

RegionScout: Exploiting Coarse Grain Sharing in Snoop-Based Coherence

Andreas Moshovos
Electrical and Computer Engineering
University of Toronto
www.eecg.toronto.edu/aenao

Abstract

It has been shown that many requests miss in all remote nodes in shared memory multiprocessors. We are motivated by the observation that this behavior extends to much coarser grain areas of memory. We define a region to be a continuous, aligned memory area whose size is a power of two and observe that many requests find that no other node caches a block in the same region even for regions as large as 16K bytes. We propose RegionScout, a family of simple filter mechanisms that dynamically detect most non-shared regions. A node with a RegionScout filter can determine in advance that a request will miss in all remote nodes. RegionScout filters are implemented as a layered extension over existing snoop-based coherence systems. They require no changes to existing coherence protocols or caches and impose no constraints on what can be cached simultaneously. Their operation is completely transparent to software and the operating system. RegionScout filters require little additional storage and a single additional global signal. These characteristics are made possible by utilizing imprecise information about the regions cached in each node. Since they rely on dynamically collected information RegionScout filters can adapt to changing sharing patterns. We present two applications of RegionScout: In the first RegionScout is used to avoid broadcasts for non-shared regions thus reducing bandwidth. In the second RegionScout is used to avoid snoop induced tag lookups thus reducing energy.

1 Introduction

There are workload, technology and cost trends that make shared memory multiprocessors an increasingly popular architecture [16,34]. Today there are applications (e.g., databases, file and mail servers, multimedia/entertainment and communications) with sufficient parallelism that shared memory multiprocessors can exploit. In addition, the increasing levels of chip integration make single chip/module multiprocessors viable and an attractive alternative for utilizing additional on-chip resources [4,14,32]. Also, the increased cost and complexity of building modern processors make using multiple identical cores a cost-effective option for high integration devices. Finally, there is a proliferation of devices that look increasingly like small-scale shared memory multiprocessors (e.g., cell phones and game consoles). Accordingly the focus of this work is on small-scale shared multiprocessors.

With the increased popularity of shared memory multiprocessors comes renewed interest in techniques for

improving their efficiency and performance, particularly in techniques related to coherence and for both small and large scale systems. This is exemplified by the recent advances in coherence speculation (e.g., [17,19,22,23,24,28,29]) and in power-aware coherence [12,27,30]. The same cost and technology trends that favor using multiple identical cores also favor techniques that are as little intrusive as possible at the hardware and software levels. At the same time, the abundance of on-chip resources creates opportunities for novel coherence implementations [4].

We propose RegionScout a technique that exploits coarse grain sharing patterns in snoop-based shared memory multiprocessors with potential applications in reducing bandwidth, latency and energy. We observe that many shared memory requests not only do not find a matching block in any remote node but also they do not find a block in the same region (where a region is a continuous, aligned memory area whose size is a power of two).

RegionScout comprises a family of filters that dynamically observe coarse grain sharing and allow nodes to detect in advance that a request will miss in all remote nodes. Such information is not available in conventional snoop- or directory-based coherence. With RegionScout when a node sends a request in addition to block-level sharing information it receives region-level sharing information. If a region is identified as not shared subsequent requests for *any* block within the region from the same node are identified as non-shared without having to probe any other node. Normal coherence activity allows the detection of regions that become shared preserving correctness. RegionScout filters utilize imprecise information about the regions that are cached in each node in the form of a hashed bit vector [5]. For this reason, they detect many but not all requests that would miss in all other nodes. This loss in coverage allows RegionScout to be completely transparent to software and avoids imposing artificial constraints on what can be cached simultaneously. RegionScout filters require no changes to the underlying coherence mechanisms: when possible they simply return a “not shared” response without having to consult the existing coherence protocol. Their implementation is flexible allowing a trade-off between filter storage cost and accuracy. To demonstrate the utility of RegionScout filters we investigate two applications. In the first, RegionScout filters are used to avoid snoop-induced tag lookups thus reducing energy in the memory hierarchy. In the second, RegionScout filters are

used to avoid broadcasts for some requests. Techniques for avoiding broadcasts are becoming important even for small-scale multiprocessors. This is because technology trends favor switch-based point-to-point interconnects as opposed to busses. Avoiding broadcasts improves performance in several ways: it reduces latency directly since the originating node does not have to wait for every other node to respond, it reduces contention on the interconnect thus reducing latency indirectly and finally it frees bandwidth which may be used more fruitfully.

The rest of this paper is organized as follows: In Section 2 we detail our simulation environment and methodology. In Section 3 we show that many requests do not find a block in the same region in remote nodes for various shared memory multiprocessor configurations. We comment on related work in Section 4. We present the principle of operation for RegionScout filters in Section 5. We also discuss their implementation and additional benefits that come “for free”. We explain how RegionScout fundamentally differs from directories in Section 5.4. In Section 6, we show experimentally that RegionScout filters are effective and can reduce bandwidth and energy. Finally, we summarize this work in Section 7.

2 Methodology

Our simulator is based on SimpleScalar v2.0 [6] with Manjikian’s shared memory extensions [21]. We made extensive changes to the simulator and the shared memory support libraries since Manjikian’s model is limited to direct-mapped caches, uses pseudo system calls for synchronization and does not implement all PARMACS macros necessary to run all the SPLASH2 benchmarks. For synchronization we modelled the *load linked* and *store conditional* MIPS instructions which we used to implement all synchronization primitives including MCS locks [25]. We added support for set-associative caches, shared or private L2 caches and a bus model. We used a modified version of GNU’s gcc v2.7.2 to produce optimized PISA binaries for the SPLASH2 benchmarks [33] shown in Table 1. We do not include Water as a result of a math library bug.

We use the two models of shared multiprocessor systems shown in Figure 1. Under model *LocalL2* each node has private L1 and L2 caches and there is a shared L3 or main memory. This is representative of conventional SMPs. Under model *SharedL2* nodes have private L1 caches and there is a shared L2. This is representative of simple CMPs [14] and to a lesser extent of processors like Power4 [32] that are popular as building blocks for larger SMP systems. We include both models since energy- and bandwidth-wise they behave differently (with the differences being primarily a function of block and cache sizes). We use a MESI coherence protocol.

We modelled systems with two, four, eight or 16 processors. Unless otherwise noted, LocalL2 systems have L2 caches of the specified size that are eight-way set-associative with 64-byte blocks and 32Kbyte L1 caches that are four-way set-associative with 32-byte blocks. In SharedL2 systems the L1 caches use 32-byte blocks and are four-way set-associative and of the specified size.

Table 1. SPLASH2 applications and input parameters.

Benchmark	Input Parameters
<i>barnes</i>	16k particles
<i>cholesky</i>	tl29.O
<i>fft</i>	256k complex data points
<i>fmm</i>	16k particles
<i>lu (contig.)</i>	512x512 matrix, B=16
<i>ocean (contig.)</i>	258x258 grid w/ defaults
<i>radiosity</i>	-batch -room
<i>radix</i>	2M keys
<i>raytrace</i>	balls4.env
<i>volrend</i>	256x256x126 voxels

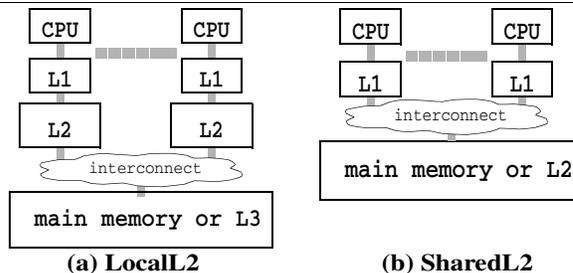


Figure 1: Two shared memory multiprocessor models. (a) *LocalL2*: local L1 and L2 caches and a shared L3. (b) *SharedL2*: local L1 caches and a shared L2. The interconnect is functionally a bus but may be implemented via other components [11].

3 Motivation

For clarity, we first define the terms *region*, *region miss* and *global region miss*. A *region* is an aligned, continuous section of memory whose size is a power of two. A request for a block B in a cache C results in a *region miss* if C holds no block in the same region as B including B itself. We were motivated by the observation that often memory requests result in region misses in *all* remote caches, or, in a *global region miss* or *GRM* for short. This is a generalization of observations made by earlier work in snoop energy reduction. Specifically, earlier work has shown that this property holds for cache blocks [27] or pages [12]. Intuitively, one would expect that this property would apply to other region sizes as well. This intuition holds true as shown in Figure 2 where we measure the global region miss ratio for systems with different number of nodes and cache sizes. The *global region miss ratio* is the fraction of all coherent memory requests that result in a global region miss. We consider regions of 256 bytes up to 16K bytes and systems with two, four, eight and 16 processors. In part (a) the per node L2 cache size is varied from 512K bytes to four megabytes. In part (b) the per node L1 cache size is varied from 32Kbytes to 128Kbytes. We do not model SharedL2 systems with 16 processors since they do not appear to be a reasonable design (i.e., the L1 caches do not effectively reduce coherence traffic for such a system). In general, the average global region miss ratio is inversely proportional to private cache sizes, node count and region size. (An anomaly is observed in (a) when the number of nodes increases from four to eight. This is primarily a result of our barrier implementation and of the memory allocation of related structures.) Even if we consider the worst case (four node LocalL2 with four Mbyte L2 caches) more than one in

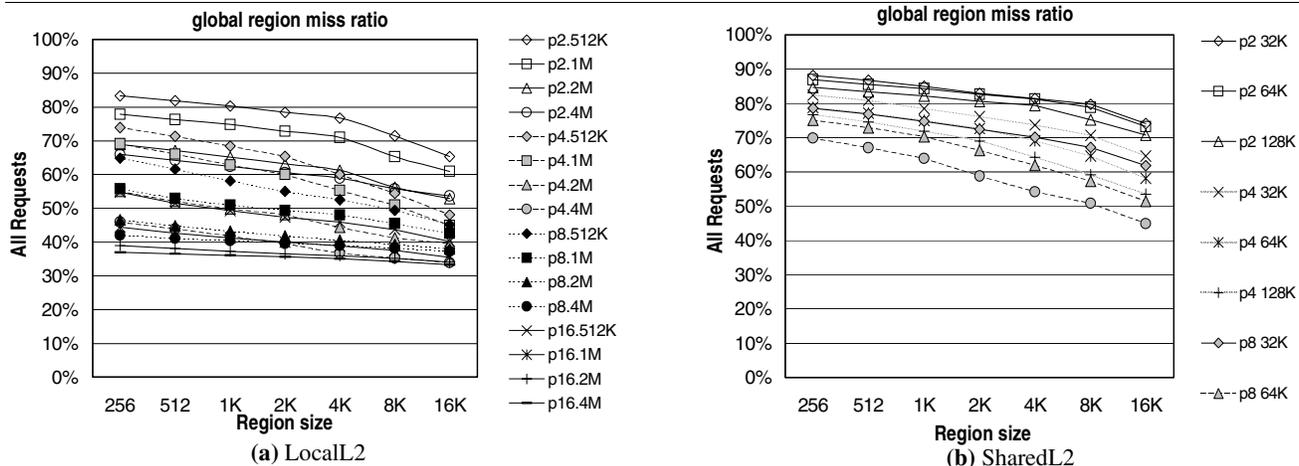


Figure 2: Average global region miss ratio for various shared memory multiprocessor systems (different curves) and region sizes (X-axis). (a) Systems with local L1 and L2 caches and shared L3. (b) Systems with local L1 caches and a shared L2. Labels are $pN.S$ where N is the number of nodes and S is the L2 and L1 cache size for parts (a) and (b) respectively.

three requests results in a global region miss for a region as large as 16K bytes. For typical organizations the fraction of global region misses is much higher. For example, for the four node LocalL2 system with 512K L2 caches, the average global region miss ratio varies from 74% down to 48% depending on the region size. For the SharedL2 four node system with 64K L1 caches the average global region miss ratio varies from 79% down to 58%. Individual program behavior varies greatly as we will see in Section 6.1. Global region misses decrease as the node count increases. But, the relative decrease gradually diminishes suggesting that potential may exist even with more nodes.

These results suggest that if we could detect or predict *a priori* that a shared memory request would result in global region miss then we could avoid broadcasts for anywhere between 88% and 34% of all requests depending on the specific organization. There are potential energy, latency and bandwidth benefits:

1. We could reduce the bandwidth requirements by avoiding the snoops that would result in a global region misses. Such snoops are unnecessary.
2. We could reduce energy by avoiding the tag lookups in all remote caches similarly to previous work [12,27]. Furthermore, if the interconnect permits it, we could avoid communicating with all remote nodes reducing energy even further. For example, in systems where broadcasts are implemented over a switch-based interconnect this may be possible [11,22]. This may also be possible in bus-based CMPs. or systems that use a separate bus for snoops.
3. Finally, by avoiding some snoops we could reduce the latency of the corresponding memory requests.

Of the aforementioned potential applications we limit our attention to bandwidth and tag lookup energy reduction.

Previous work has also shown that a significant fraction of requests that hit in some remote caches rarely hit in all or even many of them [12,27,30]. In Section 5.3, we validate this observation for region sizes other than a block or a page and

explain how our techniques can exploit this behavior to avoid many of these tag lookups.

4 Related Work

Previous work on snoop energy reduction relies on similar phenomena as RegionScout. In *Jetty* proposed by Moshovos *et. al.*, each node avoids many snoop-induced lookups that would otherwise miss [13]. Nodes maintain two structures that respectively represent a subset of blocks that are not cached (exclusive Jetty) and a superset of blocks that are cached (inclusive Jetty). The key difference is that with RegionScout a requesting node can determine *in advance* that a request would miss in all other nodes. With Jetty every node still snoops all requests. Advance knowledge of global region misses allows optimizations (such as reducing bandwidth and energy in the interconnect) that are *impossible* with Jetty. As we explain in Section 5.3, a structure used by RegionScout can be used as simplified inclusive Jetty avoiding tag-lookups for requests that hit in some but not all remote caches. Because RegionScout can use regions that are much larger than a block, much smaller structures are sufficient to capture most of the benefits as we demonstrate in Section 6.4.

The *Page Sharing Table* (PST), proposed by Ekman *et. al.*, uses vectors that identify sharing at the page level [12]. It can be thought of as a partial, distributed page-level directory. Every node keeps precise information about the pages it is caching. This information is used to form a page-level sharing vector in response to coherence requests. Subsequent requests are snooped only by those nodes that do have blocks within the same page and thus energy is reduced. Additional bus lines are required for broadcasting and collecting the sharing vectors. The PST is tightly coupled with the TLB but its page size can be smaller (pages of 1K and 4K were considered since this allows the PST to work with virtual addresses). In rare cases when it does not have enough space to track all locally cached pages the PST has to be turned off. Recovering from this state requires flushing the caches. While the PST utilizes precise page-level information, RegionScout targets

the frequent and often the common case of regions that are not shared at all. For this reason and because it utilizes imprecise information, RegionScout requires a single additional global signal as opposed to a sharing vector and requires only surgical changes to the existing infrastructure. RegionScout never has to be turned off for correct operation and disassociates the choice of region size from the rest of the design provided that physical addresses are used.

Saldanha and Lipasti proposed serial snooping to reduce energy in shared multiprocessors [30]. Ekman *et. al.*, evaluated Jetty and serial snooping for chip multiprocessors demonstrating little or no benefit [13]. Li *et. al.*, proposed the *thrifty barrier* to reduce processor energy. It uses wait time prediction to selectively place a processor into a lower power state while it waits on a barrier [20].

Recently, several coherence predictors have been proposed [17,19,22,23,24,28,29]. While some predictors may be able to capture the same behavior (e.g., by predicting the destination set or by using an “all or nothing” policy [23]) RegionScout obviates the need to predict the destination set for many requests thus freeing bandwidth for more aggressive optimizations and reducing the working set of accesses that a predictor will have to track. While the potential for synergy exists, an investigation of combining RegionScout with coherence predictors is beyond the scope of this paper.

A preliminary evaluation of RegionScout appears in [26]. Cantin, Lipast and Smith have also proposed exploiting coarse sharing for snoop coherence bandwidth reduction [8].

5 RegionScout Filters

RegionScout filters allow each node to *locally* determine that a request will result in a global region miss and thus avoid the corresponding remote snoops and transactions. Informally, whenever a node issues a memory request it also asks all other nodes whether they hold any block in the same region. If they do not, it records the region as not shared. Next time the node requests a block in the same region it knows that it does not need to probe any other node. Correctness is maintained since whenever another node requests a block in the same region, it will broadcast its request invalidating the “not shared” region records held by other nodes. What allows RegionScout to be effective yet inexpensive is that it works for *most* and not all requests that would result in a global region miss.

Formally, the RegionScout filters comprise two structures local to each node: (a) a “not shared” region table, or *NSRT*, and (b) a *cached region hash*, or *CRH*. The NSRT records non-shared regions. The CRH is a Bloom filter (similar to the inclusive Jetty [27]) that records a *superset* of all regions that are locally cached. Example organizations of CRH and NSRT are given in Section 5.2.

Here is how RegionScout works: Initially, all caches, CRHs and NSRTs are empty. Whenever a node *N* issues a memory request, the other nodes respond normally via the existing coherence protocol but in addition, using their CRH, they report (see Section 5.2.3 for a discussion) whether they cache *any* other block in the block’s region. If the region is not shared node *N* then record it in its NSRT. Next time node *N* requests a block it first checks its NSRT. If a valid record is

found then it knows that no other node holds this block and can avoid broadcasting this request. To ensure correctness it is imperative to invalidate NSRT entries whenever any other node requests a block in the corresponding region. This is easily done as part of the existing protocol actions. Specifically, if another node *N'* requests a block in the same region, it too will check its own NSRT and since it will not find a valid record it will broadcast its request to all other nodes. The NSRT of node *N* will then invalidate the corresponding entry and subsequent requests will be broadcast as they should.

Key to RegionScout’s success is the ability at each node and given a block to determine whether there is any other block in the same region that is locally cached. One possibility would be to use a table to keep a *precise* record of all regions that are locally cached similarly to what is done for pages in PST [12]. The size of the table would artificially limit what can be cached and special actions would be required to avoid exceeding these limits. RegionScout avoids these issues by using the CRH, an *imprecise* record of all locally cached regions.

Without the loss of generality we limit our attention to the LocalL2 model. CRH works as follows: Whenever a block is allocated in or evicted from the L2, we use the region part of its address and hash into the CRH. Because there are much fewer CRH entries than regions many regions may map onto the same CRH entry. Each CRH entry counts the number of locally cached blocks in the regions that hash to it. Accordingly, when a block is allocated in the L2 we increment the corresponding CRH entry and when a block is evicted we decrement it. These updates are local at each node. Given a remote request for a block, the CRH can be used to indicate whether it would result in a region miss. If the corresponding CRH counter is zero then we know for sure that no block within the region is cached locally. Otherwise, blocks in the same region *may be* cached. It is the uncertainty of the latter response that allows us to use a small structure effectively. Figure 3 shows an example of RegionScout at work.

5.1 RegionScout as a Layered Extension

RegionScout can be implemented as a layered extension over existing coherence protocols and multiprocessor systems. No changes are required in the underlying coherence protocol, cache hierarchy and software (with the exception of reporting clean replacements to the CRH). RegionScout can operate in parallel with existing structures, inhibiting snoops and broadcasts by intervening when necessary. To existing mechanisms it appears as if the coherence protocol reported a miss. RegionScout is also completely transparent to software and the operating system. It does not impose any artificial limits on the choice of region size or on what can be locally cached. Finally, since it uses dynamically collected information it can adapt to changing sharing behavior identifying regions that are only temporarily not shared. Since every request has to probe the NSRT prior to being issued on the interconnect the latency of handling such requests will increase. The NSRT is comparatively small (we consider NSRTs of up to 64 entries) and as such its latency will be

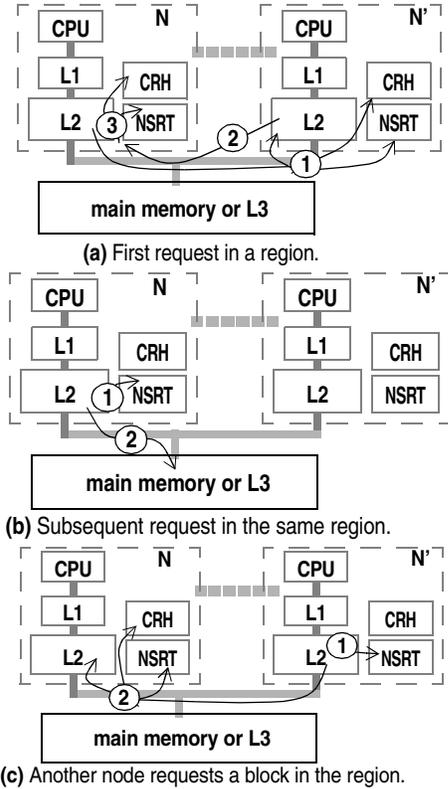


Figure 3: An example illustrating how RegionScout works. Without the loss of generality we use the LocalL2 model and show only two nodes. (a.1) Node N requests block B in region RB. The request is broadcast to all nodes (N first checked its NSRT and since it was empty it found no matching entry for region RB). (a.2) All remote nodes probe their CRH and report that they do not cache any block in region B. (a.3) Node N records RB as not shared and increments the corresponding CRH entry. (b.1) Node N is about to request block B' in region RB and first checks that its NSRT. (b.2) Since an entry is found the request is sent only to main memory. (c.1) Node N' requests block B'' in region RB. It first checks its NSRT. (c.2) Since the region is not recorded in its NSRT, it broadcasts its request. Node N invalidates its NSRT entry since now RB is shared.

comparatively small. As we explain in Section 5.3 the CRH can also be used as a simplified inclusive Jetty filter filtering many snoop-induced tag lookups that would otherwise miss.

5.2 Structures

5.2.1 NSRT and CRH. Figure 4 illustrates the NSRT and CRH organization and how physical addresses are used to index them (we assume a 42 bit physical address space and 16K regions). NSRT is a simple table with entries comprising a valid bit (V) and a region tag (the upper part of the address). The NSRT can be set-associative. Prior to issuing a coherent request each node checks whether a matching record exists in its NSRT. If so, it knows that this block is not shared. NSRT entries are evicted either as a result of limited space or when a remote node requests a block in a matching region.

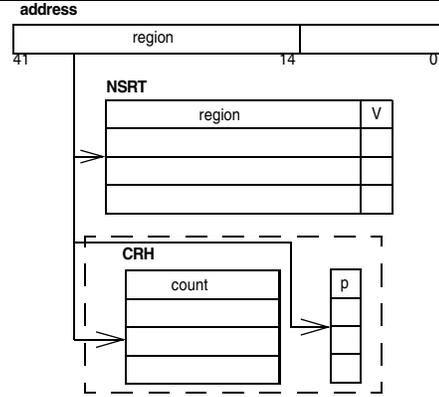


Figure 4: NSRT and CRH organization for 16Kbyte regions.

CRH is a table whose entries comprise a counter (*count*) field and a present bit (*p*). Essentially it is an inclusive Jetty with just one array [27] or a Bloom filter implementation [5,18]. The count field counts the number of matching blocks in the regions that map onto it. In the worst case, all blocks in the cache would belong to the same region. Hence, the count field needs $\lg(\text{Cache Size}/\text{Block Size})$ bits. The p-bit indicates whether count is non-zero. We use the p-bits to reduce energy and delay when probing the CRH [27]. Updating the CRH is done when blocks are allocated or evicted from the L2 or the L1 for models LocalL2 and SharedL2 respectively. Only when a count entry changes value from or to zero the p-bits are updated.

Small CRHs and NSRTs are sufficient for our purposes. For example, a 256 entry CRH needs 256 bits for the p-bit array and less than 4Kbits for the counter array assuming an 1Mbyte cache with 64 byte blocks. A 64 entry NSRT requires 1Kbits.

In principle the choice of the CRH indexing function can impact how well it manages to differentiate amongst cached and uncached regions. In this work, we simply use the low order bits of an address.

5.2.2 Simplified CRH Counter Array. In the inclusive Jetty design the count fields are updated arithmetically. We propose a simpler design that replaces the adder with a reversible linear feedback shift register or LFSR [2]. Appropriately designed n-bit LFSRs generate sequences of $2^n - 1$ states. LFSRs are used as pseudo-random number generators in many applications including testing and communications. LFSRs are much simpler, faster and energy efficient than arithmetic counters. For example, for a 4Mbyte L2 cache (worst case scenario we considered) we need 2^{16} states or 17 bit LFSRs. Each of these requires just eight XOR gates or few tens of transistors. The key advantage of the LFSR-based design is that LFSRs can be embedded in the SRAM array thus drastically reducing power. Due to space limitations we do not describe the complete design here.

5.2.3 Communicating Region Sharing Information and Inhibiting Snoops. In a bus interconnect, an additional, wired-OR bus signal, *RegionHit*, can be used to identify global region misses. This represents a small overhead compared to typical bus implementations that use several tens of signals

(e.g., approximately 90 in MIPS R10000 [1]). Prior to issuing a request RegionHit is de-asserted. In response, a node whose CRH reports a region “hit” asserts RegionHit. On global region misses no node will assert RegionHit. To inhibit snoops in a true bus implementation RegionHit can be overloaded as follows: Prior to issuing a request that would otherwise result in a global region miss (as reported by the NSRT) a node asserts RegionHit. Other nodes can sample RegionHit prior to snooping and thus avoid snooping altogether. For other requests RegionHit is de-asserted as before so that other nodes can snoop and report region hits. The information necessary for detecting global region misses and inhibiting snoops is a single bit. For this reason we expect that it will represent a small overhead for other interconnect architectures also. Moreover, it may be possible to embed this information in the existing protocol (i.e., in the control information) with no physical overhead.

5.3 Avoiding Tag Lookups for Non-Global Region Misses

As described RegionScout can avoid broadcasts for those requests that would result in a global region miss. Additional benefits are possible “for free” for requests that would result in region misses in *some* but not all remote nodes. The CRH can be used as a simplified inclusive Jetty as follows: Whenever node N makes a request, all other nodes probe their CRHs to determine whether they will report a region miss. If a node N’ determines that it has no matching block in the same region then it does not need to probe its local L2 or L1 tag array (for LocalL2 and SharedL2 respectively). By avoiding these lookups we can further reduce tag energy and tag bandwidth requirements. This comes at the expense of increased latency while probing the local tag arrays in response to a remote request. Since, only the small p-bit array is probed, the latency penalty will be small.

How much potential is there for this optimization? In the interest of space we limit our attention to the 16K byte region and to four node LocalL2 and SharedL1 systems with 512K L2 and 64K L1 caches respectively. Table 2 reports the remote region hit count distribution. The first column corresponds to global regions misses. The fraction of requests that incur a region miss in some remote caches (columns “1” and “2”) is significant for all programs. In barnes, radiosity, radix, raytrace, volrend and to a lesser extent fmm many requests result in a region hit in all other nodes (column “3”). This may seem to contradict previous findings that most requests miss in remote caches but we emphasize that here we look at much larger regions.

5.4 RegionScout and Directories

RegionScout provides region-level sharing information that is not available in a conventional directory. It is possible to track such coarse grain information in a directory and our results of Section 3 may serve as good motivation for doing so. But there are bandwidth and complexity trade-offs since nodes would have to communicate block evictions in addition to regular coherence traffic. For this reason, even in this case opting for RegionScout may be preferable offering some of the benefits of a full-blown region-level directory at a

Table 2. Remote region hit count distribution for four node LocalL2 and SharedL2 systems with 512K L2 and 64K L1 caches respectively. The region is 16Kbytes. Column “0” corresponds to global region misses.

	Four Nodes, 16Kbyte Regions			
	LocalL2 Region Hit Count (%)			
	0	1	2	3
barnes	9.38	13.89	12.84	63.89
cholesky	87.87	6.22	4.88	1.03
fft	95.22	4.47	0.19	0.12
fmm	35.20	28.22	19.69	16.9
lu	48.78	44.95	3.73	2.55
ocean	85.52	11.77	1.66	1.05
radiosity	37.75	35.56	4.15	22.54
radix	33.32	15.15	25.99	25.54
raytrace	16.01	10.86	28.23	44.90
volrend	30.94	5.07	37.87	26.13
average	47.99	17.61	13.92	20.46
	SharedL2 Region Hit Count (%)			
	0	1	2	3
barnes	15.12	14.03	10.61	60.24
cholesky	89.36	7.88	2.27	0.49
fft	99.14	0.74	0.02	0.09
fmm	47.00	21.32	5.61	26.06
lu	58.45	39.28	1.50	0.77
ocean	94.75	3.73	1.26	0.26
radiosity	31.10	19.07	1.97	47.86
radix	51.55	28.13	15.24	5.09
raytrace	61.16	23.24	9.80	5.80
volrend	33.31	9.72	26.96	30.01
average	58.09	16.71	7.52	17.66

minimal cost. In particular, RegionScout can be used to inhibit broadcasts for non-shared regions. It is also straightforward to extend RegionScout by including per node sharing information similarly to what was done in the PST [12]. In this case, RegionScout can act as an imprecise, distributed region-level directory. The results of Section 5.4 serve as motivation for doing so. However, this investigation is beyond the scope of this paper.

6 Evaluation

In the interest of space and unless otherwise noted we limit our attention to the LocalL2 and SharedL2 models with 512K L2 and 64K L1 caches respectively. In Section 6.1 we show the global region miss behavior of individual programs as a function of region size. In Section 6.2 we demonstrate that practical RegionScout filters can capture many global region misses. We chose a large enough NSRT and focus on region and CRH size. We identify the trade offs among region size, detected global region misses and RegionScout filter storage requirements. In Section 6.3 we demonstrate that RegionScout can be used to avoid broadcasts in snoop coherence. In Section 6.4 we show that it can be used to avoid many snoop induced tag lookups thus reducing tag energy and also compare it to Jetty. The two applications also serve to demonstrate that we can tailor the RegionScout filters to meet different trade-offs. The specific choice depends on what is the primary consideration: energy or performance. In Section 6.5 we demonstrate that the potential exists for RegionScout to be useful under commercial workloads.

6.1 Per Program Global Region Miss Behavior

Since average behavior can be misleading we also look at individual program behavior. Figure 5 reports the global region miss ratio per program for regions of 256 through 16K bytes. Parts (a) and (b) are for models LocalL2 and SharedL2 respectively. In the interest of space we restrict our attention to four node systems. In addition to the intrinsic sharing characteristics of each program the observed global miss ratios are mostly inversely proportional to the region and cache sizes. Programs can be classified in three categories based on how sensitive they are to the choice of region size. Cholesky, fft and ocean exhibit high global region miss ratios and are mostly insensitive to region (and cache) size. In barnes, fmm, raytrace and volrend the global region miss ratio decreases almost linearly as the region size increases. Finally, radix, radiosity and to a lesser extent, lu, exhibit abrupt changes in their global miss ratio when the region size increases above 2K, 8K and 4K respectively. For the smallest region size of 256 bytes the global miss ratio is above 37% for all programs under both models. With the largest region of 16K bytes the miss ratio remains above 30% for all programs except barnes and raytrace under the LocalL2 model (and under the SharedL2 for barnes). This result is important as it suggests that sufficient global misses occur even when we look at very coarse grain sharing. Using large regions is attractive as smaller structures could be used to track them.

6.2 Can Practical RegionScout Filters Detect Many Global Region Misses?

We have seen that sufficient global region misses occur in the programs we studied. The question we answer affirmatively in this section is whether practical RegionScout filters can capture most of them. For this purpose we use an NSRT with 64 entries (16 sets, 4-way set associative) and measure the filter rate for different CRHs. We define the *filter rate* as the fraction of *all* requests that are detected as global region misses by the RegionScout filter. A global region miss is detected when the originating node finds a matching entry in its NSRT prior to issuing the request. We found that an NSRT of 64 entries approximates one with infinite entries. We consider CRHs of 256 through 2K entries and regions of 2K through 16K bytes. The resulting average filter rates are shown in Figure 6. Each of the curves corresponds to a different region size and to a system with a different number of nodes. The curves are identified as $pN.C.RS$ where N is the number of nodes, C is the cache size (L2 for LocalL2 and L1 for SharedL2) and S is the region size. While we have seen that using smaller regions typically results in higher global miss ratios here we observe a trade-off between region and CRH size. In most cases using a larger region results in a higher filter rate with the difference being greater for the smaller CRHs (i.e., 256 or 512 entries). The CRH is an imprecise representation of cached regions. Using larger regions results in fewer regions and improves the CRH's ability to separate among them. Only when the size ratio "CRH over cache" becomes large enough we see a benefit in using smaller region sizes. This can be seen under the

SharedL2 model when the local caches are only 64K and for the 2K entry CRH

Figure 7 shows the per program filter rates for various CRH sizes. We restrict our attention to the 16K byte regions and to four node systems with 512K L2 and 64K L1 caches. Under the SharedL2 model (part (b)) it is conflict misses that dominate traffic and this results in much higher temporal and spatial locality in the request stream. For this reason and given that there are fewer regions per node even the 256 entry CRH can capture most global region misses. While the filter rates increase with CRH size these improvements are modest. Under the LocalL2 model in part (a), the filter rate is more sensitive to CRH size. This is because a much larger portion of data and thus more regions remain resident in the caches. Under LocalL2, larger CRHs result in significantly higher filter rates for some programs (e.g., ocean, cholesky and fmm). Lu, barnes and raytrace are mostly insensitive to CRH size. An anomaly is observed for radiosity where the filter rate decreases when the CRH entries are increased from 256 to 512. In this case we have thrashing in the NSRT: more regions are identified as non-shared but their base addresses are such that they trash few sets in the NSRT. This thrashing persists for the 2K CRH but the resulting filter rate is higher since additional non-shared regions are identified which map to different sets in the NSRT.

In summary, we have found that practical RegionScout filters can detect most global region misses. While there are more global region misses for smaller regions, detecting them requires larger RegionScout filters. For the configurations we studied using a large region (e.g., 16K) is typically better. For a given region size, filter rates improve with CRH size.

6.3 Bandwidth Reduction

We demonstrate that RegionScout can reduce the broadcasts in snoopy coherence systems and hence reduce bandwidth demand. We use a first-order approximation model to identify trends in the reduction in average processing time for snoops in a bus-based system. In lack of an accurate timing model we do not demonstrate any other tangible benefits.

Emerging technology trade-offs favor point-to-point links over buses for high performance interconnects [11,22]. Accordingly, we modelled two systems. The first is based on Sun's Starfire architecture [9]. It uses a bus for control information (all requests appear on this bus) and a switch-based interconnect for data (Starfire, targeted at a much larger number of processors uses four buses for snoops). The second model is for a hypothetical system where nodes are connected via a switch for both data and control. For this application we can afford to use the larger RegionScout filters as performance is our primary goal. Accordingly, in all experiments the NSRT has 64 entries and is four-way set-associative and the CRH has 2K entries.

Bus-Based Model: Figure 8(a) reports the relative average traffic ratio ($TrafficWithRegionScout$)/($TrafficWithoutRegionScout$) on the control bus for the Starfire-like model, for regions of 2K up to 16K and for LocalL2 and SharedL2 systems with different number of nodes. The systems are identified as $pN.C$ where N is the number of nodes and C is the cache size (512K

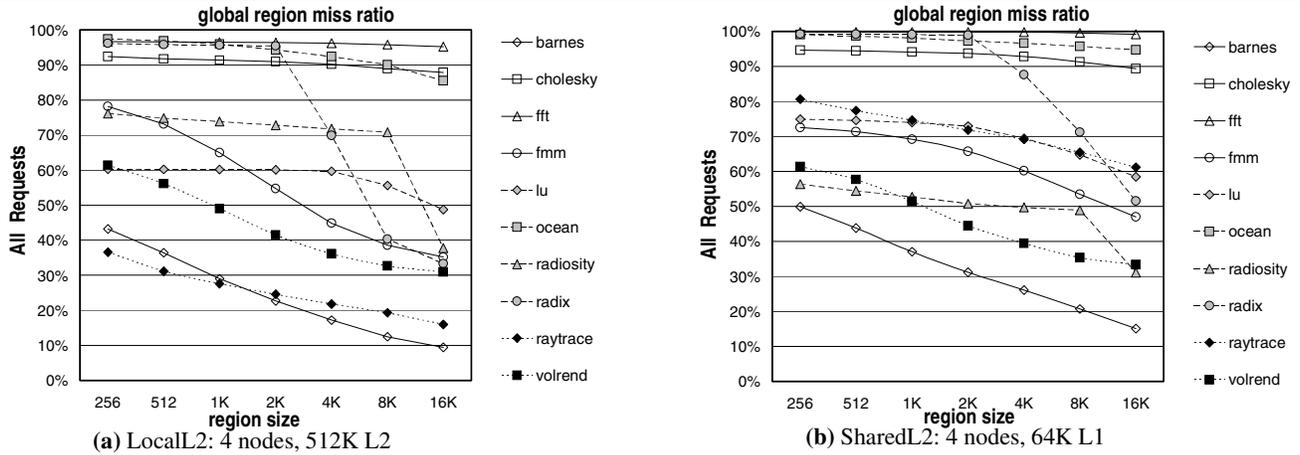


Figure 5: Global Region Miss ratios per program for various region sizes.

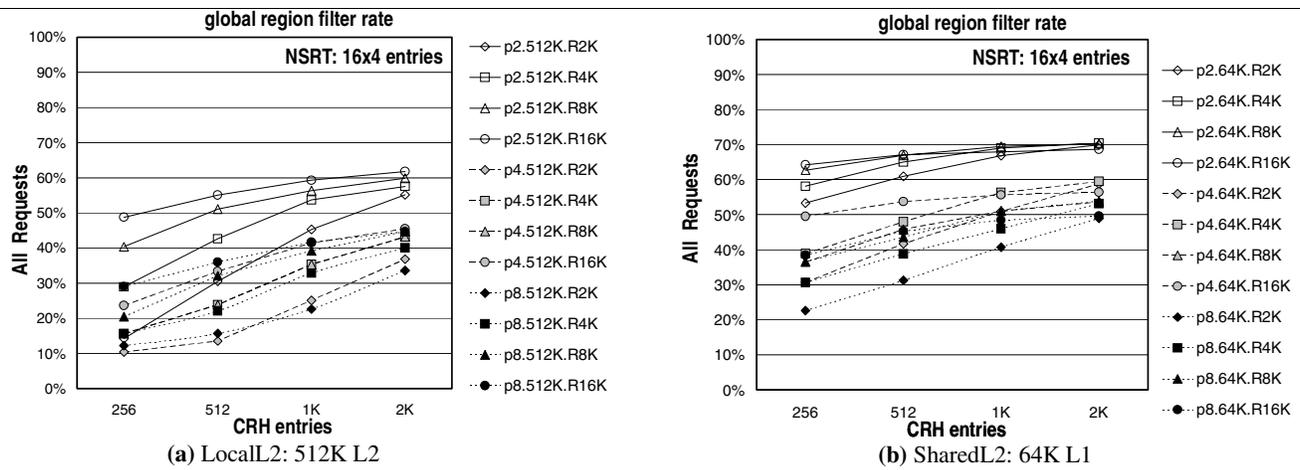


Figure 6: Average global region miss filter rates (see text for the definition) for various CRH sizes (X-axis). The NSRT has 64 entries organized in 16 sets of four entries each. The curves correspond to systems with different number of nodes N and to different region sizes S as identified by the $pN.C.RS$ labels. C is the size of the L2 and L1 caches for models LocalL2 (a) and SharedL2 (b) respectively.

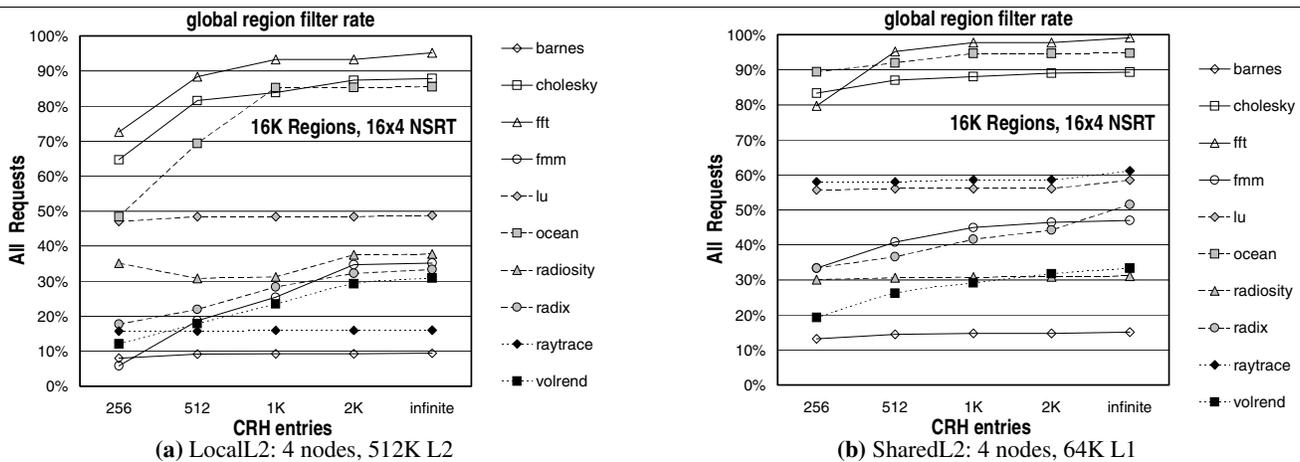


Figure 7: Filter rates per program for 16K byte regions on four node systems with 512K L2 (a) and 64K L1 (b) caches respectively. The NSRT is fixed at 64 entries (16 sets, 4-way set-associative) and the number of CRH entries varies as shown on the X-axis.

L2 and 64K L1 for LocalL2 and SharedL2 respectively). The reduction is proportional to the filter rate taking into account the additional traffic due to writebacks. With eight processors RegionScout reduces traffic to 73% and 65% of the base system for the LocalL2 and SharedL2 models respectively.

We model the control bus as an M/M/1 queue. We restrict our attention to the 16K region and to four way systems. We estimate response times for snoop processing latencies of two, four, eight and 12 *processor* cycles (labels c2, c4, c8 and c12 respectively). These latencies are optimistic given existing systems and the relative scaling properties of bus vs. processor cycle times. Using optimistic latencies is acceptable in this experiment since the slower the bus the higher the impact of reducing bandwidth demand. To estimate the arrival rate we used an average IPC value of 1.4 as per published measurements for scientific applications on commercial systems [10]. Larger IPC values would have lead to higher benefits from RegionScout. Figures 8(b) and 8(c) report results for the LocalL2 and SharedL2 systems respectively. Shown is the ratio of the average snoop response time with RegionScout over that without RegionScout. In general, the slower the bus the higher the benefits with RegionScout. With the exception of ocean and to a lesser extend cholesky and raytrace, there is little change in response time for the LocalL2 system. This suggests that the shared control bus is under-utilized and hence reducing bandwidth by itself would not lead to performance improvements. A much higher reduction in average response time is observed for the SharedL2 system. With a 12 processor cycle snoop processing latency a reduction of 15% or more is observed for cholesky, fft, lu, ocean, radix and raytrace. While we do not report these results, the reduction in response time in higher for the eight way SharedL2 system. This result suggests that snoop-based coherence coupled with RegionScout can be a viable low cost and high performance alternative for on-chip CMPs.

Point-to-Point Interconnect Model: Under the assumptions for the point-to-point interconnect a broadcast requires $N-1$ messages in an N node system (more would have been required for other interconnects such as a torus). Each link is capable of transferring 64-bits of data per message plus address and control information. Transferring a 32 byte block requires five messages at a minimum (one is for the request and the other four are for the data). With RegionScout it is not necessary to probe all other nodes for some requests. We count the number of messages sent during execution with and without RegionScout and report the ratio:

$$\frac{MessageCountWithRegionScout}{MessageCountWithoutRegionScout}$$

Figure 9(a) reports the message ratio for the point-to-point interconnect model for the same RegionScout filters as in Figure 8. We observe a reduction in messages in the range of 6% and up to 34%. The reduction is higher for the SharedL2 system (grey marks) for two reasons: Coherence messages are a higher percentage of all messages due to the smaller cache block size and the filter rates are higher as compared to the LocalL2 systems. In some cases, using larger regions results

in higher benefits since it allows the CRH to better separate among the fewer regions. For the configurations we studied using smaller regions does not result in significantly higher benefits even though we have seen that there are more global misses with smaller regions. This suggests that the CRH configurations we studied are incapable of separating among many small regions and better suited for larger regions.

For completeness we report per program message ratios for the four node SharedL2 system in Figure 9(b). Individual program behavior varies as the resulting behavior is a function of the filter rate and the fraction of messages used for coherence. Programs with a higher fraction of coherence messages are more sensitive to changes in the filter rate. We do not attempt an analytical model of response time since this interconnect does not behave as a simple queue.

6.4 Tag Energy Reduction

Previous work has shown that snoop induced tag lookups represent a significant fraction of cache energy [12,13,27]. We restrict our attention to 16Kbyte regions and the LocalL2 and SharedL2 systems with 512K L2 and 64K L1 caches. Since the RegionScout filters consume energy they represent an overhead which should be amortized by the benefits. Accordingly, we select small RegionScout filters. The NSRT has 16 entries and is direct mapped and the CRH has 256 entries. Using larger RegionScout filters resulted in lower energy benefits and in some cases in an increase in energy.

6.4.1 Comparing to Jetty. We first compare to the previously proposed JETTY filters [27]. Specifically, we simulated hybrid Jetty filters comprising an exclusive Jetty of 16 entries with 32 bit vectors and one out of four inclusive Jetty filters. The first three inclusive Jetty filters use three tables of 1K, 512 or 256 entries and we will use the acronyms 10x3, 9x3 and 8x3 to refer to them respectively. The fourth inclusive Jetty filter uses four tables of 128 entries each (acronym 7x4). We chose these four configurations after experimenting with several others. The specific sizes are in par with what was reported in the original Jetty study [27] except in that we use fewer sub-arrays and hence incur less energy overhead. We do so since in our simulation environment only the lower 32-bits of addresses are non-zero. For clarity we use a “ $pN.filter$ ” naming scheme where N is the number of nodes and *filter* is the snoop filter which can be a Jetty (10x3, 9x3, 8x3 or 7x4) or the aforementioned RegionScout configuration (RS). Figure 10(a) reports the average filter rate for SharedL2 configurations with two, four and eight nodes. The filter rate here is measured as a fraction of all snoop-induced tag lookups and is different from the filter rate reported in earlier sections (previously we were concerned with coherence requests where here we are interested in the tag lookups these induce in remote nodes). The grey bars correspond to Jetty filters while the white bars to RegionScout. It can be seen that the filter rate with RegionScout is comparable to the 8x3 Jetty. The latter utilizes at least three times as much resources. Figure 10(b) compares the various filters for the LocalL2 systems. Here RegionScout performs even better offering filtering rates that are higher than those possible even with the 9x3 Jetty which uses at least six times as much

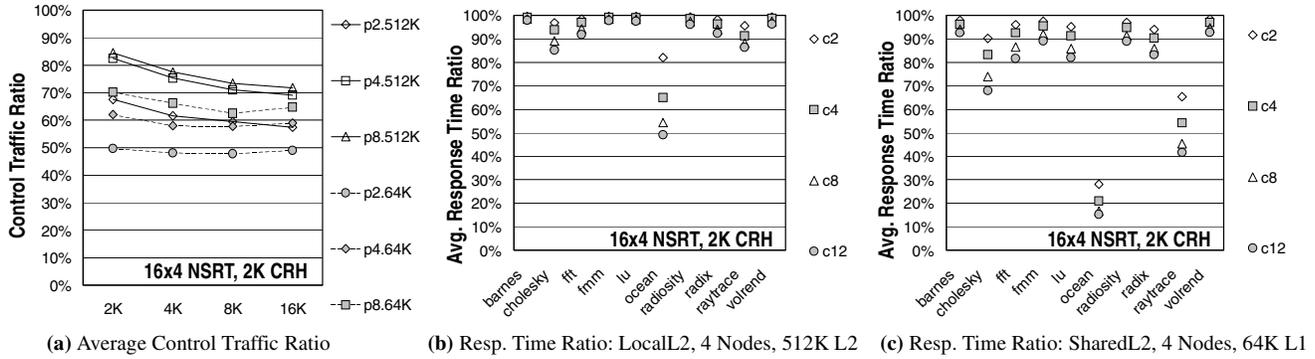


Figure 8: (a) Relative control traffic for the bus-based interconnect model. Reported is the ratio of the average control traffic for the system with RegionScout over the base system. (b) and (c): Estimated average snoop response time ratio with RegionScout over the base system for LocalL2 and SharedL2 respectively. We report results for various snoop processing latencies. Per snoop latency is modelled in N processor cycles (labels cN) where N is 2, 4, 8 or 12. The NSRT has 64-entries and is 4-way set-associative and the CRH has 2K entries.

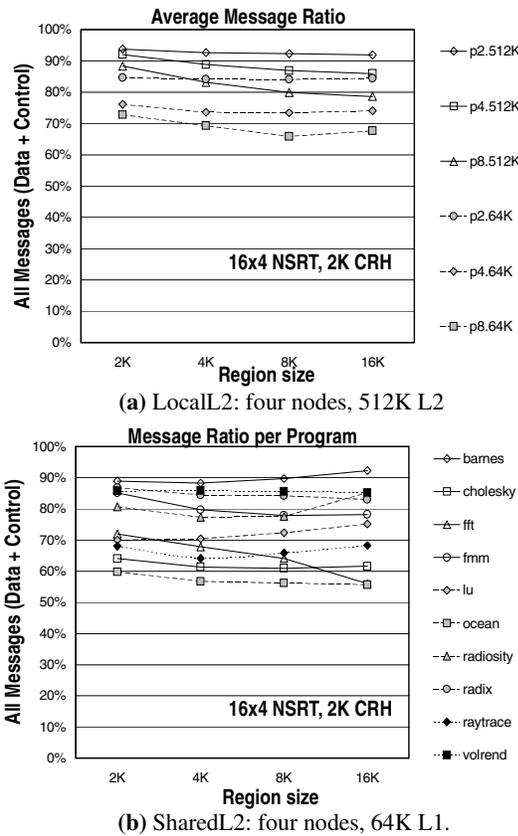


Figure 9: Avoiding broadcasts with RegionScout for various region sizes. NSRT is 64-entries, 4-way set-associative and the CRH has 2K entries. Shown is the ratio of messages sent with RegionScout over those sent without it. (a) Average message ratio for LocalL2 and SharedL2 systems with 512K L2 and 64K L1 caches respectively. (b) Message ratio per program for the SharedL2 system with 64K L1 caches.

resources than RegionScout. Due to the lack of space we do not report per program results noting that the trends are similar for all programs.

6.4.2 Energy Reduction. We measure the energy consumed during tag lookups in the L2 and L1 caches (all accesses, local or snoop induced) and report the energy ratio:

$$\text{EnergyRatio} = \frac{\text{TagEnergyWithRegionScout} + \text{RegionScoutEnergy}}{\text{TagEnergyWithoutRegionScout}}$$

To measure energy we use the WATTCH models [7]. The L1 and L2 caches were auto-partitioned using CACTI and modelled as SRAM arrays. We modelled the NSRT and the CRH as SRAM arrays also. The NSRT comprises a small SRAM while the CRH comprises two SRAMs, one for the present bits and one for the LFSR counters. We take into account the energy overhead of all RegionScout probes and updates. In the interest of space we not report energy results for Jetty. In most cases, the Jetty filters did not reduce energy for the SharedL2 systems (this corroborates earlier findings [13]). For the LocalL2 systems, Jetty filters did reduce energy but the reduction was smaller or comparable to that possible with RegionScout. Jetty utilizes much larger structures hence it incurs a much larger overhead than RegionScout.

Figure 11 reports the tag energy ratio per program and for systems with different number of nodes. The RegionScout filters energy overhead is included in the ratio and also reported separately. The resulting energy savings are primarily a function of (1) the fraction of tag accesses that are snoop-induced versus those that are generated locally, (2) the global region miss filter rate and (3) the filter rate of each CRH for those requests that are not identified as global region misses. In some cases the energy savings are higher than the filter rate for global region misses. This suggests that the CRHs filter some snoops that are not identified as global region misses by RegionScout. Most of these tag lookups are for requests that result in region misses in some but not all nodes (as per the discussion of section 5.3) and the others are for global region misses that are not identified by the requestor's NSRT. The LocalL2 systems benefit more from RegionScout since there snoop-induced tag lookups represent a higher fraction of overall tag accesses. This result corroborates previous findings [12,13]. RegionScout overhead is typically low.

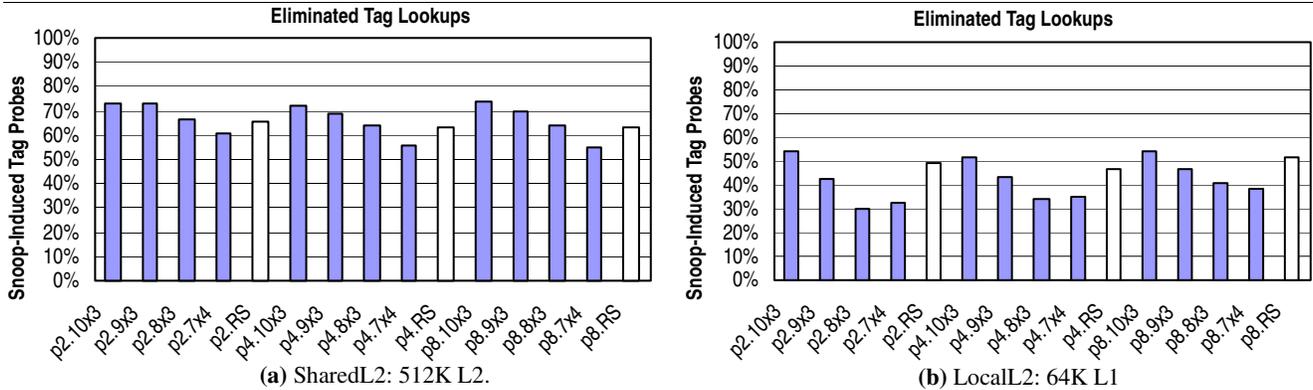


Figure 10: Snoop-induced tag lookups avoided with Jetty and RegionScout. The RegionScout filter used comprises a 256-entry CRH and a 16-entry 2-way set-associative NSRT and uses 16K regions. See text for a description of the four hybrid Jetty filters. We use a pN.filter naming scheme where N is the number of nodes and filter can be a Jetty (10x3, 9x3, 8x3 or 7x4) or RegionScout (RS).

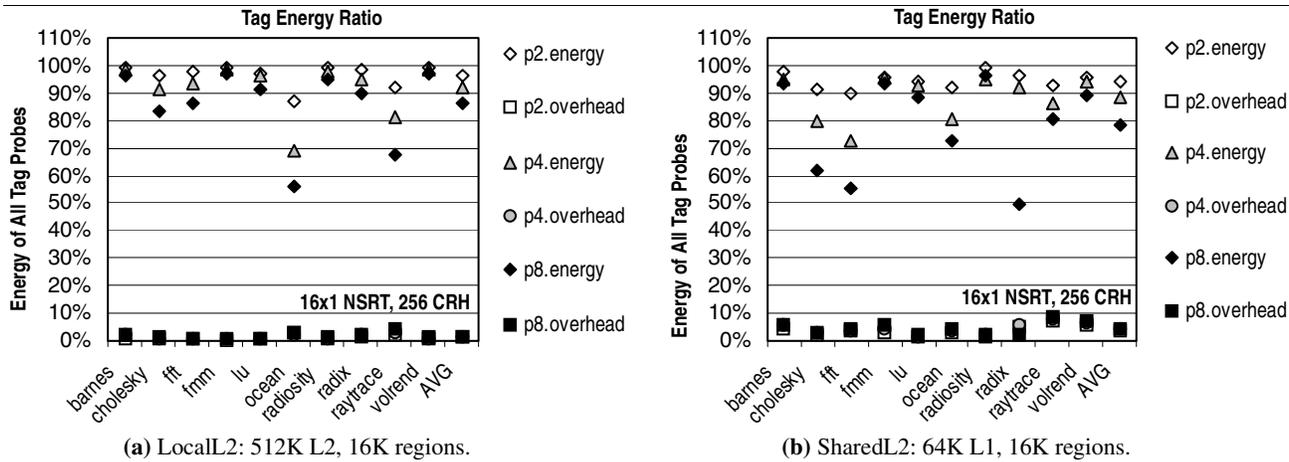


Figure 11: Tag energy ratio and RegionScout energy overhead. 16K Region, 16-entry direct-mapped NSRTs and 256-entry CRHs.

Compared to Jetty, RegionScout offers competitive energy reduction at a fraction of the cost. Moreover, RegionScout could also save energy on the interconnect. The reduction of messages reported in Section 6.3 serves as a first order approximation of the energy savings if we assume that energy consumption is proportional to the number of messages sent. We can also save energy in a bus-based CMP implementation by disabling the bus transceivers of remote nodes (an optimization not possible with Jetty). The overhead of an additional control signal would be negligible given typical bus signal counts.

6.5 Commercial Workloads

We finally, consider a few non-numerical applications and demonstrate that a large number of global region misses exists. For this purpose we use traces generated by the SimFlex full-system simulator [15]. The simulated system has 16 processors. Using these traces we simulated nodes with 512K local L2 caches to gather region miss statistics. All accesses are included including those from the operating system. We consider three workloads [31]: (1) SpecWEB99 comprising two web servers Apache and Zeus (2000 client connections), (2) SpecJBB2000 (Sun HotSpot JVM 1.4.2 with

16 warehouses and 16 clients), and (3) IBM's db2 (online transaction processing workload with 100 warehouses and 400 clients with zero think time). The traces include the first 100 million requests after initialization. As it can be seen in Figure 12 significant potential exists for all the workloads even with large regions.

7 Summary

We observed that many requests in shared memory multiprocessors not only do not hit in any remote node but also do not find any other block in a much larger surrounding region (global region miss). We proposed RegionScout a family of small and effective filters that can detect most of the requests that would result in a global region miss. RegionScout filters utilize imprecise information about the regions that are cached at each node in the form of hashed bit vector. This has several advantages as it requires only surgical additions over the existing infrastructure of conventional shared multiprocessors. This is especially important as the complexity of modern processors and the variability in behavior amongst different applications favors building systems out or pre-existing components using glue logic. RegionScout filters are transparent to software and do

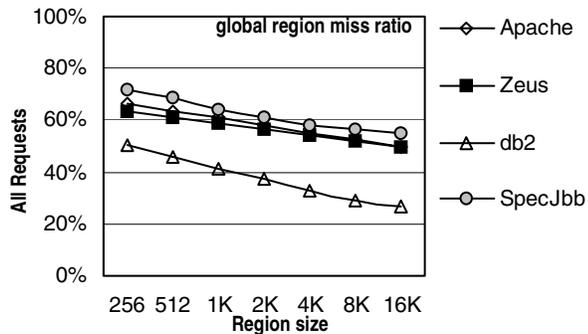


Figure 12: (a) Global Region Miss ratios for four commercial application workloads with the LocalL2 model and 512K L2 caches and 16 processors.

not impose any artificial limits on what can be cached and when. We have demonstrated that RegionScout filters can be used to reduce bandwidth by avoiding broadcasts for some requests. Moreover, we have shown that can reduce energy by avoiding some of the tag lookups for snoops that would miss. RegionScout filters are fundamentally different than previous snoop filters and directory-based coherence. We expect that region-level information tracking and imprecise information tracking will have applications beyond coherence.

Acknowledgements

I would like to thank the reviewers of this and previous conferences for their valuable comments. My understanding of multiprocessors and the ideas presented in this paper have benefited significantly from discussions with Angelos Bilas, Babak Falsafi and Dionisios Pneumatikatos. I would like to thank Babak Falsafi also for explaining how to develop the analytical performance estimation model and all the members of the CMU Impetus group for providing the commercial workload traces. I would also like to thank Mark Hill for pointing out the similarities between inclusive Jetty filters and Bloom filters. This work was supported by the Semiconductor Research Corporation under contract 901.001 and by an NSERC Discovery Grant.

References

- [1] —, MIPS R10000 Microprocessor User's Manual v2.0, MIPS Technologies, Inc., January 1997.
- [2] M. Abramovici, M. A. Breuer, A. D. Friedman, Digital Systems Testing & Testable Design, Wiley-IEEE Computer Society Press, January 1993.
- [3] P. Bannon, B. Lilly, D. Asher, M. Steinman, D. Webb, R. Tan, and T. Litt. Alpha 21364: A Single-Chip Shared Memory Multiprocessor, Government Microcircuits Applications Conference 2001, Digest of Papers, Defense Technical Information Center, Belvoir, Va., March 2001.
- [4] L. A. Barroso, K. Gharachorloo, R. McNamara, A. Nowatzyk, S. Qadeer, B. Sano, S. Smith, R. Stets, and B. Verghese. Piranha: A Scalable Architecture Based on Single-Chip Multiprocessing. In *Proc. of the 27th Annual International Symposium on Computer Architecture*, June 2000.
- [5] B. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of ACM*, pages 13(7):422-426, July 1970.
- [6] D. Burger and T. Austin. The SimpleScalar Tool Set v2.0, *Technical Report UW-CS-97-1342*. Computer Sciences Department, University of Wisconsin-Madison, June 1997.
- [7] D. Brooks, V. Tiwari M. Martonosi. Watch: A Framework for Architectural-Level Power Analysis and Optimization. In *Proc. of the 27th Annual International Symposium on Computer Architecture*, June 2000.
- [8] J. F. Cantin, M. H. Lipasti and J. E. Smith. Improving Multiprocessor Performance with Coarse-Grain Coherence Tracking. In *Proc. of the 32nd Annual International Symposium on Computer Architecture*, June 2005.
- [9] A. Charlesworth. Starfire: Extending the SMP Envelope. *IEEE Micro*, vol. 18, No 1, Jan/Feb 1998.
- [10] Z. Cvetanovic: Performance Analysis of the Alpha 21364-Based HP GS1280 Multiprocessor. In *Proc. of the 30th Annual International Symposium on Computer Architecture*, June 2003.

- [11] W. J. Dally and J. W. Poulton. Digital Systems Engineering. Cambridge University Press, 1998.
- [12] M. Ekman, F. Dahlgren, and P. Stenström. TLB and Snoop Energy-Reduction using Virtual Caches for Low-Power Chip-Multiprocessors. In *Proc. of ACM International Symposium on Low Power Electronics and Design*, August 2002.
- [13] M. Ekman, F. Dahlgren, and P. Stenström: Evaluation of Snoop-Energy Reduction Techniques for Chip-Multiprocessors. In *Proc. of the First Workshop on Duplicating, Deconstructing, and Debunking*, May 2002.
- [14] L. Hammond, B. Hubbert, M. Siu, M. Prabhu, M. Chen, and K. Olukotun. The Stanford Hydra CMP, *IEEE MICRO Magazine*, March-April 2000.
- [15] N. Hardavellas, S. Somogyi, T. F. Wenisch, R. E. Wunderlich, S. Chen, J. Kim, B. Falsafi, J. Hoe and A. G. Nowatzyk, SimFlex: A Fast, Accurate, Flexible Full-System Simulation Framework for Performance Evaluation of Server Architecture, *SIGMETRICS Performance Evaluation Review*, Vol. 31, No. 4, pp. 31-35, March 2004.
- [16] J. Huh, D. Burger, and S. W. Keckler. Exploring the design space of future CMPs. In *Proc. 10th International Conference on Parallel Architectures and Compilation Techniques*, September 2001.
- [17] S. Kaxiras and C. Young. Coherence Communication Prediction in Shared-Memory Multiprocessors. In *Proc. of the Sixth IEEE Symposium on High-Performance Computer Architecture*, Jan. 2000.
- [18] R. E. Kessler and R. Jooss and A. Lebeck and M. D. Hill. Inexpensive implementations of set-associativity. In *the Proceedings of the 16th Annual International Symposium on Computer Architecture*, 1989.
- [19] A.-C. Lai and B. Falsafi. Memory Sharing Predictor: The Key to a Speculative Coherent DSM. In *Proc. of the 26th Annual International Symposium on Computer Architecture*, May 1999.
- [20] J. Li, J. F. Martínèz and M. C. Huang, The Thrifty Barrier: Energy-Aware Synchronization in Shared-Memory Multiprocessors, In *Proc. of the 10th Annual International Symposium on High Performance Computer Architecture*, Feb. 2004.
- [21] N. Manjikian, Multiprocessor Enhancements of the SimpleScalar Tool Set, *ACM Computer Architecture News*, Vol. 29, No. 1, March 2001.
- [22] M. M. K. Martin, M. D. Hill, and D. A. Wood. Token Coherence: Decoupling Performance and Correctness. In *Proc. of the 30th Annual International Symposium on Computer Architecture*, June 2003.
- [23] M. M. K. Martin, P. J. Harper, D. J. Sorin, M. D. Hill, and D. A. Wood. Using Destination-Set Prediction to Improve the Latency/Bandwidth Tradeoff in Shared-Memory Multiprocessors. In *Proc. of the 30th Annual International Symposium on Computer Architecture*, June 2003.
- [24] M. M. K. Martin, D. J. Sorin, M. D. Hill, D. A. Wood, Bandwidth Adaptive Snooping, In *Proc. of the 8th International Symposium on High- Performance Computer Architecture*, January 2002.
- [25] J. M. Mellor-Crummey and M. L. Scott. Algorithms for Scalable Synchronization on Shared-Memory Multiprocessors, *ACM Transactions on Computer Systems*, February 1991.
- [26] A. Moshovos, Exploiting Coarse Grain Non-Shared Regions in Snoopy Coherent Multiprocessors, *Technical Report*, Computer Engineering Group, University of Toronto, Dec. 2003.
- [27] A. Moshovos, G. Memik, B. Falsafi, and A. Choudhary. Jetty: Filtering snoops for reduced energy consumption in SMP servers. In *Proc. of the 7th International Symposium on High- Performance Computer Architecture*, January 2001.
- [28] S. S. Mukherjee and M. D. Hill. Using Prediction to Accelerate Coherence Protocols. In *Proc. of the 25th Annual International Symposium on Computer Architecture*, June 1998.
- [29] J. Nilsson, A. Landin, P. Stenström. Coherence Predictor Cache: A Resource Efficient Coherence Message Prediction Infrastructure. In *Proc. of the 6th IEEE International Symposium on Parallel and Distributed Processing Symposium*, April 2003.
- [30] C. Saldanha and M. H. Lipasti, Power Efficient Cache Coherence, *High Performance Memory Systems*, edited by H. Hadimiouglu, D. Kaeli, J. Kuskin, A. Nanda, and J. Torrellas, Springer-Verlag, 2003.
- [31] S. Somogyi, T. F. Wenisch, N. Hardavellas, J. Kim, A. Ailamaki and B. Falsafi, Memory Coherence Activity Prediction in Commercial Workloads, *IEEE Workshop on Memory Performance Issues*, June 2004.
- [32] J. M. Tendler, J. S. Dodson, J. S. Fields, Jr., H. Le, and B. Sinharoy. POWER4 system microarchitecture. *IBM Journal of Research and Development*, vol. 46, No 1, January 2002.
- [33] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta. The SPLASH-2 programs: Characterization and methodological considerations. In *Proc. of the 22nd Annual International Symposium on Computer Architecture*, June 1995.
- [34] D. A. Wood and M. D. Hill. Cost-effective parallel computing. *IEEE Computer Magazine*, 28(2), Feb. 1995.