Hierarchical Extensions to Cache Coherence in SCI

Stefanos Kaxiras and James R. Goodman
Computer Sciences Department
University of Wisconsin-Madison
1210 W. Dayton St.
Madison, WI 53706
{goodman, kaxiras}@cs.wisc.edu

Abstract*—The cache coherence scheme of the Scalable Coherent Interface (SCI) offers performance for some operations degrading linearly with higher degrees of sharing. The currently proposed Kiloprocessor Extensions to SCI define tree protocols that create, maintain and invalidate binary trees without taking into account the underlying topology. We propose a different approach to Kiloprocessor Extensions that takes into account the underlying topology and offers the possibility of a simpler implementation. We give extended examples using* k*-ary n-cubes and Omega topologies. In order to implement new schemes we decouple data from directory information. In this way we can cache directory information on various nodes without the corresponding data. Regarding the distributed directory information, we define it in terms of a cache tag with multiple pointers. The actual implementation of such a tag can be in a form of a linked list (or other structure) so we can have dynamic pointer allocation. Set associative caches can be used for fast dynamic allocation and deallocation of pointers. We present several variants of the protocols without comparing them in terms of performance or cost.*

1 Motivation

The current [2] proposal for Kiloprocessor Extensions to Scalable Coherent Interface [1], which we will refer to as STEM, has achieved its goal of a logarithmic-time algorithm to build, maintain and invalidate a sharing binary tree without taking into account the topology of the interconnect network. The complexity of the algorithm however, is very high. It requires complex transactions that generate significant traffic even for modestly shared cache lines. Most significantly this traffic is non-local and consumes large portions of system bandwidth. The complexity of the algorithm impacts negatively on the marketability of SCI to high performance computing manufacturers, not to mention academic environments. Clearly if we want to do better, that is, attain comparable or even higher performance with considerably less complexity, a new approach is needed.

The starting point for STEM for a binary sharing tree is that an $O(\log_2 N)$ algorithm is needed for logarithmic time reading and writing shared data. However, this algorithm was proposed under the assumption that the latency of any message is unit latency (i.e. all messages have the same latency, regardless of the distance they travel). This is not true in a real system consisting of multiple rings and switches. When we take into account the increase of message latencies as a function of system size, it is not clear if the logarithmic nature of the algorithm can be translated into logarithmic performance, regardless of topology and system size.

Combining [3] has been proposed to reduce the bandwidth requirements of the STEM algorithm and possibly capture some geographical locality. Combining takes two requests that happen to be at the same time in a queue, and generates a single new request and a response to one of the original requests. Combining is not guaranteed to work with any load in the network. In fact it may only work in the presence of congestion in the network, i.e. with high loads when combining is mostly needed.

Large high performance systems are likely to have topologies where requests going towards a memory, and data distributed back to nodes, naturally form trees. We started from the assumption that in order to achieve high performance we need to be able to map any sharing tree on top of any reasonable topology, capturing the geographical locality available from the topology. By geographical locality we mean that nodes in physical proximity are neighbors in the tree, regardless of the timing of their requests. Capturing geographical locality means that neighbors in the logical sharing tree will also be physical neighbors. Sending a message from one node to the next (following a pointer), means that the message only travels a short distance with very low latency, consuming little bandwidth. Such a tree is impossible with STEM, because its structure is depend-

ent on the timing of the requests. In STEM we have temporal locality in the tree. Requesting nodes, in close proximity in time, are neighbors in the tree. In this work we present schemes that create trees that map well on the underlying topology. In our protocols messages are exchanged primarily between nodes on the same ring with very low latencies. Furthermore, we guarantee reduction of traffic, unlike the case where combining is used.

2 Hierarchical Extensions to SCI

We propose building a tree that closely matches the underlying topology. The tree is made of linear SCI-like sharing lists. Each such list is confined to one physical ring. Although the tree will not be optimal in the theoretical sense (some parts are linear) the operations involved (creation, rollout, invalidation) are going to be fast because messages are generally confined to a single ring.

Our approach in building such trees is a form of hierarchical caching. Nodes in the same ring will always send their requests for a given memory line, through the same switch node, which we call an *agent* node. If at this agent we have a copy of the line, we can satisfy these requests locally. The requesting nodes will form a small linear list. The head of this list will point to the agent, and not the home node of the line. The agent, along with other nodes in the next ring towards the home node, will also send their requests to their common agent and they will form their own small list. This process continues, until we reach the ring of the home node.

In order to be able to take full advantage of the wide variety of topologies that can be used to build SCI systems we propose *k*-ary logical sharing tree of lists (in contrast to STEM that specifies a binary tree). The degree *k* is dependent on the topology used. In general we specify a degree *k*, if a node is connected to *k+1* rings.

An example of a logical tree is shown in figure 1. The tree is drawn in such a way as to stress the lists as building blocks. In this figure memory is represented by the square in the top left corner, cache lines in nodes are represented by circles. The horizontal strings of cache lines are the linear lists. Several of the nodes in the above example have three children lists. A linear list and its parent are contained in one physical ring. The nodes that have child-lists, are the agents.

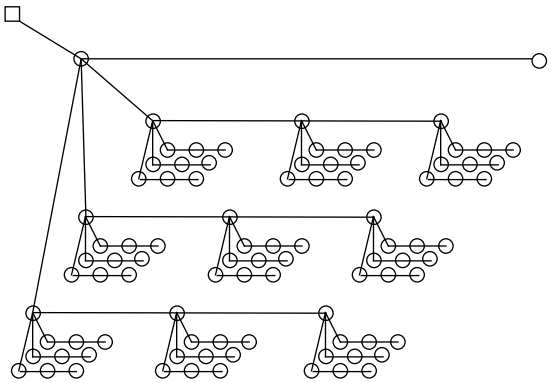

FIGURE 1. Logical *k*-ary tree of lists

We say that all nodes in a list belong to a level of hierarchy. The nodes of the list that is connected to memory, are in the highest level of hierarchy. The nodes that cannot have children are in the lowest level. The nodes of the children lists of a node in level *L* are in level *L-1* or lower. In figure 1, the tree has three levels of hierarchy and the lists in the lowest level are shown to have three members. In summary the main points of our proposal are:

- We build the logical *k*-ary tree having as a blueprint the underlying topology
- The *k*-ary tree is made of linear lists. A linear list is confined inside a single ring.
- The cache lines that are members of the same linear list have a common agent on the ring that by default routes requests for the cached address to other rings.
- Such agents have a cached line that is part of the tree and it can point to a number of child lists.

We are also investigating a more radical departure from the base SCI protocol that offers higher performance. The scheme depends on the observation that entire lists created under the proposed scheme are generally contained within a single ring. This fact, and the expectation that a modest number of nodes will be attached to a single ring, suggests the possibility of storing directory information related to a single list not in the form of doubly-linked list, but as a bitmap. Storing the information as a bitmap opens up a range of alternative algorithms, many using broadcast, and in particular, permits a pruning cache implementation in the style of that proposed by Scott and Goodman [9]. We note that the SCI transport layer provides a limited broadcast mechanism that is not used by the coherence layer.

3 Examples

In trying to come up with examples of topologies we skipped the obvious topologies where our scheme would trivially fit. These include trees, snowflakes, dense snowflakes and star topologies [4], the lens [5], and various other topologies. Although these are interesting topologies, we think that they are not mainstream and some of them also give an "unfair" advantage to our proposals by being so hierarchical in nature.

To illustrate our approach we chose two topologies we believe are likely candidates for implementation The first is *k*-ary *n*-cube topology made of rings, suggested in the SCI standard (p.12). We introduce the second topology, which we call Omega/Flip. This topology is a refinement of the butterfly topologies suggested also in the SCI standard (p.11). We believe that Omega/Flip topologies are a good fit for SCI but this is an appropriate discussion for a different document.

We found that our schemes compare favorably to STEM, while being much less complex. In several cases, our schemes in the two example topologies outperform STEM. More details about performance will be given in the appropriate sections.

3.1 *K*-ary *n*-cube

In this example the routing of requests is fixed in order to have deadlock free routing [6]. For example, in a *k*-ary 3-cube the request messages are routed first along *x*, then *y* and finally *z* dimensions. At the place where requests switch to get into the next dimension they are intercepted. A new request is generated and continues in the next dimension. The response to an intercepted request travels back along the same path the request used.

In figure 2 we show a 4-ary 3-cube which is equivalent to a 3-D Torus. This topology has four nodes per ring and three rings per node. The total number of nodes is 64. Without loss of generality we assume that all the nodes access a line in some node, labelled "Home node."

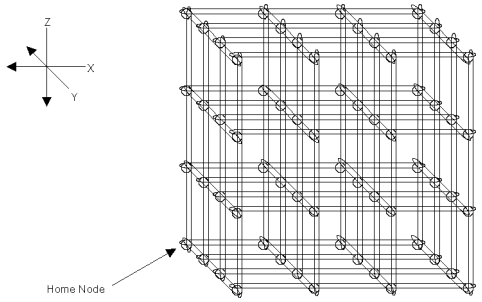

FIGURE 2. 4-ary 3-cube

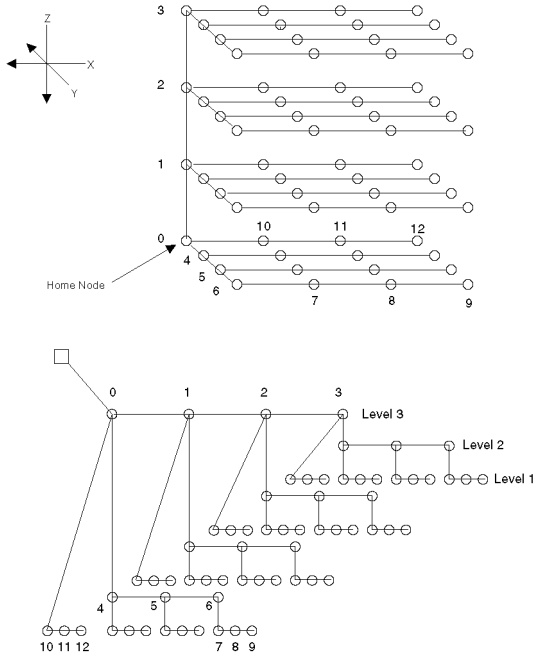

FIGURE 3. Tree on 4-ary 3-cube

In the sharing trees of this topology, not all nodes have the same number of child lists. The higher the hierarchical level of a node, the more child lists it has. For this example, nodes in the lowest level have none, in the next level they have one, and in the top level they have two. In general, in a *k*-ary *n*-cube, the logical tree will have *n* hierarchical levels. At the highest level (the list that is directly connected to memory) the nodes will have *n-1* child-lists. The child-lists do not necessarily belong to the same hierarchical levels. Moving one level

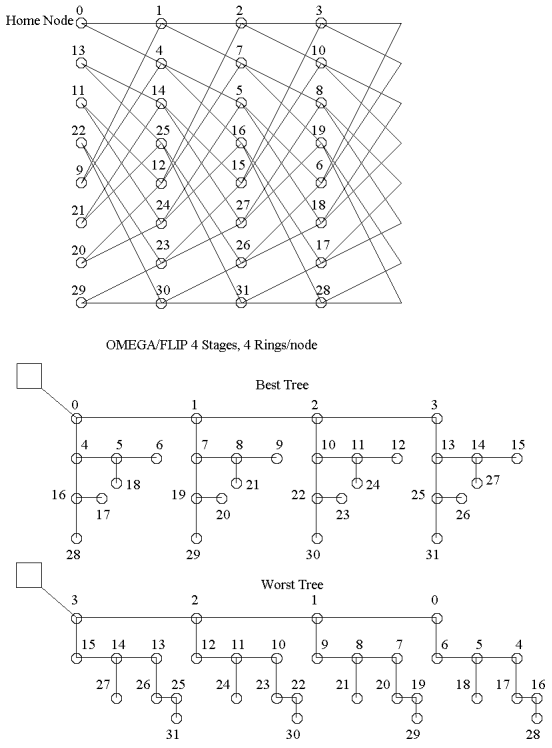

FIGURE 4. Omega–best and worst trees.

down, the number of child-lists decreases by one, until the lowest level, where the nodes do not have any child-lists.

In figure 3 we show how a *k*-ary tree is mapped on the 4-ary 3-cube. In the second representation of the tree, in figure 3, we draw it as a tree of lists. We have numbered a few nodes to correlate the two representations.

3.2 OMEGA/FLIP Topologies

In his thesis Ross Johnson suggested building butterfly topologies with rings [2]. A disadvantage of the topologies he suggested, is that the rings have a varying number of nodes. We propose two variations of these topologies, the Omega [7] and the Flip. The important characteristic of these two topologies is that they are built of rings of the same size. The ring-based Omega topology, has *R* rings per node, *S* stages, and *N* nodes per stage. The nodes of last stage are connected to the nodes of the first stage, making the topology a cylinder. The rings run once around the cylinder, and they contain *S* nodes. The Flip is the reverse of the Omega: a counter-clockwise Omega is a clockwise Flip. The Omega/Flip is made of one Omega and one Flip network superimposed (bidirectional Omega). The relations that define the Omega topologies we suggest are:

$$TotalSystemSize = S \times N = S \times R^{S-1}$$

$$S = log_R(N)+1$$

$$R^{S-1} = N$$

T*otalSystemSize* is the total number of nodes. The independent variables are *S* and *R*. Variable *N* is constrained by the above relations to have specific integer values.

Figure 4 shows a 4-stage, 4 ring per node (2 for the Omega and 2 for the Flip), Omega/Flip network. The structure of the our trees in this topology is depicted in the two trees below the network. We have numbered the nodes in the topology and the trees to show the mapping. These two trees are the worst and best case, in terms of height (longest path). In Omega/Flip networks the trees have as many hierarchical levels as number of stages.

4 Tag Storage

Two kinds of information are contained in the tag information of a cache memory. First is the data necessary to identify uniquely the address of the data in the cache, and its validity. This information is required for any cache, even a uniprocessor. Second is data that is related to maintaining the consistency of the data, which we will refer to generically as “directory information.” In a snooping cache, this directory information may be as little as the state of the line, though a separate tag memory is often maintained to permit simultaneous access from the processor and the bus. In a more general, directory-based scheme, this directory information also includes indications about what other nodes in the system have copies of the data. The basic SCI protocol, for example, includes forward and backward pointers, indicating up to two other nodes that have copies of the line. Though this data is usually associated with a cache, it is in fact directory information. The concept of distributed directory information is contained both in snooping protocols and in SCI. In both cases, critical information about where copies of the line reside are distributed along with the cached copies.

In general, directory information can be distributed around the system, including nodes where it is neither cached nor present in memory. In particular, it is not necessary that a valid copy of the cached data be present at a node containing directory information. At such a node, the problem then is to store the information in a way that it can be accessed and updated as needed. Depending on the scheme used, it may be desirable that

the amount of directory information stored vary over time. The problem of dynamically allocating space for directory information was nicely solved by Simoni and Horowitz [8] for a conventional directory scheme, where space is allocated at a centralized location. We propose that such techniques might also be applied at other, critical points in the system, where directory information might be strategically placed to minimize the distance necessary to acquire the information. Thus the problem is one of allocating storage for directory information dynamically, and independent of cached copies, which might or might not be present.

A plethora of techniques might be used to store directory information. Among those techniques are to maintain a linked list, as suggested by Simoni and Horowitz. Another technique is to store the data associatively, using conventional cache techniques of set-associativity, but eliminating the cache data. Finally, it may be useful to store information about nearby nodes that contain data using boolean variables to indicate which nodes contain data. Techniques similar to those proposed for pruning caches [9] can then be applied to store the critical directory information.

In this section we discuss various aspects of tag storage. We first present a simple implementation of the tags needed for *k*-ary trees. We then discuss how, by using multiple types of tags, we decouple the data and the corresponding directory information of the tag. We continue with alternative implementations of tags with multiple pointers in the form of linked lists or trees and finally we examine how, by taking into account geographical locality, we can minimize storage requirements with *short pointers* and bit maps.

4.1 Single Tag with Multiple Pointers

We defined previously, the degree *k* of the *k*-ary trees to be one less than the number of rings *R*, connected to a node. Since each node belongs to a list in a hierarchical level and can have up to *R - 1* child lists, we require at minimum *R + 1* pointers: a *forwid*, a *backid*, and *R - 1* *childid* pointers.

A simple minded implementation can have *R + 2* pointers. Each of the R-1 child pointers handles the child list inside the corresponding ring. Since both *forwid* and *backid* will be used to connect to the high-level list in one ring, the corresponding child pointer for this ring will be left unused. Using R+2 pointers, we avoid doing any dynamic binding between them and the rings. We can distinguish messages that refer to *forwid* and *backid* and we know which *childid* pointer we will use, by knowing the ring that carried the message. Since R will be a small number, the extra pointer is a significant overhead. A more cost effective solution is described in [13].

4.2 Multiple Types of Tags

It is clear from the two examples given in previous sections that in hierarchical trees, an agent node might be in the tree without the agent's processor accessing the cache line. The agent has a cached copy of the line just for the benefit of its children. This may create many conflict misses because local processor accesses might conflict with the tree cache line. This situation would result in many tree rollouts for the agent nodes. Tree rollouts are generally very expensive and we would like to avoid them. To avoid this we allow tags to be cached in agents without the corresponding data. Thus we are able to have pointers in an agent without competing with the agent's processor for data storage. In order to cover the possible cases, we define four types of tags summarized in table 1.

TABLE 1.

Types	Description	Pointers	Data
Type SCI	Standard SCI	2	Yes; accessed locally
Type Tree No Data (TND)	Tree	multiple	No
Type Tree Optional Data (TOD)	Tree	multiple	Yes; can be discarded
Type Tree Data (TD)	Tree	multiple	Yes; accessed locally

When a local access results in choosing a victim line with its associated tag we take different action depending on the type of the tag. Table 3 summarizes the actions in this situation. The goal is to keep tree tags (with multiple pointers) cached in an agent despite a conflict for the data line. When we eventually have to choose a tree tag to replace, because of the tag storage replacement algorithm, we will have a tree rollout as described in a later section. Notice that because a Type TND tag does not have an associated data line it is never chosen for replacement due to local accesses. It can only be chosen by the replacement algorithm for the tag stor-

age. Additionally we do not allow tags of Type TND and Type TOD to exist without any children.

TABLE 2.

	Data line Victim	Action
Type SCI		Replace (SCI rollout)
Type TND	Never	None
Type TOD		Convert to Type TND
Type TD		Convert to Type TND

Finally notice that when we use multiple types of tags there is not 1-to-1 correspondence between the tags and the data lines.

4.3 Multiple Tags

We have describe tags with multiple pointers to implement *k*-ary trees. However, we can assemble the required number of pointers by associating many tags (each one with a number of pointers). This idea is described by R. Simoni and M. Horowitz [1]. The multiple tag entries (instead of one with many pointers) can be linked together and form a linear list. Alternatively they can be unrelated (not linked) in which case they can be searched and found individually. Additional details are available elsewhere [13].

4.3.1 Multiple Linked Tag Entries

There is a range of implementations with varying number of pointers per tag entry. One end of the spectrum is to have all necessary pointers to implement *k*-ary trees in only one tag. Moving away from this end, we can have tag entries with a smaller number of pointers so we would need to have two or more of them, linked together, to implement *k*-ary trees. Another option is to have tags with a different number of pointers in a linked list. For example the first tag entry can have three pointers and the rest only two. Two interesting points in this spectrum are the case with three pointers per entry, which is a STEM requirement, and the case of just two pointers per entry, which is a base SCI requirement. Three pointers per tag entry are enough for a doubly linked list of tags where each tag has a child pointer. Two pointers per tag entry are sufficient for a singly linked list of tags where again each tag has a child pointer. Of course we are not constrained in linked lists and we can easily implement the tags as trees. In order to access all the tags we have to find only the first tag entry and then follow the pointers. This means that the SCI address has to be kept only for the first tag entry in order to be matched. Traversing these internal lists is not expected to be a performance problem because a) they are internal to a node, b) the maximum length is not great (the constant k is usually small) and c) their average length is usually very small.

4.3.2 Multiple Tag Entries

Multiple tag entries can also be used without linking them together as long as we can access all of them. This means that the SCI address has to be kept for each individual tag entry for matching purposes. Again there is a spectrum of implementations depending on the number of pointers per tag entry. Standard SCI tags with two pointers can be used. One of the tags would implement a *forwid* and *backid* pointers to connect to the high level list and each of the other tags can implement one or two child pointers. We can also use just one pointer per entry and allocate enough entries for a *forwid* pointer a *backid* pointer and the required number of child pointers. This in effect is equivalent to providing multiple memory directory entries in an agent that represents the remote memory to its children. To implement such a scheme with multiple unlinked tags we can use a set associative cache with associativity equal to or greater than the maximum number of tag entries needed to implement *k*-ary trees. For example, to implement 4-ary trees with standard SCI tags (unlinked), at minimum three tags are needed: one to implement the *forwid* and *backid* pointers and two for the four child pointers. A 3-way (or greater) set associative cache could then be used. A set of this cache can hold three tag entries for the same address or three tags corresponding to different addresses (or any combination in between). A disadvantage of using unlinked tags is that each one must carry the SCI address so it can be matched and accessed independently.

4.3.3 Linked Tags with Address

We can also combine the previous two methods so we can have a linked list of tags where each tag carries the SCI address. In this way we can use a set associative cache to implement the linked list. The associativity of the cache must be equal to or greater than the number of tags we need to implement *k*-ary trees. The set associative cache offers transparent buffer management, automatically allocating and deleting tags according to demand. Interesting cases include tags with three and tags with two pointers as described previously in section 4.3.1. Having the list linked helps in the rollout algorithm as discussed in a subsequent section.

In general multiple tags (either linked or unlinked) offer the opportunity to allocate pointer storage on demand instead of always allocating pointer storage for the maximum number of pointers needed. Notice here that in

hierarchical trees only a few nodes usually need the maximum number of pointers. Usually, nodes in the lowest level of hierarchy do not have any children and as we go up in the hierarchy the number of children increases up to the maximum in the highest level.

The multiple types of tags (Type SCI, Type TND, Type TOD, Type TD) described previously and the multiple tags in an agent are largely orthogonal. Using multiple tags (linked or unlinked) to achieve the same effect as a tag with many pointers is orthogonal to whether the data is also present in an agent.

4.4 Short Pointers

In our *k*-ary trees, geographical locality plays an important role. If all the nodes are in their proper position in the tree, we know that pointers do not point very far. In fact a pointer associated with a ring, will only point to nodes on that ring! It is therefore possible, for every node to have aliases for the other nodes on its rings and point to them using their alias id. Such a pointer would require much less storage than a system-wide pointer. We call the pointers, that are the alias ids, *short pointers*, while we call the system-wide pointers, *full pointers*. Additional details are available elsewhere [13].

The only extra hardware needed to support short pointers, is the Alias Tables. For each ring connected to a node, we need an Alias Table. Each table must have as many entries as the nodes on the corresponding ring. Note, there are only a few such tables per node (one per ring) whereas there are many pointers per cache line and many cache lines.

4.5 Agents

The agents play a central role in our schemes. They intercept other nodes' requests and cache tree tags and possibly data. In order to use a standard SCI switch in our schemes we require some additional functionality. Specifically the agent (switch node) must be able to:

1. Intercept special requests used for the creation of the trees (see section 5). The agent must recognize these requests by examining the command field (second symbol) of an SCI request. Instead of routing these requests normally, switching them from ring to ring, the agent delivers them to the local SCI controller.
2. Be able to annotate an intercepted special request with the ring id it came from. The agent delivers the special request to the local SCI controller and also supplies the ring id of the request. In this way the local SCI controller will be able to choose the appropriate pointers of a tree tag to manipulate.

5 Protocols

Current trends in research for distributed shared-memory parallel machines, favor the use of specialized protocols for different classes of shared data [11]. Several classes of shared data have being identified, including migratory, read-only, etc. [11,12]. This approach to shared memory assumes that the programmer and/or the compiler can identify the type of a datum. Subsequently, the underlying hardware or software layer is instructed to apply the proper protocol for the datum. This approach can provide greater performance and we believe that we should have different kinds of requests that specify construction of linear sharing lists, construction of sharing trees, construction of sparse sharing trees (section 5.1.2), pairwise sharing, and QOLB. We note that STEM already has implicitly assumed that a node must ask for a non-modifiable cache copy for combining to work.

5.1 Construction

In this section, we describe the construction of the *k*-ary tree. We briefly outline a method to construct a sparse tree when there is not enough sharing to justify the overhead of the full tree.

A node that wishes to read a remote line, that is expected to be widely shared, sends a special *mread_tree* request towards the node that has the data in its local memory (home node). The request is intercepted at the first agent where it is taken off the ring. The intercept causes a local lookup to find information about the requested line. If the lookup results in a hit, the requesting node is instructed to prepend to the appropriate child list. The node will either get the data from the previous head of the child-list or the agent. If the lookup results in a miss, the agent sends its own request for the line towards the home node. The agent does not forward intercepted requests but rather it replaces all of them by its own request. The requesting node is instructed to prepend to the appropriate child list. As soon as the interceptor gets a copy of the line it will pass it to the waiting node.

In the time it takes for an agent to get its own copy of the data, new requests may arrive. Each of the new requesting nodes is instructed to prepend to the appropriate list and wait for the data. We have three options of how to distribute the data to a list where all nodes are waiting for the agent to get the line:

1. As is currently defined in SCI. A node asks the agent (memory in SCI) for the data. The agent (or memory in SCI) instructs the node to prepend to the child-list

and ask the data from the old head. The very first node that requested the data is the waiting tail of the child-list. It will eventually get the data from the agent. The agent may keep a temporary pointer to the first requesting node (the storage for the data can be used for temporary pointers).

2. If a requesting node prepends to a child list, where the old head is waiting for the data, then it assumes responsibility to forward the data. When the agent gets the data it will pass it to the waiting heads of its child-lists. The heads will then forward the data towards the waiting tails. This distribution happens in the opposite direction of the one described above (standard SCI).
3. Broadcast the line on the ring. This is a broadcast confined in a single ring. All nodes that are waiting will receive the broadcast and consequently the data. Here we may need to define a confirmation protocol to make sure that all waiting nodes actually got the data.

In the case of a miss in the agent, its own request will be treated the same way in the next agent. Eventually the agent just before the home node will get the data from memory and it will pass it to its child lists.

Construction of the tree when all nodes simultaneously read a line, after a global barrier, is quite fast. All the linear lists can be constructed in parallel. Copies of the cache line are distributed down the tree concurrently with list construction.

5.1.1 Construction with Multiple Types of Tags

In this section we describe in more detail the construction taking into account the four types of tags. We remind here that the types of tags we have defined: Type SCI, Type TND, Type TOD, Type TD. Additional details are available elsewhere [13].

When a node requests a line and its request is intercepted in an agent we can have either a miss or a hit. In case of a miss the agent can allocate either a Type TND or a Type TOD tag. When the data returns the agent will forward it to the waiting tails of its newly created child-lists. If the allocated tree tag was of Type TOD the data is also stored in the agent to service subsequent requests more efficiently.

In case of a hit the tag in the agent can be of any of the four types. The agent services the request differently depending on the type of the tag. The four cases are described below:

1. Hit on Type SCI tag. In this case the agent has to convert the SCI tag to a tree tag. Since the data is used locally the tree tag will become Type TD. The agent replies immediately to any request by putting the requesting node head of the appropriate child-list and passing it the data.
2. Hit on Type TND tag. Because the tag is already a tree tag there is no need to allocate a new one. However the data are not local so the requests cannot be satisfied with a single response. Since Type TND tags are not allowed to exist without children there is at least one child-list where the data can be found. The agent instructs the requesting node to prepend to the appropriate child list. If the requesting node prepends to an existing child-list then it gets the data from the old head of the child-list. If the requesting node becomes the first head on a new child-list the agent will acquire the data from one of its other child-lists and supply it to the waiting head.
3. Hit on Type TOD. In this case the requesting node is made head of the appropriate child-list and it is given the data in a single response.
4. Hit on Type TD. The same as a Type TOD hit.

5.1.2 Construction of Sparse Trees

Sparse trees are trees constructed without invoking every agent on the way to memory thus minimizing the number of superfluous nodes in the tree. Restructuring of the trees is very complex and we do not have a solution to build the tree top-down. We do have, however, a proposal to create the tree slowly, bottom-up.

For sparse tree construction it is not necessary to invoke all the agent nodes on the way to memory. Instead a requesting node will invoke only the agent at the next hierarchical level. If this node accepts a number of additional requests, it will invoke the next agent and so on. When an agent is first invoked by one of its children it will send a new *mread* request to memory. This request, in contrast to the *mread_tree* request, does not generate lookups in other agents. When an agent decides to invoke its own agent it connects to it and forces it to join the tree if it has not already. After the new agent joins the tree the requesting agent disconnects from the list it is in (with an SCI-like rollout) and rejoins in its correct position under its agent. By changing only the *forwid* and *backid* and never the child pointers of the nodes, the tree restructures and eventually it will assume the structure of a full tree.

While constructing sparse trees it is possible for an agent to have a copy of the requested line, but nevertheless ignore a passing request. For this case, we propose a

'hint' cache. The hint cache would reside in the SCI linc hardware (not in the SCI cache controller hardware) and would nudge the node to do a lookup if the address is believed to have been seen recently in an *mread* message.

5.2 Rollout

Rollouts are a concern in tree protocols because they require a replacement in order to keep the structure of the tree intact. In our trees rollouts can be more of a problem because lines may be in the tree structure just for the benefit of other nodes, without being actively accessed by the local processor. In hierarchical caching schemes a property that is sometimes useful is multilevel inclusion. This means, that the contents of a cache are the superset of all the contents, in the *mread* category, of all the caches it serves. Multilevel inclusion is desirable (but not enforced) in our case only for the cache lines that are in tree structures.

In the simple cases where a node does not have children the rollout is the same as in the standard SCI. A childless node is only connected to a linear list with its *forwid* and *backid*. A childless node always has the data cached; otherwise it would not be in the tree. Remember that childless Type TND and Type TOD are not allowed to exist. Whenever all the children of a Type TND or a Type TOD tag roll out the tag is also required to roll out.

However, when using multiple types of tags, agents may not have the data although they take part in the tree. If all nodes that do cache the data roll out the last valid copy of the line might be lost. In order to preserve a valid copy we require that the data be written back by the rolling out nodes when they are alone in their list. This is the same behavior as in the standard SCI where the last node of the linear list writes the data back to memory when it rolls out. For our case when any node rolls out and it is the only node in a child-list, it is required to write the data back to the agent. If the agent has other children or if it is head or in the middle in its list, it ignores the write. When a node has children, it is certain that one of its descendents has the data, because childless Type TND or Type TOD tags are not allowed. On the other hand, if the agent is head or in the middle in its list, then another node in the list will be responsible for the data, specifically the one which will eventually be alone in the list. However if the agent is alone in its list and it does not have any other children but the one that rolls out, it accepts the write. If the tag in the agent is a Type TND or a Type TOD, it will write the data back to its own agent and immediately roll out.

In this section we propose three approaches to handle rollouts. The first is very simple but destroys parts of the tree, the second is similar to STEM and the third is to patch the tree locally. All schemes pose problems, and this is a topic of ongoing research. Additional details are available elsewhere [13].

5.2.1 Simple Rollout

The simple solution, borrowed from hierarchical caches, is to maintain multilevel inclusion [10]. When a node rolls out it invalidates all of its descendents. In this case we should try to reduce conflict misses in the caches to a minimum. Tags of Type TND and Type TOD help minimize conflict misses. Notice here, that by using four types of tags we maintain multilevel inclusion only for the tags and not the data.

The above simple scheme has a great appeal because of its simplicity. If we can really keep the conflict misses to a minimum then this might be one the best ways to handle rollouts. One disadvantage is that this rollout method is especially prone to thrashing behavior.

5.2.2 STEM-like Rollout

A different method to handle rollouts is similar to STEM. When a node with children rolls out we search among its descendents for a childless node and swap the two nodes. The STEM-like rollout protocol in more detail is defined as follows. If the node does not have children the rollout is the same as in normal SCI lists. If the node does have children it must find a childless node, swap places with it, and then roll out. As in STEM, a node has to walk down the tree, following pointers in order to find a replacement node.

When a childless node is found it can be reserved for the subsequent swap. This will protect against problems with concurrent deletions. When the result of the search is reported back to the originating node, it starts the swap process with the replacement node. It instructs the replacement node to roll out and patch the list it is in. It then sends in one message all of its pointers to the replacement node. Finally it notifies its siblings (or parent and sibling) to point to the replacement node and rolls out.

A node that rolls out has to search for a replacement node among its descendents. If a descendent happens to roll out at the same time, it has priority to choose a replacement, over its own set of descendents (as in STEM). If a parent node and one of its children roll out simultaneously, the search of the parent node will result in negative answer from the child. The replacement node must be found in the rest of the subtree of its other

children. If a node gets a negative answer from all its children (i.e. all of them roll out) it has to retry later. Forward progress can be guaranteed with retries, since all its children will roll out and the node will eventually succeed.

While, for the STEM tree, rollouts do not affect the performance of subsequent operations (invalidation, write-updates) this is not true for our proposal. Geographical locality in the STEM tree remains as poor as it was. In our tree it is degraded from the optimal case (where pointers point to nodes in the same ring). Furthermore a replacement node cannot act as an agent for the children and would-be children of the replaced node. It can only take its place but it cannot assume its role. Subsequent read requests from would-be children, will not get intercepted in the replacement node but rather they will invoke again the replaced node in the head of the replacement node's list. As the number of rollouts is increasing, and as time goes by, both locality and structure of the tree are degraded. This is expected to happen to an excessive degree only in the cases where the sharing tree exists for a long time. It is, therefore, acceptable for the invalidation of the tree or subsequent write-updates to be slower.

5.2.3 Local Patching of the Tree

Here we present rollout protocols that are not based on finding a childless node to swap places with. Instead they try to patch up the tree *locally,* manipulating only the tree structure around the agent that rolls out. The protocols described below give priority to maintaining geographical locality instead of preserving the structure of the tree.

In the first protocol the agent that rolls out instructs the heads of its child-lists to chain together in a list. It then connects the chain of the heads of the child-lists in its place. In effect it raises the heads of the child-lists to its own hierarchical level. This method preserves the geographical locality of the tree (as possible), over the quality of its structure. The tree is degraded because each node that rolls out is replaced by a linear list of the heads of its child-lists.

Another way to chain the children together is to concatenate all the child-lists together. The agent sends messages to the heads of the child-lists that are distributed towards the tails. The child-lists are chained head to tail. The resulting linear list is then raised to the position of the agent.

In the rollout protocols we have described so far (including STEM-like rollout) there is the problem of an agent rolling out and later rejoining the tree. In these situations the relation of the agent and its children is lost at the point of the rollout. If the agent rejoins the tree it will do so in another position in the linear list of the appropriate hierarchical level. New requesting nodes that would otherwise become heads in the old child-lists, will create new child-lists under the agent.

5.2.4 Bitmaps and Rollout

The bitmap scheme mentioned earlier could be used to store in an agent all the siblings for each of its child lists. The pruning cache concept is based on the observation that if the bitmap information is lost, the tree can be reconstructed conservatively, that is, by including all possible siblings in the child list. While nodes not in the tree may therefore be inadvertently inserted into the tree, they need not contain data, so no harm is created by adding them (except that the tree is unnecessarily larger than it need be, and extra traffic may be generated in destroying it). This approach suggests a fifth type of tag: Absent. This virtual tag is not stored at all, but is implicit in situations where a node is informed by its parent that it is in a list.

One of the benefits of the bitmap scheme is that a node can be rolled out for free. If it is simply deleted, it automatically becomes a Type Absent (virtual) tag, and nothing further need be done. Note that this feature meets a basic requirement of the bitmapped scheme: all lists must be contained within a single ring. This scheme therefore always forms the perfect tree, and avoids the problems of patching ill-formed trees.

5.3 Destruction (Invalidation)

Destruction (invalidation) of the tree is defined similarly to STEM. When a node wants to write a line, it first becomes the head of the top level list. Invalidation messages are then forwarded down the tree in parallel. At each stage a node can accept the invalidation message (with a response to the node that sent it) assuming the responsibility to forward it. Alternatively, if the node is childless, it can return its *forwid* pointer to the sender. In the latter case the sender repeats the invalidation message to the node pointed by the returned *forwid* (as in the base SCI protocol). When the invalidation messages reach the tails of the lists, acknowledgment messages are sent the reverse way. The tails invalidate themselves and send their acknowledgment back. Nodes wait for the acknowledgment of all the invalidation messages they forwarded, invalidate themselves, and return their own acknowledgment.. See reference [13] for further details.

Invalidation is fast because all the messages are exchanged locally, between nodes within the same ring. Even if the longest path of our tree is larger than that of the STEM tree, invalidation in many cases is faster because of the low message latency.

We note here that confining lists to a single ring potentially offers very significant reductions both in message traffic and destruction latency for a scheme that can exploit broadcast effectively. The bitmap scheme, in particular, can invalidate an entire tree very efficiently by employing recursive broadcast, once per ring.

5.4 StoreUpdate

5.4.1 StoreUpdate by the Head of Tree

StoreUpdate in the tree, is defined similarly to STEM. A node wishing to update all other entries in the sharing tree, leaves the tree (performing a rollout as described in the previous section). It then prepends to the top level list. Being the head, directly connected to memory, it can initiate a StoreUpdate. The StoreUpdate happens in two phases. The first is the initiation phase where the new data propagate throughout the tree in the same way as the invalidation messages. A node receiving a StoreUpdate from its parent or left sibling forwards it in parallel to all its children and right sibling. Alternatively a childless node can simply return its *forwid* to its parent, giving it the responsibility to forward the StoreUpdate to the next node. When the data reach the tails of the lowest level lists, the tails return confirmation of their update. Nodes in the tree wait for confirmation from all their children and right sibling, before they return their own confirmation messages to their parent.

5.4.2 StoreUpdate with Bimodal Behavior

A method of doing StoreUpdate that is currently discussed in Kiloprocessor Extensions to SCI working group is for a node to have bimodal (or schizophrenic) behavior. This behavior allows a node to be in its proper position in the tree and also temporarily be head in the highest level list. A node in this position has exclusive access to the tree and can perform a StoreUpdate as described above.

6 Acknowledgements

We wish to thank David V. James, Ross E. Johnson and Stein Gjessing for their suggestions, contributions and constructive comments on this work. We also thank many other members of the working group for their suggestions and their encouragement.

7 References

[1] IEEE Standard for Scalable Coherent Interface (SCI) 1596-1992, IEEE 1993.

[2] Ross E. Johnson. "Extending the Scalable Coherent Interface for large-Scale Shard-Memory Multiprocessors," PhD Thesis, University of Wisconsin-Madison, 1993.

[3] Gregory F. Pfister and V. Alan Norton, "'Hot Spot' Contention and Combining in Multistage Interconnection Networks." *Proceedings of the 1985 International Conference on Parallel Processing*, pp. 790-797, August 20-23 1985.

[4] Marvin H. Solomon and Raphael A. Finkel, "Processor interconnection strategies," *IEEE Transactions on Computers*, Vol. **C-29**, pp. 360-371, May 1980.

[5] Raphael A. Finkel and Marvin H. Solomon, "The lens interconnection strategy," *IEEE Transactions on Computers*, Vol. **C-30**, no. 12, December 1981.

[6] William J. Dally and Charles L. Seitz, "Deadlock-Free Routing in Multiprocessor Interconnection Networks," *IEEE Transactions on Computers*, Vol. **C-36**, no. 5, pp. 547-553, May 1987.

[7] Duncan H. Lawrie, "Access and Alignment of Data in an Array Processor." *IEEE Transactions on Computers*, Vol. **24**, no. 12, pp. 1145-1155, December 1975.

[8] R. Simoni and M. Horowitz, Dynamic Pointer Allocation for Scalable Cache Coherence Directories. *Proceedings of the International Symposium on Shared Memory Multiprocessing*, 1991.

[9] Steve L. Scott and James R. Goodman, "Performance of Pruning-Cache Directories for Large-Scale Multiprocessors," *IEEE Transaction on Parallel and Distributed Systems*, Vol. **4**, no. 5, May 1993.

[10] J. L. Baer and W. H. Wang, "Architectural choices for multi- level cache hierarchies," *Proc. 16th International Conference on Parallel Processing*, pp. 258-261, 1987.

[11] John Carter, John Bennet and Willy Zwaenopoel, "Munin: Distributed Shared Memory Based on Type-Specific Memory Coherence," *Proceedings of the Conference on the Principles and Practices of Parallel Programming*, 1990.

[12] Wolf-Dietrich Weber and Anoop Gupta, "Analysis of Cache Invalidation Patterns in Multiprocessors," *ASPLOS III*, pp. 243-256, 1989.

[13] James R. Goodman and Stefanos Kaxiras, "Hierarchical Extensions to SCI," University of Wisconsin Computer Sciences Technical Report #1235, July 1994. Available through anonymous ftp at ftp.cs.wisc.edu.