In these benchmarks a significant number of pages are not mapped as shared on a lock acquire, and thus are not invalidated. While these overheads are a non-negligible cause of some performance loss in three benchmarks, they do not affect the others as badly as one might expect.

6.6 Multi-Level On-Chip Cache Hierarchies

The design evaluated so far assumes only a single level of cache per tile. We also evaluated systems with a 128 KByte L2 cache per tile and a write-through L1 cache with the same size as before. The total L2 capacity on chip of 4 MBytes and the relatively small capacity per core are in line with what could be expected from a CMP with 32 cores. Each L2 has 20 cycle access time.

Figure 11 shows the speedup results of such a system for *Dir-Coh* and *NUCA-Dist A & B* for the Splash-2 benchmarks. *NUCA-Dist A* is a system where the RCA is connected to the L1 and *NUCA-Dist B* is a system where the RCA is connected to the L2 (see Fig. 4). Note that these speedup numbers are not directly comparable to those in Fig. 5, because they are normalized to different sequential execution times. The figure shows that it makes a difference where the RCA is connected to the cache hierarchy.

The figure shows that the performance gap between *NUCA-Dist A* and *Dir-Coh* remains mostly the same as for systems without the second-level cache (the gap range is now 1%–32% and the average gap is 15%). However, for the system where the RCA is connected to the L2 cache (*NUCA-Dist B*), the performance gap increases (the gap range is now 1%–41% and the average gap is 20%). The gap between *NUCA-Dist A*




Fig. 11 Speedups for 32 tiles with L2 caches compared to the execution time of a single tile also with L2 cache




Fig. 12 Average latencies for local, remote, and all loads for *NUCA-Dist A* and *NUCA-Dist B*. The average latencies for *Dir-Coh* are shown as a comparison

and *NUCA-Dist B* can be as much as 18%, and 6.4% on average. This demonstrates it is crucial to connect the RCA to the right level of cache and that our scheme also works with the addition of a second level cache.

The difference in performance between these two multi-level cache configurations can be explained by looking at the load latency. Figure 12 shows that the remote latency increases, as expected since remote requests are now being serviced by the slower L2 cache. It also shows that for some benchmarks (*fft*, *barnes*, *ocean* and *water-nsq*) the increase in remote access latency has little or no impact on the average latency. In these cases, the performance is not affected by the higher latencies of remote accesses. The overall conclusion that can be drawn from this experiment is that the remote access latency should be as short as possible.

7 Related Work

The work in [4] extended the original uniprocessor NUCA proposal of [21] for CMPs. Unlike our work, that work focused on a large shared L2 and assumed that L1 coherence is maintained through directories.

Closer to our work, [9, 11, 40] considered the tradeoffs in organizing the L2 caches in a tiled CMP where L2 is physically distributed along with each tile. Similarly to

ours, those works considered the option of organizing these distributed L2 caches as a logically single L2 cache. They differ from ours in the following ways: firstly, the L1 caches are private to each tile and allow replication of data, such that coherence is always required; secondly, those works propose techniques that allow replication of data in the L2 caches that is at the line level and is controlled by the hardware. Our work emphasizes simplicity and only allows a very restricted degree of replication that is totally controlled by the OS and, thus, forgoes hardware coherence mechanisms.

Our work is also similar in spirit to attempts to migrate most of the cache coherence management to software [10,26]. Like those systems, our proposal benefits from the possibility to modify, and fix, the protocol with software modification and without any hardware changes. Those systems, however, run a full-blown coherence protocol in a dedicated protocol processor or a dedicated processor context. Even closer to ours are recent works that attempt to transfer some of the coherence burden to the OS/software [38,39]. Unlike such previous trap-based schemes, however, the small hardware extensions that we propose minimize the need for OS and trap handler activity. In our scheme, only the processor's first load or store to data in a page requires trap handler intervention and only the system's first load or store to data in a page requires full OS intervention. Another important difference is that all those schemes focused on coherence mechanisms for multi-chip systems.

There have been several proposals for tiled CMP architectures [6,8,24,33,34]. Most of these have focused on novel execution paradigms to exploit ILP and DLP in single-threaded applications. In the few studies with parallel applications, it is assumed that there is some hardware mechanism for cache coherence, but no details are given. Closer to our architecture, [8] does not provide hardware cache coherence, but, unlike ours, relies on the programmer/compiler to maintain coherence.

Our work is related to previous work on OS directed page migration and replication in CC-NUMA environments, such as [36]. Those differ from ours in that the hardware cache coherence mechanism of CC-NUMA machines supports fine-grain caching of memory lines, so that the page-level migration and replication is only necessary when the workloads overflow the private caches.

While most past shared-memory systems offered cache coherence in hardware, the Cray T3D and T3E are notable exceptions [32]. Unlike our system, those machines did not support remote cache accesses and did not offer OS control of caching. Thus, avoiding incorrect local caching of shared data was left to the responsibility of the programmer/compiler. Hardware supported remote memory accesses were also proposed in the M-Machine [14]. However, that system also allowed indiscriminate private caching of data and no details are given on how coherence would be maintained, and it leaves the decision of caching versus remote accesses to the programmer/compiler.

Finally, our work is also related to previous work on software DSM systems, such as [7,20,23,28]. Similarly to our proposal, those also tried to avoid the costs of hardware coherence by using the OS page mechanism to enforce coherence, but unlike ours, the majority of those systems supported full-blown coherence in software with full replication and multiple readers and writers. Our proposal, on the other hand, allows only a single writer at a time and relies on the relatively short communication delays on chip to perform efficient remote cache accesses. While [28] enforced a single-writer policy, it allowed ownership to move across nodes instead of enforcing remote

accesses, which can lead to significant traffic. In addition, while [23] supported remote writes, it did not support remote reads, which had to be implemented by a tortuous mechanism by which the remote node performs remote writes on request. Those works also differ from ours in that they were tailored to multi-computer systems, where no hardware-supported single address exists.

8 Conclusion

In this paper, we presented and evaluated a novel cost-effective software/hardware mechanism to support shared-memory parallel applications that forgoes hardware maintained cache coherence. Our mechanism treats all caches in the tiled CMP as a single logical cache and is based on the key idea that mapping of lines to physical caches is done at the page level with OS support. We extend a tiled CMP architecture with this mechanism and evaluate it on the Splash-2 and ALPBench benchmarks against an SGI-Origin-like cache coherent system. We extend our system by adding different multi-level cache hierarchies and investigate in more detail the impact of different page sizes to the system. The performance of our system is within 16% on average for 16 and 32 processors of the directory coherent system across all benchmarks. This is an impressive result considering that the directory coherence system uses a very aggressive hardware implementation and that our architecture requires only simple hardware support.

Acknowledgments This work was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) under grant GR/S79572/01 and by the European Commission (EC) under grants IP 27648 (FP6) and HiPEAC IST-004408.

References

1. Abts, D., Scott, S., Lilja, D.J.: So many states, so little time: verifying memory coherence in the Cray X1. In: Proceedings of the International Parallel and Distributed Processing Symposium. (2003). doi:[10.1109/IPDPS.2003.1213087](https://doi.org/10.1109/IPDPS.2003.1213087)
2. Adve, S.V., Gharachorloo, K.: Shared memory consistency models: a tutorial. *IEEE Comput.* **29**(12) (1996). doi:[10.1109/2.546611](https://doi.org/10.1109/2.546611)
3. Allen, E., Chase, D., Hallett, J., Luchangco, V., Maessen, J.W., Ryu, S., Steele, G.L. Jr., Tobin-Hochstadt, S.: The Fortress Language Specification Version 1.0 β . Sun Microsystems, Inc., <http://research.sun.com/projects/plrg/Publications/fortress1.0beta.pdf> (2007)
4. Beckmann, B.M., Wood, D.A.: Managing wire delay in large chip-multiprocessor caches. In: Proceedings of the 37th Annual IEEE/ACM International Symposium on Microarchitecture, pp. 319–330. (2004). doi:[10.1109/MICRO.2004.21](https://doi.org/10.1109/MICRO.2004.21)
5. Burger, D., Austin, T.M., Bennett, S.: Evaluating future microprocessors: the SimpleScalar tool set. Technical Report CS-TR-1996-1308, University of Wisconsin-Madison (1996)
6. Burger, D., Keckler, S.W., McKinley, K.S., Dahlin, M., John, L.K., Lin, C., Moore, C.R., Burrill, J., McDonald, R.G., Yoder, W.: The TRIPS Team: Scaling to the end of silicon with EDGE architectures. *IEEE Comput.* **37**(7), 44–55 (2004). doi:[10.1109/MC.2004.65](https://doi.org/10.1109/MC.2004.65)
7. Carter, J.B., Bennett, J.K., Zwaenepoel, W.: Implementation and performance of munin. In: Proceedings of the 13th Symposium on Operating Systems Principles, pp. 152–164 (1991). doi:[10.1145/121133.121159](https://doi.org/10.1145/121133.121159)
8. Caşcaval, C., Castañón, J.G., Ceze, L., Denneau, M., Gupta, M., Lieber, D., Moreira, J.E., Strauss, K., Warren, H.S. Jr.: Evaluation of a multithreaded architecture for cellular computing. In: Proceedings of

- the 8th International Symposium on High-Performance Computer Architecture, pp. 311–322 (2002). doi:[10.1109/HPCA.2002.995720](https://doi.org/10.1109/HPCA.2002.995720)
9. Chang, J., Sohi, G.S.: Cooperative caching for chip multiprocessors. In: Proceedings of the 33rd Annual International Symposium on Computer Architecture, pp. 264–276 (2006). doi:[10.1109/ISCA.2006.17](https://doi.org/10.1109/ISCA.2006.17)
 10. Chaudhuri, M., Heinrich, M.: SMTp: an architecture for next-generation scalable multi-threading. In: Proceedings of the 31st Annual International Symposium on Computer Architecture, pp. 124–137 (2004). doi:[10.1109/ISCA.2004.1310769](https://doi.org/10.1109/ISCA.2004.1310769)
 11. Chishti, Z., Powell, M.D., Vijaykumar, T.N.: Optimizing replication, communication, and capacity allocation in CMPs. In: Proceedings of the 32nd Annual International Symposium on Computer Architecture, pp. 357–368 (2005). doi:[10.1109/ISCA.2005.39](https://doi.org/10.1109/ISCA.2005.39)
 12. Cray: Chapel Language Specification 0.785. Cray Inc., <http://chapel.cray.com/spec-0.785.pdf> (2009)
 13. Fensch, C., Cintra, M.: An OS-based alternative to full hardware coherence on tiled CMPs. In: Proceedings of the 14th International Symposium on High-Performance Computer Architecture, pp. 355–366 (2008). doi:[10.1109/HPCA.2008.4658652](https://doi.org/10.1109/HPCA.2008.4658652)
 14. Fillo, M., Keckler, S.W., Dally, W.J., Carter, N.P., Chang, A., Gurevich, Y., Lee, W.S.: The M-machine multicomputer. *Int. J. Parallel Programm.* **25**(3), 183–212 (1997). doi:[10.1007/BF02700035](https://doi.org/10.1007/BF02700035)
 15. Hagersten, E.: Personal Communication regarding the verification of the coherence protocol of Sun Microsystems' Enterprise Servers E3000, E4000, E5000 and E6000 (2007)
 16. Hill, M.D.: Multiprocessors should support simple memory-consistency models. *Computer* **31**(8), 28–34 (1998). doi:[10.1109/2.707614](https://doi.org/10.1109/2.707614)
 17. Ifode, L., Singh, J.P., Li, K.: Understanding applications performance on shared virtual memory systems. In: Proceedings of the 23rd Annual International Symposium on Computer Architecture, pp. 122–133 (1996). doi:[10.1145/232973.232987](https://doi.org/10.1145/232973.232987)
 18. Intel: Intel Core2 Extreme Processor X6800 and Intel Core2 Duo Desktop Processor E6000 and E4000 Sequence Specification Update. Intel, document No: 313279-016 (2007)
 19. Kalla, R., Sinharoy, B., Tendler, J.M.: IBM Power5 chip: a dual-core multithreaded processor. *IEEE Micro* **24**(2), 40–47 (2004). doi:[10.1109/MM.2004.1289290](https://doi.org/10.1109/MM.2004.1289290)
 20. Keleher, P., Cox, A.L., Dwarkadas, S., Zwaenepoel, W.: TreadMarks: distributed shared memory on standard workstations and operating systems. In: USENIX Winter 1994 Technical Conference Proceedings, pp. 115–131 (1994)
 21. Kim, C., Burger, D., Keckler, S.W.: An adaptive, non-uniform cache structure for wire-delay dominated on-chip caches. In: Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems, pp. 211–222 (2002). doi:[10.1145/605432.605420](https://doi.org/10.1145/605432.605420)
 22. Kongetira, P., Aingaran, K., Olukotun, K.: Niagara: a 32-way multithreaded sparc processor. *IEEE Micro* **25**(2), 21–29 (2005). doi:[10.1109/MM.2005.35](https://doi.org/10.1109/MM.2005.35)
 23. Kontothanassis, L.I., Hunt, G., Stets, R., Hardavellas, N., Cierniak, M., Parthasarathy, S., Meira, W. Jr., Dwarkadas, S., Scott, M.L.: VM-based shared memory on low-latency, remote-memory-access networks. In: Proceedings of the 24th Annual International Symposium on Computer Architecture, pp. 157–169 (1997). doi:[10.1145/384286.264163](https://doi.org/10.1145/384286.264163)
 24. Krashinsky, R., Batten, C., Hampton, M., Gerding, S., Pharris, B., Casper, J., Asanović, K.: The vector-thread architecture. In: Proceedings of the 31st Annual International Symposium on Computer Architecture, pp. 52–64 (2004). doi:[10.1109/ISCA.2004.1310763](https://doi.org/10.1109/ISCA.2004.1310763)
 25. Kumar, R., Zyuban, V., Tullsen, D.M.: Interconnections in multi-core architectures: understanding mechanisms, overheads and scaling. In: Proceedings of the 32nd Annual International Symposium on Computer Architecture, pp. 408–419 (2005). doi:[10.1109/ISCA.2005.34](https://doi.org/10.1109/ISCA.2005.34)
 26. Kuskin, J., Ofelt, D., Heinrich, M., Heinlein, J., Simoni, R., Gharachorloo, K., Chapin, J., Nakahira, D., Baxter, J., Horowitz, M., Gupta, A., Rosenblum, M., Hennessy, J.L.: The stanford FLASH multiprocessor. In: Proceedings of the 21st Annual International Symposium on Computer Architecture, pp. 325–337 (1994). doi:[10.1109/ISCA.1994.288140](https://doi.org/10.1109/ISCA.1994.288140)
 27. Laudon, J., Lenoski, D.: The SGI Origin: a ccNUMA highly scalable server. In: Proceedings of the 24th Annual International Symposium on Computer Architecture, pp. 241–251 (1997). doi:[10.1145/384286.264206](https://doi.org/10.1145/384286.264206)
 28. Li, K.: IVY: a shared virtual memory system for parallel computing. In: Proceedings of the 1988 International Conference on Parallel Processing, vol. 2, pp. 94–101, Pennsylvania State University Press (1988)

29. Li, M., Sasanka, R., Adve, S.V., Chen, Y.K., Debes, E.: The ALPBench benchmark suite for complex multimedia applications. In: Proceedings of IEEE International Symposium on Workload Characterization, pp. 34–45 (2005). doi:[10.1109/IISWC.2005.1525999](https://doi.org/10.1109/IISWC.2005.1525999)
30. Martin, M.M.K., Hill, M.D., Wood, D.A.: Token coherence: decoupling performance and correctness. In: Proceedings of the 30th Annual International Symposium on Computer Architecture, pp. 182–193 (2003). doi:[10.1109/ISCA.2003.1206999](https://doi.org/10.1109/ISCA.2003.1206999)
31. McNairy, C., Bhatia, R.: Montecito: a dual-core, dual-thread titanium processor. *IEEE Micro* **25**(2), 10–20 (2005). doi:[10.1109/MM.2005.35](https://doi.org/10.1109/MM.2005.35)
32. Scott, S.L.: Synchronization and communication in the T3E multiprocessor. In: Proceedings of the 7th International Conference on Architectural Support for Programming Languages and Operating Systems, pp. 26–36 (1996). doi:[10.1145/237090.237144](https://doi.org/10.1145/237090.237144)
33. Swanson, S., Michelson, K., Schwerin, A., Oskin, M.: WaveScalar. In: Proceedings of the 36th Annual International Symposium on Microarchitecture, pp. 291–203 (2003). doi:[10.1109/MICRO.2003.1253203](https://doi.org/10.1109/MICRO.2003.1253203)
34. Taylor, M.B., Lee, W., Miller, J., Wentzlaff, D., Bratt, I., Greenwald, B., Hoffmann, H., Johnson, P., Kim, J., Psota, J., Saraf, A., Shnidman, N., Strumpfen, V., Frank, M., Agarwal, A., Amarasinghe, S.: Evaluation of the raw microprocessor: an exposed-wire-delay architecture for ILP and streams. In: Proceedings of the 31st Annual International Symposium on Computer Architecture, pp. 2–13 (2004). doi:[10.1109/ISCA.2004.1310759](https://doi.org/10.1109/ISCA.2004.1310759)
35. Vachharajani, M., Vachharajani, N., August, D.I.: The liberty structural specification language: a high-level modeling language for component reuse. In: Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation, pp. 195–206 (2004). doi:[10.1145/996893.996865](https://doi.org/10.1145/996893.996865)
36. Verghese, B., Devine, S., Gupta, A., Rosenblum, M.: Operating system support for improving data locality on CC-NUMA compute servers. In: Proceedings of the 7th International Conference on Architectural Support for Programming Languages and Operating Systems, pp. 279–289 (1996). doi:[10.1145/237090.237205](https://doi.org/10.1145/237090.237205)
37. Woo, S.C., Ohara, M., Torrie, E., Singh, J.P., Gupta, A.: The SPLASH-2 programs: characterization and methodological considerations. In: Proceedings of the 22nd Annual International Symposium on Computer Architecture, pp. 24–36 (1995). doi:[10.1145/223982.223990](https://doi.org/10.1145/223982.223990)
38. Zeffer, H., Hagersten, E.: A case for low-complexity MP architectures. In: Proceedings of the Conference on Supercomputing (2007). doi:[10.1145/1362622.1362648](https://doi.org/10.1145/1362622.1362648)
39. Zeffer, H., Radović, Z., Karlsson, M., Hagersten, E.: TMA: a trap-based memory architecture. In: Proceedings of the 20th Annual International Conference on Supercomputing, pp. 259–268 (2006). doi:[10.1145/1183401.1183438](https://doi.org/10.1145/1183401.1183438)
40. Zhang, M., Asanović, K.: Victim replication: maximizing capacity while hiding wire delay in tiled chip multiprocessors. In: Proceedings of the 32nd Annual International Symposium on Computer Architecture, pp. 336–345 (2005). doi:[10.1109/ISCA.2005.53](https://doi.org/10.1109/ISCA.2005.53)