



(22) Date de dépôt/Filing Date: 1991/05/10

(41) Mise à la disp. pub./Open to Public Insp.: 1991/11/19

(45) Date de délivrance/Issue Date: 2001/12/18

(30) Priorité/Priority: 1990/05/18 (526,396) US

(51) Cl.Int.⁵/Int.Cl.⁵ H04Q 3/42

(72) Inventeurs/Inventors:

Binder, Paul A., US;
Kelly, Dana R., US;
Frank, Steven J., US;
Morss, Stephen W., US;
Epstein, David I., US;
Rothnie, James B., US;
Burkhardt, Henry III, US


(73) Propriétaire/Owner:

SUN MICROSYSTEMS, INC., US

(74) Agent: RICHES, MCKENZIE & HERBERT LLP

(54) Titre : COMMUTATEUR D'ACHEMINEMENT DE PAQUETS

(54) Title: PACKET ROUTING SWITCH



(57) Abrégé/Abstract:

A switch is provided for selectively routing digital information packets received from at least first and second external sources to at least first and second external destinations. At least one of the first sources generates an information packet including a



(57) Abrégé(suite)/Abstract(continued):

datum (or a request therefor) and a corresponding descriptor. First and second routing interconnects have inputs for receiving packets from respective sources and outputs for transmitting packets to respective destinations. The interconnects are also coupled for transferring packets between one another. Directories within the interconnects store descriptors corresponding to data associated with the first destination, as well as requests routed from the other interconnect. A controller routes packets based on the correspondence, or lack thereof, between the descriptor in that packet and an entry signal allocated to corresponding directory.

ABSTRACT

A switch is provided for selectively routing digital information packets received from at least 5 first and second external sources to at least first and second external destinations. At least one of the first sources generates an information packet including a datum (or a request therefor) and a corresponding descriptor. First and second routing 10 interconnects have inputs for receiving packets from respective sources and outputs for transmitting packets to respective destinations. The interconnects are also coupled for transferring packets between one another. Directories within the 15 interconnects store descriptors corresponding to data associated with the first destination, as well as requests routed from the other interconnect. A controller routes packets based on the correspondence, or lack thereof, between the 20 descriptor in that packet and an entry signal allocated to corresponding directory.

PACKET ROUTING SWITCH

5 BACKGROUND OF THE INVENTION

This invention relates to digital data communications networks and more particularly to packet switched networks for use, for example, with high-speed distributed multiprocessing systems.

-2-

Multiple instruction/multiple data (MIMD) parallel processing computers can be categorized according to address space, memory organization and memory hierarchy. With respect to the former, a
5 system may be characterized as utilizing a single or a multiple address space. A single address space is commonly referred to as shared memory and implies implicit communication as part of any memory access. A multiple private address spaces implies explicit
10 communication by, for example, message passing.

Memory organization may be characterized as centralized or distributed. In a system based on centralized memory, a common memory element is centrally located, with access time to any physical
15 memory location being the same for all processors. In a distributed memory system, on the other hand, system memory is divided into modules with some placed near each processor.

-3-

Memory hierarchy may be characterized as static, mixed or dynamic. A hierarchy is classified as static where system data is divided among local and/or global memory units, with each datum having an explicitly assigned address based at least in part on the locale of the unit in which it is stored. In a fully dynamic system, on the other hand, individual datum are not assigned physical memory location-based addresses by which they are accessed. Mixed memory hierarchies contain static local and/or global

memory, where a portion of the memory hierarchy is dynamic (e.g., a cache memory) and another portion is static.

The art provides a number of architectures where a single address space, i.e., shared memory, is organized in a centralized manner. Processing units in such systems communicate via high-bandwidth shared buses or switching networks. A problem with such architectures is that the shared memory forms a bottleneck, impeding system performance, except in instances where there are relatively few processors.

To avoid this problem, Frank et al, United States Patent No. 4,622,631, discloses a multiprocessing system in which a plurality of processors, each having an associated private memory, or cache, share data contained in a main memory element. Data within that common memory is partitioned into blocks, each of which can be owned by any one of the main memory and the plural processors. The current owner of a data block is said to have the correct data for that block.

While the solution suggested in the aforementioned patent permits an increase in the number of processors which can be supported without a bottleneck, the centralized memory system is not scalable.

In order to achieve scalability, a distributed memory organization must be used since it theoretically allows parallel high bandwidth access to memory to grow in proportion to the number of processor and memory modules. The art provides two alternative programming models for such an organization: the multiple address space model and

the single address space module. Both models present a dilemma for computer system designers and programmers.

From the programmers viewpoint, the single
5 address architecture is a simpler programming model since data movement is implicit in memory operations. Whereas, in the multiple address architecture explicit message passing is required for a multiple address architecture. The multiple
10 address architecture, moreover, requires explicit data localization, explicit memory allocation and deallocation, explicit replication and explicit coherency to produce a correct parallel program. These aspects are theoretically handled implicitly in
15 a single address architecture.

As discovered by designers in the prior art, from a hardware perspective simultaneous access to a single address space, i.e., a logically shared memory, is prohibitively expensive. This is due to
20 the complexity of the switching network, as well as the lack and potential complexity of a general high performance solution for memory coherency. In consequence, most scalable distributed memory systems such as Intel/IPSC have historically implemented a
25 multiple address architecture.

The performance of a single address architecture is dependent on the memory hierarchy. A static memory hierarchy requires the programmer to explicitly manage the movement of data for optimal
30 performance in a manner similar to the multiple address architecture. Two examples of a distributed memory organization which implement a single address architecture using a static memory hierarchy are the BBN Butterfly and IBM RP3. Such implementations
35 require the programmer to explicitly manage coherency.

Mixed hierarchies necessarily include communication bottlenecks which limit scalability, while still requiring the programmer to partially manage data movement. One such hierarchical approach 5 is disclosed by Wilson Jr. et al, United Kingdom Patent Application No. 2,178,205, wherein a multiprocessing system is said to include distributed cache memory elements coupled with one another over a first bus. A second, higher level cache memory, 10 attached to the first bus and to either a still higher level cache or to the main system memory, retains copies of every memory location in the caches, if any, and system main memory, in turn, retain copies of each memory location of cache below 15 them. The Wilson Jr. et al processors are understood to transmit modified copies of data from their own dedicated caches to associated higher level caches and to the system main memory, while concurrently signalling other caches to invalidate their own 20 copies of that newly-modified data.

Notwithstanding solutions proposed by the prior art, none has achieved a high performance, fully dynamic, coherent shared memory programming environment with unlimited scalability.

25 An object of this invention is to provide such a system. More particularly an object of the invention is to provide a multiple instruction/multiple data parallel processing system utilizing a shared memory addressing model and a distributed 30 organization with improved coherency.

A further object is to provide a fully dynamic memory hierarchy achieving high performance within a distributed system utilizing a shared memory address model.

A still further object is to provide an improved digital data communications network.

Yet another object is to provide a packet switch network for use, for example, in high-speed distributed multiprocessing systems.

Still another object is to provide an improved switching mechanism for use in routing data and data requests through a digital communications network.

SUMMARY

These and other objects are attained by the invention which provides an improved digital data communications network, packet switch and associative
5 directory for use, for example, with a multiprocessing computer system.

In one aspect, the system includes two groups of processing cells and a routing cell that transfers digital information packets between those
10 groups based upon the identity of data referenced in the packets and data requested by or allocated to a first one of those processing groups, or routed between both groups.

Each processing group has a plurality of
15 processing cells, e.g., central processing units and associated local cache memories, that communicate by transferring data and request packets over an associated group bus. The individual cells have allocated to them data or data copies identifiable by
20 unique descriptors. Cache control units in the processing cells generate packets requesting data, as well as those providing data or data copies in response to such requests. Each packet includes a descriptor corresponding to the requested or supplied
25 data.

The routing element includes first input and output sections connected with the bus of the first processing group, along with second input and output sections connected to the bus of the second
30 processing group, for receiving and transmitting packets to the cells of those groups. Request and response packets received at the first and second inputs are routed to one or both of the outputs according to whether the data identified by the

descriptors in those packets are allocated to cells of the first processing group.

To facilitate determining where to route packets received by the routing cell, an associative
5 memory-based directory stores a list of descriptors corresponding to data allocated to cells of the first processing group. The directory additionally stores a list of descriptors of data requested by those cells and routed to the second processing group for
10 handling, as well as a list of descriptors of data requested by cells in the second group and routed to the first for handling. Along with these descriptors, the directory maintains information about the access state of the data. This includes,
15 for example, invalid, read-only, non-exclusive ownership, exclusive ownership, and atomic states, in addition to pending states.

According to related aspect of the invention, the aforementioned second processing group
20 itself includes a plurality of further processing groups, themselves interconnected via routing cells and network buses and each comprising a plurality of processing cells interconnected by intra-group buses. By way of example, in a two-level system,
25 each of further groups is associated with a routing cell that selectively transfers information packets to and from the individual group buses. A higher-level bus, referred to below as a level:1 bus, connects these further routing cells with the routing
30 cell associated with the first processing group.

In another related aspect, the invention provides a hierarchical digital data communications network having a plurality of packet transfer levels, each including a plurality of transfer segments made

up of interconnected processing cells. The cells have allocated to them data, or data copies, each identifiable by a unique descriptor. The cells also include cache control units that can generate request
5 and response packets referencing data by their respective descriptors. The number of transfer segments decreases at each higher transfer level, there being only one segment at the highest level.

Within the network, routing cells provide
10 communications between the processing cells at each level with those of the higher levels. The routing cells, which transfer information packets between the transfer level segments, selectively route request packets based upon the association of data requested
15 in that packet and the cells of the descendant transfer segments. Particularly, these routing cells route to higher transfer segments those packets requesting data that is not associated with a descendant segment, while routing to lower segments
20 those packets requesting data that is associated with a descendant segment.

According to another aspect of the invention, a request packet can include a requestor-ID field identifying the processing cell
25 that generated the request. A cell generating a response to that request copies the requestor-ID into the response packet. The routing cells route such a response packet back to the requestor in accord with the value of the requestor-ID field.

30 In addition, the routing cells can route to one requesting processing cell a response packet that includes data requested in a previous, unrelated packet -- that is, a request packet generated by another cell. Thus, for example, a response packet

-11-

containing a read-only copy of a datum requested independently by two processing cells may be routed to both, notwithstanding that the requestor-ID of only one of them is included in the response packet.

5 Accordingly, in one aspect, the present invention provides a switch for selectively routing digital information packets received from a plurality of external sources to a plurality of external destinations, each said external destination having
10 associated therewith zero, one or more data items, each unique one of said data items having an associated descriptor for use in referencing the associated data item, said external sources including means for generating information packets, each
15 information packet including a said descriptor of a said data item, said switch comprising first and second routing interconnect elements, each said routing interconnect element including

A. a source input means for receiving an
20 information packet from at least a respective one of said external sources,

B. a destination output means for transmitting an information packet to at least a respective one of said external destinations,

25 C. internal input means for receiving an information packet from the other of said routing interconnect elements, and

D. internal output means for transmitting an information packet to the other of said routing
30 interconnect elements,

-11a-

E. directory means having allocated therein entries each respectively containing a descriptor corresponding to at least one of

- 5 i) one of said data items associated with said at least a respective one of said external destinations,
- ii) a data item requested in an information packet received by the internal input means of the respective routing interconnect element and
10 routed for transmission by the destination output means of the respective routing interconnect element, and

F. control means for selectively transferring an information packet received by at
15 least one of the source input means and the internal input means of the respective routing interconnect element, to at least one of the destination output means and the internal output means of the respective routing interconnect element based on the
20 correspondence, or lack thereof, between the descriptor in that information packet and an entry containing a corresponding descriptor in said directory means.

In another aspect, the present invention
25 provides a method of routing digital information packets received from a plurality of external sources to a plurality of external destinations, each said external destination having associated therewith zero, one or more data items, each unique one of said data
30 items having an associated descriptor for use in referencing the associated data item, said external

-11b-

sources including means for generating information packets, each information packet including a said descriptor of a said data item, said method comprising:

5 A. providing first and second routing interconnect elements, each said routing interconnect element including

10 i) a source input means for receiving an information packet from at least a respective one of said external sources,

 ii) a destination output means for transmitting an information packet to at least a respective one of said external destinations,

15 iii) internal input means for receiving an information packet from the other of said routing interconnect elements, and

20 iv) internal output means for transmitting an information packet to the other of said routing interconnect elements,

 B. allocating in a directory means associated with each routing interconnect element entries each respectively containing a descriptor corresponding to at least one of

25 i) one of said data items associated with said at least a respective one of said external destinations,

30 ii) a data item requested in an information packet received by the internal input means of the respective routing interconnect element and routed for transmission by the destination

-11c-

output means of the respective routing
interconnect element, and

C. transferring an information packet
received by at least one of the source input means and
5 the internal input means of the associated routing
interconnect element, to at least one of the
destination output means and the internal output means
of based on the correspondence, or lack thereof,
between the descriptor in that information packet and
10 an entry containing a corresponding descriptor in said
directory means.

These and other aspects of the invention are
discussed in the detailed description which follows in
the attached drawing.

BRIEF DESCRIPTION OF THE ILLUSTRATED EMBODIMENT

A fuller understanding of the invention may be attained by reference to the drawings, in which:

Figure 1 depicts the structure of a
5 preferred multiprocessing system constructed in
accord with the invention;

Figure 2A depicts a preferred memory
configuration providing data coherence in a
multiprocessing system of the type shown in Figure 1;

10 Figure 2B illustrates datum movement in
conjunction with an ownership request by processing
cell which does not already have access to that datum
in a preferred multiprocessing system constructed in
accord with the invention;

15 Figure 3 depicts a preferred configuration
for an exemplary level:0 segment of Figure 1;

Figure 4 depicts a preferred structure for
processing cells used in practice of the invention;

20 Figure 5 depicts a preferred structure of a
cache directory in a processing cell constructed in
accord with the invention;

Figure 6 depicts a preferred ring routing
cell constructed in accord with the invention;

25 Figure 7 depicts a preferred ring
interconnection unit for a ring routing cell used in
practice of the invention;

Figures 8 and 9 depict the organization of a
routing directory constructed in accord with a
preferred practice of the invention;

30 Figure 10 depicts a format of routing
directory entries used in a preferred practice of the
invention;

Figure 11A depicts a preferred structure of the InterRingBuffer section of a ring routing cell constructed in accord with the invention;

Figure 11B depicts a preferred arrangement
5 of InterLinkDram buffers in a ring routing cell constructed in accord with the invention;

Figure 12 depicts a preferred InterRingLink frame format used in a system constructed in accord with the invention;

10 Figure 13 illustrates a preferred request routing sequence in a system constructed in accord with the invention;

Figures 14A and 14B depict a preferred invalidate sequence in a system constructed in accord
15 with the invention;

Figure 15 depicts a preferred deallocate sequence used in a system constructed in accord with the invention;

Figure 16 depicts a full ring freeze in a
20 system constructed in accord with the invention;

Figure 17 is a table showing RDU:0 states in an alternative embodiment of the invention;

Figure 18 is a table showing RDU:1 states in an alternative embodiment of the invention; and

25 Figures 19A-19D depict preferred and alternate ring interleaving embodiments of the invention.

DETAILED DESCRIPTION OF THE ILLUSTRATED EMBODIMENT

Overview

A preferred digital data processing system
5 utilizing a switching network constructed in accord
with the invention comprises a plurality of
processing cells arranged in a hierarchy of rings,
each including one or more routing interconnects
coupled by way of a ring-like bus to zero, one or
10 more processing cells. Routing cells linking the
interconnects provide pathways for selectively
transferring information packets between the rings.
Directories in the routing cells track data (and
requests) transferred through the associated cells,
15 thereby determining routing pathways for subsequent
routing of subsequent packets.

In one preferred embodiment, the processing
cells include central processing units coupled with
local memory elements, or caches, each having a
20 physical data and control signal store, a directory,
and a control element. Groups of cells are
interconnected along unidirectional intercellular bus
rings, forming units referred to as segments. These
segments together form a larger unit referred to as
25 "information transfer level:0."

Communications between cells of different
level:0 segments are carried out over higher level
information transfer levels, e.g., information
transfer level:1 and information transfer level:2.
30 These higher levels are themselves made up of one or
more segments, each comprising a plurality of routing
cells coupled via a unidirectional bus ring. Each
routing cell is connected with an associated segment
of the next lower information transfer level. These

connected lower segments are referred to as "descendants."

Every information transfer level includes fewer segments than the next lower level. Apart from 5 the single segment of the system's highest level, signals are transferred between segments of each information transfer level via segments of the next higher level.

The routing cells each include a directory 10 of descriptors corresponding to the data allocated to the memory elements of the descendant segments. By way of example, a routing cell connecting a level:0 segment to a level:1 segment includes a directory of all data allocated to cells in that level:0 segment, 15 while a routing cell connecting a level:1 segment to a level:2 segment includes a directory of data allocated to all cells in the level:0 segments connected to that level:1 segment.

Along with descriptors, the directories 20 store the access states of data allocated in the descendant segments. These states include, among others, an ownership state, a read-only state, and an invalid state. The ownership state, like the atomic state, is associated with datum which can be modified 25 only by the local processor; however, unlike atomic state data, ownership state data can be accessed -- e.g., for generating read-only copies -- by other processing cells. The read-only state is associated with data which may be read, but not modified, by the 30 local central processing unit. The invalid state is associated with invalid data copies.

In this regard, it will be appreciated that a preferred digital data processing system constructed in accord with the invention does not

utilize a main memory element, i.e., a memory element coupled to and shared by the multiple processing cells. Rather, data is distributed, both on exclusive and shared bases, among the local memory
5 elements associated with each of those processors. Modifications to datum stored exclusively in any one processing cell need not be communicated along the network to the other processing cells. As a result, only that data which the cells dynamically share is
10 transmitted on network, i.e., along the ring buses and routing cells, thereby reducing bus contention and bottlenecks.

The routing cell directories also store descriptors corresponding to data requested in
15 packets routed by the associated routing cells. As above, along with these descriptors, the directories store pending state information indicating requested access states of the data.

In view of the above, it will be appreciated
20 that the contents of the routing cell directories -- to wit, the allocated data and pending request lists -- reflect the prior history of routings of data and data requests routed between the processing groups.

The routing cells use their directories of
25 descriptors and states to determine whether specific information packets travelling along the level segments should be passed onwards to other processing cells in the segment in which the packet originated or, instead, routed upwards or downwards to other
30 segments.

To further understand this routing function, it will be appreciated that in a preferred multiprocessor system constructed in accord with the invention, data access requests are handled within

the processing cell that originated them whenever possible. Particularly, a controller coupled with each memory monitors the cell's internal bus and responds to local processor requests by comparing the
5 request with descriptors listed in a local cache directory. If found, matching data from the cache itself is transmitted back along the internal bus to the requesting processor.

In the event a datum request cannot be
10 resolved locally, the processing cell generates an information request packet that is transferred by a local cache control unit to local ring bus. As that packet travels along the ring, control elements of processing cells adjacent the requester check their
15 own directories, passing the requested data, if found, in a response packet.

When an unresolved request passes the routing element of the segment to which the requesting processing cell belongs, that element too
20 checks its directory. If that directory shows that the requested data is present in local ring in the proper access state, the routing cell permits the request packet to continue along the local ring bus. If not, the routing cell extracts the packet and
25 passes to the associated level:1 segment.

The unresolved request packet travels along the level:1 segment in a similar fashion, i.e., it is compared with the directories of the associated level:1 routing cells. If one of those directories
30 lists the requested data in the proper access state in a descendant level:0 segment, the request is passed to that segment. Otherwise, the request packet may pass to a higher level, if any, in the hierarchy or be returned to the requester unresolved.

Packets containing requested data are routed back to requesting cells by differing mechanisms. A first mechanism relies on the address, or ID, of the requesting cell. Here, each requesting cell includes
5 within its request packet an ID uniquely identifying that cell. Once that packet reaches the responding cell, that cell copies the requestor-ID, along with the datum and corresponding descriptor, into the response packet. As the response packet travels
10 along the segments rings, the routing cells examine the requestor-ID to determine whether the requesting cell is in a descendant or parent segment and route the packet accordingly.

A second mechanism is used in conjunction
15 with response packets that include data requested by processing cells but not specifically generated in response to those requests. By way of example, the mechanism applies in instances where two or more requesting cells generate requests for read-only
20 copies to a specific datum held in a remote cell.

Presuming, in accord with one preferred practice of the invention, that the network prevents at least some, but not all, of those requests from reaching a processing cell having a copy of that
25 datum, the responding cell generates response packets bearing only the requestor-ID of a request packet that reached it.

As the response packet is routed along the network, routing cells compare the packet's
30 descriptor and access-state information with pending request entries in their directories. If that comparison reveals that routing cell previously received a request for datum of the type in the packet, that packet (or a copy thereof) may be routed

to one of the requesting cells. In this manner, a single packet containing, for example, a read-only copy of a datum may effectively serve as a response to a plurality of outstanding read-only requests.

5

System Structure

Figure 1 depicts a preferred multiprocessing system utilizing a communications network constructed in accord with the invention. The illustrated system 10 includes three information transfer levels: level:0, level:1, and level:2. Each information transfer level includes one or more level segments, characterized by a bus element and a plurality of interface elements. Particularly, level:0 of the 15 illustrated system 10 includes six segments, designated 12A, 12B, 12C, 12D, 12E and 12F, respectively. Similarly, level:1 includes segments 14A and 14B, while level:2 includes segment 16.

Each segment of level:0, i.e., segments 12A, 20 12B, ... 12F, comprise a plurality of processing cells. For example, segment 12A includes cells 18A, 18B and 18C; segment 12B includes cells 18D, 18E and 18F; and so forth. Each of those cells include a central processing unit and a memory element, 25 interconnected along an intracellular processor bus (not shown). In accord with the preferred practice of the invention, the memory element contained in each cells stores all control and data signals used by its associated central processing unit.

30 As further illustrated, each level:0 segment may be characterized as having a bus element providing a communication pathway for transferring information packets between the cells of the segment. Thus, illustrated segment 12A is

characterized by bus 20A, segment 12B by 20B, segment 12C by 20C, et cetera. As described in greater detail below, digital information packets are passed between the cells 18A, 18B and 18C of exemplary
5 segment 12A by way of the memory elements associated with each of those cells. Specific interfaces between those memory elements and the bus 20A are provided by cell interface units 22A, 22B and 22C, as shown. Similar direct communication pathways are
10 established in segments 12B, 12C and 12D between their respective cells 18D, 18E, ... 18R by cell interface units 22D, 22E, ... 22R, as illustrated.

As shown in the illustration and noted above, the remaining information transfer levels,
15 i.e., level:1 and level:2, each include one or more corresponding level segments. The number of segments in each successive segment being less than the number of segments in the prior one. Thus, level:1's two segments 14A and 14B number fewer than level:0's six
20 12A, 12B ... 12F, while level:2, having only segment 16, includes the fewest of all. Each of the segments in level:1 and level:2, the "higher" levels, include a bus element for transferring packets within the respective segments. In the illustration, level:1
25 segments 14A and 14B include bus elements 24A and 24B, respectively, while level:2 segment 16 includes bus element 26.

The routing cells themselves provide a mechanism for transferring information between
30 associated segments of successive levels. Routing cells 28A, 28B and 28C, for example, provide a means for transferring information to and from level:1 segment 14A and each of level:0 segments 12A, 12B and 12C, respectively. Similarly, routing cells 28D,

28E and 28F provide a means for transferring information to and from level:1 segment 14B and each of level:0 segments 12D, 12E and 12F, respectively. Further, routing cells 30A and 30B provide an
5 information transfer pathway between level:2 segment 16 and level:1 segments 14A and 14B, as shown. As discussed in further detail below, the routing cells each comprise two like processing sections.

The routing cells interface their respective
10 segments via interconnections at the bus elements. Thus, routing cell 28A interfaces bus elements 20A and 24A at ring interconnects 32A and 34A, respectively, while element 28B interfaces bus elements 20B and 24B at ring interconnects 32B and
15 34B, respectively, and so forth. Similarly, routing cells 30A and 30B interface their respective buses, i.e., 24A, 24B and 26, at ring interconnects 36A, 36B, 38A and 38B, as shown.

Figure 1 illustrates further a preferred
20 mechanism interconnecting remote levels and cells in a digital data processing system constructed in accord with the invention. Cell 18R, which resides at a point physically remote from bus segment 20F, is coupled with that bus and its associated cells (18P
25 and 18O) via a fiber optic transmission line, indicated by a dashed line. A remote interface unit 19 provides a physical interface between the cell interface 22R and the remote cell 18R. The remote cell 18R is constructed and operated similarly to the
30 other illustrated cells and includes a remote interface unit for coupling the fiber optic link at its remote end.

In a like manner, level segments 12F and 14B are interconnected via a fiber optic link from their

parent segments. As indicated by dashed lines, the respective sections of the routing cells 28F and 30B are physically separated, yet coupled for example by a fiber optic link. With respect to routing unit 5 28F, for example, a first part is linked directly via a standard bus interconnect with routing interface 34F of segment 14B, while a second part is linked directly with routing interface 32F of segment 12F. As above, a physical interface between the routing 10 unit parts and the fiber optic media is provided by a remote interface unit (not shown).

Figure 2A illustrates a preferred mechanism for maintaining data coherence in a multiprocessing system of the type described above. The illustrated 15 system includes plural central processing units 40(A), 40(B) and 40(C) coupled, respectively, to associated memory elements 42(A), 42(B) and 42(C). Communications between the processing and memory units of each pair are carried along buses 44A, 44B 20 and 44C, as shown. Network 46, representing the aforementioned level segments and routing cells, transfers information packets (passed to the network 46 over buses 48(A), 48(B) and 48(C)) between the illustrated processing cells 42A - 42C.

25 In the illustrated embodiment, the central processing units 40A, 40B and 40C each include an access request element, labelled 50A, 50B and 50C, respectively. These access request elements generate requests for access to data stored in the memory 30 elements 42A, 42B and 42C. Among access requests signals generated by elements 50A, 50B and 50C is the ownership-request, representing a request for exclusive, modification access to a datum stored in the memory elements. In a preferred embodiment,

access request elements 50A, 50B and 50C comprise a subset of an instruction set implemented on CPU's 40A, 40B and 40C. This instruction subset is described below.

5 The memory elements 40A, 40B and 40C include control elements 52A, 52B and 52C, respectively. Each of these control units interfaces a data storage area 54A, 54B and 54C via a corresponding directory element 56A, 56B and 56C, as shown. Stores 54A, 54B
10 and 54C are utilized by the illustrated system to provide physical storage space for data and instruction signals needed by their respective central processing units. Thus, store 54A maintains data and control information used by CPU 40A, while
15 stores 54B and 54C maintain information used by central processing units 40B and 40C, respectively. The data maintained in each of the stores are identified by unique descriptors corresponding to system addresses. Those descriptors are stored in
20 address storage locations of the corresponding directory. While the descriptors are considered unique, multiple copies of some descriptors may exist among the memory elements 42A, 4B and 42C where those copies themselves identify copies of the same data
25 element.

Access requests generated by the central processing units 40A, 40B and 40C include, along with other control information, a descriptor or SVA ("system virtual address") request portion matching
30 that of the requested datum. The control elements 52A, 52B and 52C respond to access requests generated their respective central processing units 40A, 40B and 40C by determining whether the requested datum is stored in the corresponding storage element 54A, 54B

and 54C. If so, that item of information is transferred for use by the requesting processor. If not, the control unit 52A, 52B, 52C transmits a packet including the request to the network 46 along 5 lines 48A, 48B and 48C.

The control units 52A - 52C also monitor the network 46 to determine whether they can satisfy remote access requests (i.e., requests received from other processing cells). As above, these control 10 units compare the descriptor and state of an access request received on the network 46 with descriptors stored in their directories. If the requested datum is found in the requested state, it is transferred in a response packet back onto the network 46 for 15 routing to the requesting unit. If a requested datum does not reside in any of the system's processing cells, the operating system can search the system peripheral devices.

Data coherency is maintained by cooperative 20 action of the processing cells in response to data requests and transfers. More particularly, concurrent with generation of an ownership-access request packet by a first processing cell, the associated memory allocates in its store physical 25 space to hold the requested data. Likewise, concurrent with the transfer of a requested datum from the processing cell in which it was previously stored, the associated memory deallocates physical storage space which had been previously allocated for 30 storage of the requested datum.

These cooperative actions are illustrated in Figures 2A and 2B. In the first of those drawings, DATUM(0), DATUM(1) and DATUM(2) are retained in the store of the memory element 42A partnered with CPU

40A. Descriptors "foo," "bar" and "bas" which correspond, respectively, to those data stored in directory 56A. Each such descriptor includes a pointer indicating the location of its associated
5 information signal in the store 42A.

The memory element 42B partnered to CPU 40B stores DATUM(3) and DATUM(2). Corresponding to each of those data elements are descriptors "car" and "bas," retained in directory 56B. DATUM(2), and its
10 descriptor "bas," are copied from store 42A and, therefore, retain the same labels.

The system illustrated in Figure 2A does not store any data in the memory element 54C partnered to CPU 40C.

15 Figure 2B illustrates how a datum moves in conjunction with an ownership for it by processing cell which does not already have access to that datum. In particular, the illustration depicts the movement of DATUM(0) following issuance of an
20 unresolved request for ownership for it by the processing cell made up of CPU 40C and memory 42C. At the outset, concurrent with the request, the control unit 52C allocates physical storage space in the store 54C of memory element 42C. A response
25 packet generated by memory 42A transfers the requested data, DATUM(0), from store 54A, where it had previously been stored, to the requestor's store 54C. Concurrently the control unit 52A deallocates that space in store 54A which had previously held the
30 requested datum. At the same time, the control unit 52A effects invalidation of the descriptor "foo" in directory 56A (where it had previously been used to identify DATUM(0) in store 54A), while control unit 52C reallocates that same descriptor in directory

-26-

56C, where it will subsequently be used to identify the signal in store 54C.

In addition to descriptors, the memory elements 42A - 42C assign access state information to the data and control signals in their respective stores. These access states, which include the invalid, read-only, ownership and atomic states, govern the manner in which data may be accessed by specific processors. A datum which is stored in a memory element whose associated CPU maintains modification access over that datum is assigned an ownership state. While, a datum which is stored in a memory element whose associated CPU does not maintain priority access over that datum is assigned a read-only state. Further, a purported datum which associated with "bad" data is assigned the invalid state.

Figure 3 depicts a preferred configuration for exemplary level:0 segment 12A of Figure 1. The segment 12A includes processing cells 18A, 18B and 18C interconnected by cell interconnects 22A, 22B and 22c along bus segment 20A. Routing unit 28A provides an interconnection between the level:0 segment 12A and if parent, level:1 segment 14a of Figure 1. This routing unit 28A is coupled along bus 20A by way of cell interconnect 32A, as shown. The structure of illustrated bus segment 20A, as well as its

-27-

interrelationship with cell interconnects 22A, 22B, 22C and 32A, may be fully appreciated by reference to the above-cited related applications, e.g., United States Patent Application No. 509,480, filed 5 April 13, 1990 (Attorney Docket No. KSP-002CN).

Figure 4 depicts a preferred structure for an exemplary processing cell 18A of the illustrated embodiment. The illustrated processing cell 18A includes a central processing unit 58 coupled with 10 external device interface 60, data subcache 62 and instruction subcache 64 over processor bus 66 and instruction bus 68, respectively.

The Memory System

15 As discussed above, a multiprocessing system 10 constructed in accord with a preferred embodiment of the invention permits access to individual data elements allocated to processing cells 18A, 18B, . . . 18R by reference to a unique descriptor, based on a 20 system virtual address (SVA), associated with each datum. In the illustrated system 10, implementation of this capability is provided by the combined actions of the processing cell memory units in connection with their transfer of request and 25 response packets over the network 46. In the

discussion which follows, this is collectively referred to as the "memory system."

A complete understanding of the structure and operation of the memory system may be attained through recognition of its architectural features, enumerated below:

- 10 (1) Data Storage -- The memory in each cache is divided into pages, each of which may be dynamically assigned to some page of SVA space. The memory system maintains usage and status information about the data in each cache to facilitate efficient migration to and from secondary storage.
- 15 (2) Data Locality -- The memory system keeps data recently referenced by a processor in the subcache or cache in the same cell of that processor.
- 20 (3) Data Movement -- The memory system moves data to the cache of the processor referencing it.
- 25 (4) Data Sharing -- The memory system keeps copies of SVA data in more than one cache to facilitate efficient data sharing by parallel programs.
- 30 (5) Data Coherence -- The memory system implements the strongly ordered coherent memory model and the transaction model.
- 35

Those skilled in the art will appreciate that a system is "sequentially consistent" if the result of any execution is the same as if the operations of all the processors were executed in some sequential order, and the operations of each individual processor appear in this sequence in the order specified by its program.

Moreover, storage accesses are considered "strongly ordered" if accesses to data by any one processor are initiated, issued and performed in program order and; if at the time when a store by processor I is observed by processor K, all accesses to data performed with respect to I before the issuing of the store must be performed with respect to K. By contrast, storage accesses are weakly ordered if accesses to synchronizing variables are strongly ordered and; if no access to synchronizing variable is issued in a processor before all previous data accesses have been performed and; if no access to data is issued by a processor before a previous access to a synchronizing variable has been performed.

15 A coherent system with strong ordering of events is sequentially consistent.

The organization of SVA space within the illustrated system is a major departure from ordinary virtual memory schemes. Conventional architectures include a software controlled page-level translation mechanism that maps system addresses to physical memory addressor generates missing page exceptions. In these schemes, the software is responsible for multiplexing the page table(s) among all the segments in use. In the architecture of the illustrated system, there is no software controlled page-level translation mechanism. The memory system can handle a significant portion of the address space management normally performed by software in conventional architectures. These management responsibilities include:

20
25
30
35

- (1) maintaining page usage and status information,

- (2) reusing old pages,
- 5 (3) synchronizing and ensuring coherence of shared data access amongst multiple processors,
- 10 (4) migrating data and copies of data on a subpage basis from place to place in the system to keep data nearest to the processors that are using it most frequently,
- 15 (5) directly implementing a sparse representation of the entire SVA space, by migrating the SVA of data along with the data, and
- 20 (6) moving the page between referencing cells (no fixed home for a page).

The illustrated system's processors, e.g., processors 40A, 40B, 40C, communicate with the memory system via two primary logical interfaces. The first 25 is the data access interface, which is implemented by the load and store instructions. In data access mode, the processor presents the memory system with an SVA and access mode information, and the memory system attempts to satisfy that access by finding the 30 subpage containing the data and returning it.

The second logical interface mode is control access, which is implemented by memory system control instructions. In control access, the processor instructs the memory system to perform some side 35 effect or return some information other than the actual data from a page. In addition to the primary interfaces, system software uses control locations in SPA space for configuration, maintenance, fault recovery, and diagnosis.

Hierarchy

The memory system implements a hierarchy of storage. In the illustrated, preferred embodiment, each processing cell contains a central processing unit (or CEU) which has a subcache that stores 0.5 MBytes of data. Those processing cells also include caches, which store 32 MBytes of data. Moreover, a preferred level:0 having, for example, 15 processing cells stores a total of 480 MBytes. While, a preferred level:1 having, for example, 32 level:0's has a total of 15360 MBytes.

Cache Structure

The memory system stores data in units of 15 pages and subpages. Each individual cache describes 32 Mbytes of SVA space, subdivided into 2048 pages. Pages contain 2^{14} (16384) bytes, divided into 128 subpages of 2^7 (128) bytes. The memory system allocates storage in the caches on a page basis, and 20 each page of SVA space is either entirely represented in the system or not represented at all. The memory system shares data between caches in units of subpages.

When a page of SVA space is resident in the 25 system, the following are true:

- (1) One or more caches allocates a page of storage to the page, each subpage of the page is stored on one or more of the caches with space allocated, but 30
- (2) Each cache with space allocated for a page may or may not contain a copy of all of the page's subpages.

The associations between cache pages and SVA pages are recorded by each cache in its cache directory. Each cache directory is made up of descriptors. There is one descriptor for each page 5 of memory in a cache. At a particular time, each descriptor is said to be valid or invalid. If a descriptor is valid, then the corresponding cache memory page is associated with a page of SVA space, and the descriptor records the associated SVA page 10 address and state information. If a descriptor is invalid, then the corresponding cache memory page is logically not in use. There is no explicit validity flag associated with a descriptor; a descriptor may be considered invalid if the anchor and held fields 15 are both clear, and there are no valid subpages present for the SVA page.

Each cache directory acts as a content-addressable memory. This permits a cache to locate a descriptor for a particular page of SVA 20 space without an iterative search through all of its descriptors. Each cache directory is implemented as a 16 way set-associative memory with 128 sets. All of the pages of SVA space are divided into 128 equivalence classes, each associated with a cache 25 directory set. A descriptor for a page can only be stored in the set of a cache directory that corresponds to the page's equivalence class. The equivalence class is selected by SVA[20:14]. At any given time, a cache can describe no more than 16 30 pages with the same value for SVA[20:14], since there are 16 elements in each set.

Subpages and Data Sharing

When a page is resident in the memory system, each of its subpages is resident in one or more of the caches. When a subpage is resident in a
5 cache, the descriptor (in that cache) for the containing SVA page records the presence of that subpage in one of several states. The state of the subpage in a cache determines two things:

- 10 (1) What operations that cache's local processor may perform on the data present in the subpage; and
- (2) What responses, if any, that cache makes to requests for that subpage
15 received over the levels from other caches.

The states of subpages in caches change over time as programs request operations that require
20 particular states. A set of transition rules specify the changes in subpage states that result from processor requests and inter-cache level communications.

In order for a processor to complete an
25 instruction or data reference, several conditions must be simultaneously satisfied:

- 30 (1) For instruction references, the subblock containing the data must be present in the instruction subcache. For most data operations, the subblock containing the data must be present with the appropriate state in the data subcache;

- (2) The subpage containing the data must be present in the local cache; and
- (3) The local cache must hold the subpage in the appropriate state.

5

If the data is not present with the required state in the subcache, but is present in the local cache with correct state, the CEU obtains the data from the cache. If the local cache does not have the data in the correct state, it communicates over the levels to acquire a copy of the subpage and/or to acquire the necessary state for the subpage. If the cache fails to satisfy the request, it returns an error indication to the processor, which signals an appropriate exception.

The instruction set includes several different forms of load and store instructions that permit programs to request subpage states appropriate to the expected future data reference pattern of the current thread of control, as well as protocol between different threads of control in a parallel application. This section first describes the states and their transitions in terms of processor instructions and their effect on the caches.

25

Subpage States

The subpage states and their transition rules provide two general mechanisms to programs:

30

- (1) They transparently implement the strongly ordered sequentially consistent model of memory access for ordinary load and store accesses by the processors of the system; and

- (2) They provide a set of transaction primitives that are used by programs to synchronize parallel computations. These primitives can be applied to a variety of traditional and non-traditional synchronization mechanisms.

The basic model of data sharing is defined in terms of three classes of subpage states: invalid, read-only, and owner. These three classes are ordered in strength according to the access that they permit. invalid states permit no access. Read-only state permits load and instruction-fetch access. There are several owner states: all permit load access and permit the cache to respond to a data request from the interconnect; some permit store access. Only one cache may hold a particular subpage in an owner state at any given time. The cache that holds a subpage in an owner state is called the owner of the subpage. Ownership of each subpage moves from cache to cache as processors request ownership via store instructions and special load instructions that request ownership.

The sections below describe the state classes and how they interact to implement the strongly ordered sequentially consistent model of memory access.

Invalid States

When a subpage is not present in a cache, it is said to be in an invalid state with respect to that cache. If a processor requests a load or store to a subpage which is in an invalid state in its

local cache, then that cache must request a copy of the subpage in some other state in order to satisfy the data access. There are two invalid states: invalid-descriptor and invalid.

5 When a particular cache has no descriptor for a particular page, then all of the subpages of that page are said to be in invalid-descriptor state in that cache. Thus, subpages in invalid-descriptor state are not explicitly represented in the cache.

10 When the CEU references a subpage in invalid-descriptor state, the local cache must allocate one of the descriptors (in the correct set) to the SVA. After the descriptor allocation is complete, all subpages in the page have invalid state.

15 When the local cache has a descriptor for a particular page, but a particular subpage is not present in that cache, then that subpage is in invalid state. The local cache will attempt to obtain the subpage data by communicating with other
20 caches.

Read-Only State

 There is only one read-only state: read-only. Any number of caches may hold a
25 particular subpage in read-only state, provided that the owner of the subpage holds the subpage in non-exclusive state. If the owner of the subpage has any other state (i.e. an exclusive-ownership state, one of: exclusive, atomic, or transient-atomic), then
30 no read-only copies can exist in any cell. The CEU cannot modify a subpage which is in read-only state.

Owner States

There are two basic owner state types: non-exclusive and exclusive-ownership. When a particular cache holds a particular subpage in non-exclusive state, then some other caches may be holding that subpage in read-only state. Programmatically, non-exclusive state is the same as read-only state. The CEU cannot modify a subpage which is in non-exclusive state. Non-exclusive state is basically a book-keeping state used by the memory system; it defines the ownership of the subpage.

The exclusive-ownership states are exclusive, atomic, and transient-atomic. When a particular cache holds a particular subpage in an exclusive-ownership state, then no other cache may hold a read-only or non-exclusive copy of the subpage. If the local cache has a subpage in an exclusive-ownership state, the CEU can modify subpage data provided that the STT grants write access to the segment and the descriptor.no_write flag is clear.

Atomic State

Atomic state is a stronger form of ownership than exclusive state. Subpages only enter and leave atomic state as a result of explicit requests by programs.

Fundamentally, atomic state can be used to single-thread access to any subpage in SVA space. When a processor executes a gsp.nwt (get-subpage, no-wait) instruction to request that a subpage enter atomic state, the instruction will only complete normally if the subpage is not in atomic state already. Thus, atomic state on a subpage can be used as a simple lock. The lock is locked when the

gsp.nwt instruction completes normally by first obtaining the subpage in exclusive state and then changing state from exclusive to atomic. The lock is unlocked by executing the rsp (release-subpage) 5 instruction. The rsp instruction requires that the subpage exist in some cache with atomic or transient-atomic state. The local cache obtains the subpage and then changes the subpage from atomic or transient-atomic state to exclusive state. (If the 10 subpage has transient-atomic state, the operation is more complex, but the effect is programmatically the same.)

It is important to note that atomic state is associated only with a subpage; there is no 15 association with a particular operating system process (typically a user program) or to a particular cell. It is possible for a process to execute a gsp instruction to get a subpage in atomic state and subsequently be switched by system software so that 20 it continues execution on another cell. That process continues execution on the second cell and eventually executes an rsp instruction to release the subpage. Between those two instructions, there will only be a single copy of the subpage in the entire memory 25 system, and it will be in atomic or transient-atomic state. As various processors execute instructions which reference the subpage, that single valid copy will move from cell to cell. It is also possible for a particular process to get atomic state, and another 30 process to release atomic state.

Atomic state is simply an additional flag associated with a subpage; it is possible to implement protocols which use atomic state in addition to the data state of the subpage. Just as a

protocol which is implemented using only data can have errors, it is possible for an atomic-state protocol to be defective. The hardware does not impose any checking on the use of atomic state beyond
5 access control imposed by the STT and descriptor.no_atomic.

Transient-Atomic State

The gsp.nwt instruction always completes
10 within its defined execution time, but it can succeed or fail (depending upon the current state of the subpage in the memory system). A second form of the instruction is gsp.wt (get-subpage, wait), which will not complete until the subpage is obtained in
15 exclusive state and changed to atomic state. The gsp.wt instruction relieves the programmer of the burden of determining whether or not the gsp instruction was successful. If the subpage is already in atomic or transient-atomic state when a
20 processor executes gsp.wt, the processor will stall until the subpage is released, obtained by the local cache, and changed back from exclusive state to atomic or transient-atomic state. Use of the gsp.wt instruction can reduce the number of messages sent
25 between caches as a cell waits for the opportunity to 'lock the lock'.

Transient-atomic state is used automatically by the memory system to allow gsp.wt to function efficiently. Its use is entirely transparent to the
30 programmer. If a subpage is in atomic state and another cache executes gsp.wt on that subpage, that subpage enters transient-atomic state in the holding cache. When the subpage is later released with an rsp instruction, the transient-atomic state forces

the subpage to be expelled onto the interconnect in a special release state. The releasing caches changes its own state for the subpage to invalid. Any cache which is executing a gsp will see the subpage and
5 accept it. The accepting cache will then be able to complete its gsp instruction and the subpage will enter transient-atomic state in that cache. This operation will happen for each succeeding gsp and rsp until such time as an expelled subpage is not
10 accepted by any other cache. At that time, the cache performing the release will change its subpage state back from invalid state (set when the subpage was released) back to exclusive state.

It is also possible that the packet with
15 release state will be accepted by a cache whose CEU is performing a load or store instruction. The original cache sees that the subpage was accepted, and leaves its subpage in invalid state. The accepting cache allows its CEU to execute a single
20 instruction before it retracts the subpage, sets its own subpage state to invalid, and sends the subpage out with release state. The cache which was executing the load or store is now the owner of the page, for purposes of this release. As before, if no
25 other cache accepts the data, this cache will change the subpage state to exclusive state and retain ownership.

Preferably, there is no limit to the time which a gsp.wt instruction may wait. The process
30 which has issued the instruction will wait until the subpage is released. The instruction may be interrupted by various XIU signals; when such an event occurs, the CCUs abandon the attempt to gain atomic access. If the subpage is released in the

interim, and there is no other requestor, it will change from transient-atomic state release state and finally to exclusive state. Since the typical system software action is to service the interrupt and
5 restart the interrupted instruction, the CCU will again issue the request. It may succeed or be forced to wait, as before.

State Transition

10 This section provides some examples of subpage data and state flow.

 The basic mechanism by which data moves from some owning cache to other caches is by instruction fetches and the execution of load and store
15 instructions by the processors local to those other caches. The different load and prefetch instructions permit programs to request that their local cache acquire read-only or an exclusive-ownership state; a store instruction always requires that the subpage
20 have an exclusive-ownership state. In some circumstances, a cache may acquire a read-only copy of a subpage as it passes by on the interconnect. The post-store-subpage (pstsp) instruction broadcasts a read-only copy of a subpage to all interested
25 caches. Finally, the owning cache may send ownership state on the interconnect as part of recombining a page (described later).

 Instruction fetches and load instructions can result in the local cache requesting a read-only
30 copy of the subpage. This request is answered by the cache which owns the subpage. If the owning cache has the subpage in non-exclusive state, it supplies a read-only copy to the requesting cache, but does not change its own state. If the owning cache has the

subpage in exclusive state, it changes its own
subpage state to non-exclusive and then supplies the
read-only copy to the requester. If the owning cache
has the subpage in atomic or transient-atomic state,
5 it supplies the subpage with that state and
invalidates its own copy.

When a cache requests exclusive-ownership,
the owning cache yields its copy of the subpage. If
the subpage is owned with non-exclusive state, it is
10 possible that there are read-only copies in other
caches. All such caches react to the
exclusive-ownership request by invalidating their
read-only copies.

When a cache acquires a subpage with an
15 exclusive-ownership state in order to satisfy a store
instruction, it does not grant ownership or a
read-only copy to another cache until the store
instruction is complete. This rule provides the
strongly ordered nature of the memory system, in that
20 it ensures readers of a memory location see
modifications in the order that they are made.

When a subpage is in atomic state it may
change to transient-atomic state, but it will never
change to any other state as a result of any load or
25 store instruction. If some other cache requests the
subpage, it will always obtain the subpage in atomic
or transient-atomic state. After the subpage has
been released to exclusive state, the transitions
between exclusive and non-exclusive state may again
30 occur, and read only copies may exist when the
subpage is owned non-exclusively.

When a particular subpage is in invalid
state in a particular cache (i.e., a descriptor is
already allocated, but the particular subpage is not

present), and a copy of that subpage is available on the information transfer level interconnection due to a request from some other cache, the cache with invalid state will acquire a read-only copy of the
5 subpage. The effect of this mechanism is to accelerate parallel computations, since it can remove the latency associated with requesting a copy of a subpage from another cache.

It is important to note that the basic
10 mechanisms provide the strongly ordered memory access model to programs that use simple load and store instructions. Programs may use the forms of the load, store, and prefetch instructions that request particular states in order to improve their
15 performance, and it is expected that in many cases compilers will perform the necessary analysis. However, this analysis is optional.

In simple transactions, subpage atomic state is used purely as a lock. The data in the subpage is
20 not relevant. Some of the more sophisticated forms of synchronization mechanisms make use of the data in a subpage held in atomic state. One technique is to use atomic state on a subpage as a lock on the data in that subpage. Programs take one or more subpages
25 into atomic state, manipulate their contents, and release them.

Data Copying Strategy

The interaction among states is a tradeoff
30 between time, cache space, and information transfer level bandwidth spent when a load waits for a copy from another cache and the time and bandwidth spent invalidating copies in other caches when a store or gsp instruction is executed. When there are many

read-only copies of a subpage in the system, then the chances that a read will find the data already in the local cache are increased. However, if there are any read-only copies in the system, then the owner must
5 first invalidate those copies (by sending a single message) when modifying the subpage. The following heuristics attempt to dynamically distinguish between multiple read/writer sharing from single read/writer access on a short-term basis.

10 Multiple read/writer sharing is multiple read-only copies with high temporal locality and write updates with lower temporal locality. Retaining read-only copies is most efficient since multiple copies are read multiple times between
15 updates. The modifying CEU can gain an exclusive-ownership state with a single information transfer level operation; it can then broadcast the changed data using the pstsp instruction to distribute the data to interested caches.

20 Single read/writer access is multiple read-only copies with low temporal locality and write updates with much higher locality. Retaining read-only copies is less efficient since copies are updated multiple times between reads. A single
25 read/write copy (exclusive-ownership state) does not require a information transfer level operation for writes.

The strategy for balancing these considerations is as follows:

30

- (1) When a copy of a subpage is sent across the information transfer levels to satisfy a request, any cache that has a descriptor for the page but no copy of

the subpage picks up a read-only copy from the message. This mechanism accelerates applications with high locality of reference.

- 5 (2) When a cache with non-exclusive or read-only state needs to update a subpage, every other cache with a copy of the subpage invalidates that copy and exclusive ownership moves to the
10 modifying cell. The modifying cell can use pstsp to broadcast the change. Any subsequent attempt by another cache to read the subpage can result in copies being distributed to interested caches.

15

Cache Layout

A preferred organization of a cache directory is shown in Figure 5. When a reference to an SVA is made, the cache must determine whether or
20 not it has the required information. This is accomplished by selecting a set within the cache, and then examining all the descriptors of that set. SVA[20:14] selects a set. In the general architecture, each of the descriptors in the selected
25 set is simultaneously compared against SVA[63:21]. In a preferred embodiment having a 240 Mbyte SVA space, this implies a comparison with SVA[39:21]. If one of the elements of the set is a descriptor for the desired page, the corresponding comparator will
30 indicate a match. The index in the set of the matching descriptor, concatenated with the set number, identifies a page in the cache.

If more than one descriptor matches, the cache signals a multiple descriptor_match exception.

If no descriptor matches, the cache allocates a descriptor and requests data from the interconnect. It is possible that either the allocation or data request will fail, and the cache will indicate an
5 error to the CEU.

The use of SVA[20:14] to select a set is effectively a hash function over SVA addresses. System software must assign SVA addresses so that this hash function gives good performance in common
10 cases. There are two important distribution cases: referencing many pages of a single segment and referencing the first page of many segments. This set selector produces good cache behavior for contiguous groups of pages, since 128 contiguous
15 pages will reside in 128 distinct sets. However, this selector will produce poor hashing behavior for many pages with the same value in SVA[20:14]. System software can avoid the latter situation by varying the logical origin of data within segments. For
20 example, each per-process user stack can be started at a different segment offset.

Contents of a Descriptor

When a cell responds to a request for a
25 subpage, it supplies the subpage data and the values of certain descriptor fields in the local cache. When the response returns to the requestor it either copies those fields to descriptor fields (if it has no other valid subpages) or logically ors those
30 fields into descriptor fields. Some descriptor fields are never supplied by the responder nor updated by the requestor.

In a preferred embodiment, the descriptor fields are defined as follows:

descriptor.tag (19 bits)

5 Bits [39:21] of an SVA. This field identifies the particular page of SVA space specified by the corresponding descriptor. For a given set in a given cell, this field must be unique among all 16 descriptors. Software 'sets' this field when it creates an SVA page. (It is also set by software during cache initialization.)

10

descriptor.atomic_modified (1 bit)

15 A cache sets this bit flag to one when any subpage of this page undergoes a transition into or out of atomic state because a gsp or rsp instruction was successfully executed. It is also set when a subpage changes from atomic state to transient-atomic state. This flag is not set if a gsp fails because the subpage is already in atomic state or if a rsp fails because the subpage was not in atomic state. This flag is not set if the gsp or rsp fails because descriptor.no_atomic is set. System software sets this flag to zero to indicate that it has noted the atomic state changes. This field is propagated from cache to cache.

20

25

30

35

descriptor.modified (1 bit)

40 A cache sets this bit flag to one when any data is modified in the page. System software sets descriptor.modified to zero to indicate that it has noted the modification of the page. This flag is not set if an attempt to modify data fails because descriptor.no_write is set. This field is propagated from cache to cache.

45

50

descriptor.LRU_Position (4 bits)
The cache maintains this field
as the current position of the
descriptor in its set from Most
Recently Used (0) to Least
Recently Used (15).

5

descriptor.anchor(1 bit)
Software sets the field to
indicate that data requests from
other caches may not be honored
nor may the descriptor be
invalidated. Any read or get
request from another cache
returns unresponded to the
requestor, and is treated as if
the page were missing. This
field is set by system software
as part of creating or
destroying an SVA page, and as
part of modifying the page
descriptor.

10

15

20

descriptor.held (1 bit)
Software sets the field to
indicate that the descriptor may
not be invalidated by the cache
even if no subpages are present
in the cache.

25

30

descriptor.no_atomic (1 bit)
Software sets this field to
prevent any cache from changing
the atomic state of any subpage
of this page. An attempt to
execute a gsp or rsp fails, and
is signalled back to the
processor. The processor
signals a page_no_atomic
exception. descriptor.no_atomic
can be altered even when some
subpages have atomic state. This
flag simply prevents attempt to
change atomic state, in the same
way that descriptor.no_write
simply prevents attempts to
change data state. This field
is propagated from cache to
cache.

35

40

45

50

descriptor.no_write (1 bit)
5 Software sets this field to prevent modifications to the page by the local processor. An attempt to modify the page fails, and is signalled back to the processor. The processor signals a page_no_write exception. This flag does not
10 affect the ability of any cache to acquire the subpage in exclusive or atomic/transient-atomic state. This field is propagated from
15 cache to cache.

descriptor.summary (3 bits)
20 Summarizes subpage state field of a set of subpages. There is one three-bit summary field for each set of subpages. The summary sometimes overrides the contents of the individual
25 subpage_state fields for the subpages within the summary set.

descriptor.subpage_state (4 bits)
30 The subpage state consists of a three-bit state field and a single bit subcached status field. It is set by the cache to record the state of each subpage and to indicate whether
35 or not any portion of the subpage is present in the CEU subcache.

It is possible for flags in descriptors to
40 be incorrect or out-of-date. There are two reasons for this: latency between a CEU and its local CCUS, and latency between the CCUs of different cells. The former can occur when a page which was previously unmodified is modified by the CEU. Until the
45 modified subblock leaves the subcache, the local CCUs are not aware that a modification has occurred. The

latter can occur because many caches may each contain a descriptor for a particular SVA page, and a state change affected by one cache is not automatically broadcast to all other caches.

5 Because there is no valid bit associated with a descriptor, a descriptor always has the tag of some SVA page. However, if all the cache descriptors which have tags denoting a particular SVA page have no valid subpages in their descriptors, then the SVA
10 page does not exist in the memory system. Similarly, if no cache descriptor has a tag denoting a particular SVA page, then that page does not exist in the memory system. It is possible to read the descriptor fields for such a page, but since the SVA
15 page has logically been destroyed, the field values are not valid.

For example, consider two caches with descriptors for a page; cache A has all the subpages in exclusive state, and descriptor.modified is clear;
20 no other cache has a descriptor for the page. The CEU of cell B executes a store instruction to modify subpage data. CEU B requests the subpage with an exclusive-ownership state from its local cache. The cache allocates a descriptor for the page and then
25 requests the subpage using the ring. The owner (cell A) responds by yielding exclusive ownership to cell B. After the subpage arrives, cell B copies descriptor.modified (which is clear) from the response. Then the CEU of B loads a subblock from
30 the subpage into its data subcache and modifies the subblock. At this point, the CEU subcache indicates that the subblock has been modified, but the local cache still shows the subpage as unmodified. At some later time, CEU B will send the subblock data from

its subcache to the local cache. This can occur because the CEU requires the subcache block for other data, because the CEU is using idle cycles to write-back modified subblocks, or because some other
5 cell has requested the subpage. Then the cell B cache sets the descriptor.modified flag. Throughout this time, descriptor.modified is clear on cell A.

As a second example, assume that the system software running on cell A now destroys the SVA
10 page. The page is destroyed by gathering up all subpages in exclusive state in cache A and then altering A's descriptor such that the SVA page no longer has any valid subpages. However, cache B still has a descriptor for the page. Fields such as
15 descriptor.modified are meaningless. Even if some other cache subsequently recreates the SVA page, cache B's descriptor will be out of date until the first subpage arrives.

System software must ensure that the SVA
20 page is really present in the memory system when it attempts to use descriptor information. One method of accomplishing this is to always set the page anchor and obtain at least one valid subpage in the local cache. To be absolutely certain that the
25 modified and atomic_modified fields are not set, software must first obtain every subpage in exclusive state. The anchor prevents any other cache from asynchronously acquiring any subpage already obtained. When the operation is complete, system
30 software clears descriptor.anchor in the local cache. Then other cache requests for subpages of that page will again be honored.

System software also needs to ensure that an SVA page does not really exist in the memory system

before it creates the page. As noted above, the simple existence of a descriptor with the correct tag value does not indicate that the SVA page actually exists in the memory system. Software can verify non-presence by setting the page anchor and then attempting to fetch a subpage with exclusive state. If the fetch succeeds, software has lost a race, and the page exists in the memory system. Otherwise, software can create the page using the mpdw instruction to establish the SVA page address with all subpages owned exclusively. Note that there still needs to be a software interlock against the simultaneous use of mpdw for this purpose.

15 Cache Descriptor Usage and Replacement

The caches of the illustrated system 10 can be used by system software as part of a multilevel storage system. In such a system, physical memory is multiplexed over a large address space via demand paging. The caches include features that accelerate the implementation of a multi-level storage system in which software moves data between the caches and secondary storage in units of SVA pages.

All of the caches together make up a system's primary storage. However, for some purposes, it is necessary to consider each individual cache as an independent primary store. This is because each cache can only hold a limited number of SVA pages: 2048 pages in each cache, and 16 in any particular set. There are several automatic and manual mechanisms to balance the utilization of the caches.

Each cache maintains LRU state for all of the resident pages. The LRU data is maintained

separately for each of the 128 sets of the descriptor associative memory, and orders the 16 pages in the set according to their approximate time of last reference.

5 Each cache maintains an LRU-MRU ordering of the descriptors in each set. The ordering is maintained in descriptor.LRU_priority. Each of the descriptors in a set has a value from 0 (MRU) to 15 (LRU) in descriptor.LRU_priority. Conceptually, when
10 a page is referenced it moves to MRU and all of the pages from MRU down to the referenced page's old LRU priority then move down one step towards LRU. A descriptor is made MRU:

- 15 (1) on any load from the local CEU,
 including increasing subcache state strength
- (2) on any store from the local CEU
- (3) on any local CEU writeback which
20 unsubcaches
- (4) following an XCACHE operation (the cache requests that the local CEU change subcache state)
- (5) after deallocation from a local CEU
25 subcache

A CEU instruction may reference an SVA page which is not present in the local cache (it is in invalid-descriptor state). The cache must allocate a
30 descriptor, as described below, and then issues a request to the interconnect for the particular subpage which is required by the CEU. If the corresponding page is resident elsewhere in the system, the requesting cache will copy certain

descriptor fields and the referenced subpage data from the responding cache. Eventually, this process of allocating descriptors will fill up the elements of a particular cache set.

- 5 Each cache has capabilities for automatically reusing certain descriptors by changing an SVA page to invalid-descriptor state in the cache. As a background activity, the caches also automatically move owned subpages to other caches.
- 10 All of these automatic actions can be tuned or disabled by system software, and are inhibited when descriptor.held and/or descriptor.anchor is set. The following subsections describe the descriptor reclamation and background data movement activities
- 15 of the caches.

Sometimes, there will be one or more descriptors in a cache set which have no valid subpages. This situation can arise as a result of requests generated by other caches, or as a result of

20 recombining activity (described below). For example, assume that a cache descriptor has only one valid subpage, for which it has exclusive state. If some other cache requests ownership of the subpage, this cache no longer has any valid subpages of the page.

25 If descriptor.held and descriptor.anchor are clear in this cache, then this descriptor can be reused when the CEU references some other page which has invalid-descriptor state in the same set of this cache.

- 30 A cache will automatically drop a page which has all subpages in read-only or invalid state (a read-only page), provided that no subpage is subcached. This is possible because no information is lost by the memory system when it destroys

read-only copies. The cache has configuration options to totally disable copy dropping or to restrict it to a partial range of descriptors, according to LRU ordering.

5 A cache will automatically drop a page which has all subpages in some exclusive-ownership state and which has `descriptor.modified` and `descriptor.atomic_modified` both clear. The fact that the cache has exclusive ownership of all subpages
10 guarantees that no other cache has ownership or a copy of some portion of the data state, and allows this cache to destroy data state without communicating with any other cache. The fact that `descriptor.modified` and `descriptor.atomic_modified`
15 are both clear indicates that no data or atomic state change (including transition from atomic state to transient-atomic state as the result of a `gsp.wt` instruction) has occurred since the SVA page was created or last polled by system software. The memory
20 system presumes that system software has a copy of the page data and atomic state on secondary storage which can be used to recreate the SVA page, and destroys the page. The cache has configuration options to totally disable pure page dropping or to
25 restrict it to a partial range of descriptors, according to LRU ordering. Note that system software must never asynchronously clear `descriptor.atomic_modified` if any subpage is in transient-atomic state.

30 Note that setting `descriptor.held` does not guarantee that the individual subpages will be present in the local cache. Further, simply setting `descriptor.held` for a page does not prevent some other cache from destroying the page. System

software must take explicit action to prevent hardware from automatically dropping a pure page (e.g. set descriptor.modified).

Recombining is an operation whereby an owned
5 subpage is expelled from a cache onto the interconnect and other caches are given an opportunity to accept ownership state. Any accepting cache must already have a descriptor allocated for the containing page; a cache will not allocate a
10 descriptor in order to accept the ownership of a subpage being recombined. The goal of page recombining is to reduce the total number of descriptors which are allocated for a particular SVA page.

15 A cache issues a recombine message to recombine a subpage. If the recombine messages fail to find another cache to take over the page, the recombining cache retains the data. In effect, it has found itself as the target of the recombine. If
20 some other cache accepts ownership (which may be non-exclusive, exclusive, atomic or transient-atomic state), the issuing cache changes its subpage state to invalid.

The caches will automatically attempt to
25 recombine subpages as a background activity while memory refresh is occurring. At each refresh interval, the cache examines a particular set, and searches for an acceptable descriptor. The descriptor must have no subcached subpages, must own
30 some subpage, and must not have all subpages in exclusive-ownership state. If such a page is found, the cache issues a recombine message for some owned subpage. The cache has configuration options to totally disable background recombining or to restrict

it to a partial range of descriptors, according to LRU ordering. Background recombining makes it more likely that the cache will be able to allocate a descriptor for a newly referenced SVA page, instead
5 of causing a line-full error.

The CEU also provides an instruction whereby system software may attempt the expulsion of an owned subpage. The subpage must not be subcached, but no other restrictions are imposed before the cache
10 attempts to expel ownership.

Since pages below the working set point are less likely to be referenced, most of the recombines that actually move data will move it from a cell which is not referencing the data to a cell which is.
15 When a new descriptor in a set is needed, the cache proceeds in order through as many of the following actions as needed to find a usable descriptor:

- 20 (1) find an invalid descriptor
- (2) invalidate a read-only copy
- (3) destroy a pure SVA page
- (4) Signals a line_full exception if it
cannot allocate a descriptor by the
25 means described above.

The individual steps are explained below:

Find an invalid descriptor

30 If an invalid descriptor exists it can be used immediately. This requires:

- all subpages are in invalid state, and
 - descriptor.held and descriptor.anchor are both clear
 - no subpage of the page is described by a PRT entry
- 5

Invalidate a read-only copy

If `ctl$ccu_lru_config.cde` is 1, the cache attempts to identify a descriptor which contains only 10 read-only copies of subpages. It searches from LRU to MRU, looking for any page which has:

- all subpages in read-only or invalid state, and
- 15 - no subcached subpages, and
- an LRU value which is greater or equal to `ctl$ccu_lru_config.cdl`, and
- descriptor.held and descriptor.anchor both clear, and
- 20 - no subpage of the page is described by a PRT entry

If an acceptable descriptor is found, all subpages are changed to invalid state and the 25 descriptor is used.

Drop a pure SVA page

If `ctl$ccu_lru_config.pde` is 1, the cache attempts to identify an SVA page which can be 30 destroyed (removed entirely from the memory system). It searches from LRU to MRU, looking for a page which has:

- all subpages in (various)
exclusive-ownership states, and
- no subcached subpages, and
- descriptor.modified and
5 descriptor.atomic modified both clear,
and
- LRU value which is greater or equal to
ctl\$ccu_lru_config.pdl, and
- descriptor.held and descriptor.anchor
10 both clear, and
- no subpage of the page is described by
a PRT entry

If an acceptable descriptor is found, all
15 subpages are changed to invalid state (thereby
destroying the SVA page), and the descriptor is used.

To reduce the probability that a line_full
exception will occur, the cache periodically attempts
to reduce its ownership of subpages within
20 less-recently-used pages. These actions are
independent of descriptor allocation operations, and
occur during RAM refresh.

If ctl\$ccu_lru_config.bre is set, the CCU
searches up from LRU looking for any page which has:

- 25
- some subpages in ownership state
- some subpage not in an
exclusive-ownership state
- no subcached subpages
- 30
- an LRU value with is greater or equal
to ctl\$ccu_lru_config.br1
- descriptor.held and descriptor.anchor
both clear

- no subpage of the page is described by a PRT entry

If an acceptable descriptor is found, the
5 cache issues a recombine message to the interconnect.

System software can monitor the LRU behavior of pages within individual caches to establish a 'working-set point' for each cache line. Descriptors with lower LRU values are in the working set and
10 descriptors with higher LRU values are not in the working set. Repeated examination of the descriptor tags, LRU values and subpage state can be used to determine the working-set point.

When a modified (or atomic_modified) page
15 leaves the logical working set, system software can take action to purify the page so that the descriptor will be reallocatable in the near future.

If the descriptor holds relatively few subpages in ownership states, then software can
20 attempt to recombine those subpages to other caches using the mrcsp (recombine-subpage) instruction. Otherwise, it might decide to obtain the entire SVA page and purify or destroy it.

As the working set changes, system software
25 can modify the configuration parameters which control the cache. The parameters for pure-dropping, copy-dropping, and background recombines can be altered to ensure that hardware actions are consistent with software working-set policy.

30 The typical use of secondary storage is that each SVA page has a corresponding page allocated in secondary storage, extending the caching model provided by hardware. A page originally 'exists' only in secondary storage, and is encached in the

memory system when it is referenced by a program. If the page is modified (data state or atomic state) while in the memory system, then the copy of the page on secondary storage must be updated before the
5 containing memory system page can be re-used for some other SVA. Such modifications can be detected by examining the `descriptor.modified` and `descriptor.atomic_modified` fields. The operation of updating secondary storage is call 'purifying' the
10 SVA page. System software typically purifies SVA pages which are below the working set point, so that the cache space will be quickly available for other SVA pages.

Only some of the cells in the system have a
15 physical connection to the disk drive which provides secondary storage for an SVA page. The appropriate I/O-capable cell must be able to obtain page data during the I/O operation. System software can ensure this by allocating and holding a descriptor on the
20 I/O-cell before issuing the I/O commands. After the purification is complete, system software can clear the `descriptor.modified` and `descriptor.atomic_modified` flags or destroy the SVA page. The actual protocols used can be much more
25 complex, depending upon whether or not the system allows programs to access the page during the purification attempt.

In general, system software must not clear `descriptor.atomic_modified` if any subpage has
30 transient-atomic state unless the memory system is otherwise prevented from destroying the page. e.g. if all descriptors have the held flag set, or if `descriptor.modified` is set, the memory system will not destroy the page.

When a particular set of a particular cache is full, there is no reason for software to assume that the entire memory system is correspondingly full. Thus, it is desirable for software to respond
5 to a full set by moving a page from that set to the corresponding set of another cache.

In order to use memory efficiently, software must use some strategy for identifying an appropriate page to remove from a full set, and an appropriate
10 target cache, if any, for the page. The caches include facilities that accelerate a class of strategies for this page replacement. To accelerate software's selection of a page for replacement within a set, each cache approximately orders the pages from
15 Most Recently Used (MRU) to Least Recently Used (LRU). When a page is referenced, it moves to MRU. As other pages are referenced thereafter, it ages toward LRU. The LRU information accelerates strategies that replace the least recently used page.

20

Processor Data Accesses

A processor makes data requests to its local cache to satisfy load and store instructions and co-execution unit operations. A cache makes requests
25 to its local processor to force the processor to invalidate its copy of a subpage in subcache.

Load and Store Instructions

A processor passes load and store
30 instructions to its local cache as requests when the subblock containing the referenced address is not present in the subcache in the required state. The different types of load and store instructions pass information to the local cache about the access

patterns of the following instructions. For example, if the sequence of the instructions is a load followed by a store, and the subpage containing the data item is not yet resident in the local cache, it
5 is more efficient to acquire ownership for the load than to get a read-only copy for the load instruction and then communicate over the information transfer levels a second time to acquire ownership for the store instruction.

10 The state of subblocks in the subcaches does not always reflect the state of the corresponding subpage in the cache. The instruction subcache always obtains a read-only copy of data. The data
15 subcache may hold a subblock in read-only or exclusive state. The subcache can only have exclusive state if the cache has an exclusive-ownership state and descriptor.no_write is not set. The subcache does not distinguish between
20 exclusive, atomic, and transient-atomic subpage states. When the subcache has a subblock in exclusive state, the CEU can execute store instructions by simply placing the new data in the subcache. For all store instructions except st64.nsc, if the subblock is not described by the subcache, or
25 has invalid or read-only state, then the CEU must request exclusive state from the local cache prior to completing a store instruction. If descriptor.no_write is set or the subpage is not present in the memory system, a fault will occur.

30 When a request for a subpage arrives from another cache, the owning cache must respond. If any part of the subpage is in the data subcache, the local cache must ensure that it obtains any modifications which might be present only in the

subcache. The cache also causes the CEU to change subcache state for the subblock to read-only or invalid, depending upon the request. In certain cases, the cache will also ensure that the
5 instruction subcache invalidates its read-only copy of the subpage.

It is important to distinguish between the units of subcache management (blocks and subblocks) and the units of cache management (pages and
10 subpages). Data travels between a CEU and its local cache in subblocks. Data travels between caches in subpages. There are two subblocks per subpage.

The different forms of load and store instructions are described below. Each description
15 begins with a brief summary of the semantics of the instruction, and continues with an overview of subcache and cache actions.

load (read_only) [ld.ro]
20 load_64 (read_only, subcached) [ld64.ro.sc]
The program will continue a pattern of reading data. The minimum amount of work is done to obtain data. If the
25 containing subblock is subcached, it is used directly. If the local cache does not have the subpage, it obtains a copy. The local cache supplies the
30 subblock to the subcache with exclusive or read-only state, as appropriate.

load (exclusive) [ld.ex]
35 load_64 (exclusive, subcached) [ld64.ex.sc]
The program will write the subblock in the following instructions, and exclusive state is preferable to any other
40 state. A program would use this when the data was expected to

5 have little sharing, or when a series of writes was upcoming. This can reduce the number of interconnect messages required before the CEU can modify data.

10 A particular example of the use of load (exclusive) is per-program data such as stacks. Generally, there will be no read-only copies of such data, since the only copy will be the one in use by the program. However, if a program moves from one processor to another, the new processor's local cache will have no copy, and the old processor's local cache will continue to hold the subpage in an exclusive-ownership state. If the program uses load (read_only), the local cache acquires the subpage in read-only state (unless the subpage is in atomic or transient-atomic state, in which case that state is acquired). The subsequent store requires the cache to make another interconnect request (to obtain an exclusive-ownership state) before any CEU data modification can occur.

35 As with ld.ro, the minimum amount of work is done to obtain data. If the subblock is already present in the subcache, it is used directly. If the local cache does not have the subpage, it requests the subpage in an exclusive-ownership state. When the local cache has the subpage, the subblock is supplied to the CEU in read-only or exclusive state, as appropriate.

40

45

store [st]
store_64 (subcached) [st64.sc]
5 If the subblock is already present in the subcache in exclusive state, subcache state is unchanged; data is written to the subcache.

10 The subcache must have the subblock in exclusive state. As necessary, the subcache will request exclusive state from the local cache, and the local cache will request an exclusive-ownership state from the interconnect.

15

20 If the descriptor.no_write flag is set, an error is signalled to the CEU which generates a page_no_write exception. Otherwise, the subcache obtains the subblock in exclusive state; data is written to the subcache.

25 load_64 (read_only, unsubcached) [ld64.ro.nsc]
load_64 (exclusive, unsubcached) [ld64.ex.nsc]
30 The programmer uses the exclusive and read_only designations according to the expected reference pattern, as documented for ld. However, the number of references to the subblock is expected to be
35 small, and the subcache should not be disturbed while fetching this data.

40 If the data is present in the subcache, it is used directly. If the local cache does not have the subpage, it obtains a copy. The CEU obtains a copy of the data and loads the destination
45 registers.

store_64 (unsubcached) [st64.nsc]
50 The number of references to the subblock is expected to be small (typically one), and the

subcache should not be disturbed while storing this data.

5 If the subblock is subcached in exclusive state, the subcache state is unchanged and data is written to the subcache. If the subpage is subcached in
10 read-only state it is immediately invalidated. The CEU supplies the data to the local cache. If the cache does not have the subpage in an
15 exclusive-ownership state, it requests it from the interconnect.

20 If the descriptor.no_write flag is set, an error is signalled to the CEU which generates a page_no_write exception. Otherwise, the CEU data is written directly to the subpage in the cache.

25 instruction fetch Instruction fetches always fetch subpages specifying read-only state.

30 Subpage Atomic State Instructions

The subpage atomic instructions are the program interface to the get and release operations described above. These instructions exist in several
35 forms to permit precise tuning of parallel programs.

get subpage [gsp.nwt]
get subpage & wait [gsp.wt]
40 Get subpage requests that a subpage be set into atomic state. For both forms of the get subpage instruction, if the subpage is not in atomic state in any cache, then the local
45 cache acquires it in atomic state.

5 For gsp.nwt, the @MEM condition code indicates the success or failure of the attempt; the instruction will trap instead of changing @MEM if the trap option is present in the instruction and the subpage is already atomic.

10 The gsp.wt instruction form causes the cache to stall the CEU until the subpage can be obtained in atomic state. This reduces the amount of
15 interconnect traffic if the program must obtain atomic state before it can proceed. If the subpage is already atomic in any cache (including the local
20 cache), the instruction waits until the subpage is released. The local cache then acquires the subpage in atomic state. The @MEM condition is always
25 changed to indicate success.

release subpage [rsp]

30 Release subpage is used to remove a subpage from atomic state. If the subpage is not present in the local cache, it is first requested in via the interconnect. Once the local cache has exclusive ownership,
35 rsp proceeds. If the subpage is not in atomic state then release subpage does not change the subpage state. In this situation, the CEU will trap if
40 the trap modifier is present for the instruction. If the subpage is in atomic state, it is changed to exclusive state. If the subpage is in
45 transient-atomic state, it is changed to exclusive state and expelled onto the interconnect so that any waiting cell may acquire atomic state.
50

Other Subpage Instructions

Post Store Subpage[pstsp]

5 Post-store subpage allows the
program to expel a read-only
copy of a subpage onto the
interconnect. All caches with
descriptors for the page will
10 take a copy of the data. This
instruction can be used to
broadcast data as part of
completing some operation,
reducing the probability that
some other cache will have to
15 make a read request on the
interconnect when it needs to
use the data.

Prefetch Subpage [pcsp]

20 Prefetch Subpage requests that a
copy of a subpage be acquired on
the local cache in a specified
state. The instruction can
request read-only or exclusive
25 state. A subsequent reference
to the subpage blocks until the
prefetch subpage has completed.

30 'Manual' Control of the Memory System

As described above, the memory system is
designed to support a virtual memory system with
automatic data sharing and LRU maintenance. However,
software can take explicit control of the memory
35 system for special applications.

In normal use, all processors share SVA
space, and data automatically moves from cache to
cache in response to instructions (and control
operations). Software can dedicate some or all of
40 the memory on a cache to its local, unshared use.
Such a system must partition SVA space among the
caches, and use explicit control operations to move
such data from cache to cache.

By setting descriptor.held in every descriptor, system software can prevent the cache from ever moving or destroying a page to make room for another page. The system software can then
5 handle exceptions or perform explicit destroys as needed to multiplex each cache's memory.

In automatic mode, the memory system can be configured as a shared memory multiprocessor. When various automatic features are disabled, the memory
10 system can be configured to emulate more loosely coupled message-oriented architectures. Messages can be passed by references to special SVA ranges. Manual control of the memory system can be used to enforce a specific memory model more closely.

15

Memory System Control Instructions

Control operation permit the processor to directly manipulate the memory system. There are two classes of control instructions: data movement and
20 page state control. The data movement control instructions move pages and subpages of data from cache to cache in the hierarchy. The page state control instructions manipulate page descriptors.

CEU instructions result in cache commands
25 which execute synchronously or asynchronously, depending upon the command. A CEU cache instruction occupies an entry in the cache PRT (a hardware table) while it is in progress. The PRT has four entries, so a maximum of four cache instructions may execute
30 in parallel. Most CEU instructions result in assignment of a PRT entry which remains in use until the request is satisfied, providing synchronous behavior. For example, load/store instructions execute synchronously, so that certain

software-controlled exceptions (such as missing page or unwriteable page) can be predictably resolved.

The pcsp (prefetch-cache-subpage) and pstsp (post-store-subpage) instructions operate

5 asynchronously, as described in following subsections.

Synchronous errors typically result in the CEU executing the trap sequence.

Asynchronous errors result from actual hardware errors or are provoked by a request from
10 some other cache. Such errors are reported by memory-system interrupts.

Prefetch Instruction

The prefetch instructions request that a
15 copy of a subpage be acquired on the local cache in a specified state. Pcsp prefetches a subpage. The cache allocates a PRT entry when this instruction is detected. If the subpage is already present, the PRT entry is freed and the pcsp completes. Otherwise,
20 the cache issues a request, and then indicates instruction completion to the CEU, which proceeds asynchronously. When the message returns as a request or response, the cache accepts data (if present), and frees the PRT entry. There is no
25 indication to the CEU that the data has arrived.

Post-Store Subpage Instruction

The pstsp instruction requests that a copy, of a subpage be circulated on the interconnect so
30 that any caches having a descriptor for the containing page may acquire a read-only copy of the subpage. pstsp references a subblock within a subpage. If the subblock is subcached with exclusive state and is modified in the subcache, the CEU

requests a post-store action from the local cache; otherwise, the pstsp instruction has no effect. The cache allocates a PRT entry, and requests the subpage data from the CEU. The cache then submits the
5 post-store message to the interconnect, frees the PRT entry, and indicates instruction completion to the CEU. The CEU proceeds asynchronously. When the message returns to the issuing cache, it is discarded.

10 Fetch Subpage Instruction

The mfsva instructions permits system software to fetch a subpage in read-only or exclusive-ownership state, specifying the SVA location of the subpage. This saves system software
15 the effort of establishing a DSTT translation, as is required by pcsp.

Flush Subcached Subpage Instruction

The mflsp instruction causes the cache to
20 ensure that the specified subpage is not subcached in the local CEU. If the subpage is in invalid-descriptor state or invalid state, no descriptor is allocated and the subpage is not requested via the interconnect.

25

Recombine Subpage Instruction

The mrcsp instruction allows system software to reduce the number of active descriptors for a page by causing ownership to migrate to another cache.
30 Unlike the background recombine activity of the cache, this instruction is not controlled by cache configuration parameters.

The page state control instructions operate on individual pages of SVA space.

35

Anchor Descriptor Instruction

The mpsa instruction provides an anchored descriptor in the local cache for an SVA page. If the descriptor already existed prior to the mpsa, its anchor flag is set. Otherwise, the cache allocates a descriptor and then sets the anchor flag. Page state control operations require that an anchored descriptor for the SVA page be present on the local cache.

10

Write Descriptor Instruction

The mpdw instruction is used to create and destroy SVA pages, and to change descriptor flags of existing SVA pages. mpdw requires that system software first obtain an anchored descriptor for the page, using the mpsa instruction. The following discussion assumes that an anchored descriptor exists on the local cache.

20

Creating an SVA page

Following the mpsa, the descriptor exists, but all subpages are in invalid state. System software executes mpdw specifying that all subpage states should be set to exclusive. This causes a message to be sent on the interconnect so that any interested ring members may note the creation of the page.

The SVA page now exists, although its data values are undefined. Software must initialize the page using store instructions or l/0 before allowing the user to reference the page. For this reason, software will typically create a page at an SVA location inaccessible to user programs, initialize the page data, and then change the address of the SVA

page as described below. The page is released for general use by executing an mpdw instruction which clears the anchor.

5 Destroying an SVA page

After the mpsa, system software must obtain all subpages in exclusive state. This is done using mfsva instructions. Software then executes mpdw specifying that all subpages should be changed to
10 invalid state. This instruction causes a message to be sent on the interconnect so that any interested ring members may note the destruction of the page. The SVA page is destroyed by this operation. Software releases the descriptor for reuse by
15 executing a second mpdw which clears the anchor.

Change Descriptor Fields

The mpdw instruction is used to change various fields in a local descriptor. It can set or
20 clear the modified, atomic_modified, no_write, no_atomic, and held fields and can clear the anchor field. mpdw can also change the tag, and thus the SVA space address associated with the descriptor. (Since the index of the descriptor forms part of the
25 SVA, the new tag is in the same cache set, by definition.)

To ensure memory system consistency, system software must obey certain rules when altering the fields or tag of a descriptor. mpdw requires that
30 descriptor.anchor be set (although the instruction itself may result in clearing descriptor.anchor). Various sequences require that all subpages be present in the local cache with an exclusive-ownership state. This is accomplished by

setting descriptor.anchor and executing mfsva.ex for each subpage. Various sequences require that all subpages be unsubcached in the local cache. This is accomplished by executing mflsp for each subpage
5 which might be subcached in the local CEU.
(Executing mfsva.ex ensures that a subpage is not subcached in by the CEU of any other cell.)

The following list give the restrictions which are in force for each flag and the tag.

10

(1) anchor is set and typically cleared as part of any descriptor tag or flag modification. It may be left set for any duration, but in a shared-memory
15 system, it is cleared as soon as possible.

(2) held can be modified to hold or unhold a descriptor on a particular cell without restriction.

20

(3) changing atomic_modified and no_atomic requires that all subpages be in an exclusive-ownership state in the local cache.

25

(4) clearing modified and setting no_write require that all subpages be in unsubcached in the local cache. This ensures that the local subcache does not have any subblock in exclusive state nor does it have any modified
30 subblock. When changing modified no_write, system software may decide whether or not the change should be perceived by all cells which reference the page. Effecting a global change

- requires that all subpages be in an exclusive-ownership state in the local cache. System software will typically make a global change when it clears the modified flag. Effecting a local change does not require that any subpage be present in an exclusive-ownership state; however, this results in delayed perception of the new state.
- 5
- 10 (5) changing the SVA page number by changing the tag requires that all subpages be in an exclusive-ownership state and unsubcached on the executing cell.
- 15
- Changing single bit fields is accomplished with a single mpdw.desc instruction. This instruction contains the new value of the changed flag, the old values for the other flags and the
- 20 tag. Unless system software has some special reason to keep the page anchored, it clears the anchor flag.
- Changing the SVA page number of a descriptor is logically the same as destroying the old page and then creating a new page which happens to have the
- 25 same data. The sequence is:
- (1) anchor a descriptor for the old SVA page and obtain each subpage in an exclusive-ownership state. If any
- 30 subpage has atomic or transient-atomic state, it will be acquired by the executing cell. Once all subpages have been acquired, any access to the old SVA page will take a missing_page fault.

- 5 (2) determine the atomic state of each
subpage. This is most rapidly
accomplished by executing gsp.nwt on
each subpage, and examining the
resulting @MEM indicator. Any cell
which already executing a gsp.wt for a
subpage of the old SVA will eventually
take a timer interrupt; when the gsp.wt
is restarted, it will take a
10 missing_page fault.
- 15 (3) use mpdw.alli to change all subpages to
invalid state. This instruction causes
the local CCUs to inform the local CIUs
and the RRC (if any) that the SVA page
is being destroyed. Even though all
subpages are changed to invalid state,
the data remains in the local CCUS.
- 20 (4) use mpdw.desc to change the tag and set
flags to desired state. The new anchor
flag must be set.
- 25 (5) use mpdw.allx to change all subpages to
exclusive state. This instruction
causes the local CCUs to inform the
local CIUs and the RRC (if any) that an
SVA page is being created. The old
data is now perceived to be present in
exclusive state.
- 30 (6) restore saved atomic state of each
subpage. For each subpage which was in
atomic state, issue a gsp.nwt.
- (7) use mpdw.desc to clear the anchor flag.

The uses of descriptor.no_write include:
prevention of inadvertent modification of certain

data; support for copy_on_write/copy_on_access protocols; debugger watchpoints. In the first case, no_write is set and remains set. In the second case, when a program attempts to modify a page, system
5 software can respond by making a copy of the page available to the other users and then clearing no_write. Software can make this change local or global; in the former case, no_write faults may continue to occur on other cells which reference the
10 page. Finally, debugger watchpoints are intended to detect the modification of a particular region of context address space, e.g., find out where a global variable is being destroyed. System software can implement this by setting no_write for the page, and
15 trapping every modification attempt. For modifications outside the watched range, system software can anchor the page, clear no_write, modify the data, uncache the data, set no_write, and proceed. It is also possible to implement watchpoint
20 support by making global changes to no_write.

Find Page Local - LRU Lookup Instruction

The mfpl instruction searches a specified set in the LRU space of the cache for a descriptor
25 that matches a specified set of criteria. The search begins with the descriptor at LRU position 15, and proceeds upward until the criteria are met.

Instruction and Data System Address Space Split

30 To ensure correct operation, the cache must know when a subpage is also present in a CEU subcache. This permits the cache to request subcache invalidation when some request from the interconnect

(or even the local CEU) requires a subpage state change.

As part of this mechanism, the CEU communicates with the cache when it changes subcaching status. However, the cache does not maintain per-subcache information. As a consequence, the same subpage of SVA space must not simultaneously appear in both subcaches. In general, this results in a system software restriction that the same SVA region cannot be used as both instructions and data. Self-modifying programs or programs where code and data are part of the same context address segment are not supported.

System software must take special care in the following cases:

- (1) changing an instruction, as when inserting or removing a breakpoint
- (2) reading an instruction, as part of trap analysis or program disassembly by a debugger
- (3) reading a page from an I/O device which will become an instruction page

To read an instruction as data, system software must:

- (1) construct a DSTT entry which describes the SVA
- (2) ensure the subpage is not present in the instruction subcache (use mflsp)
- (3) read the instruction as data (use ld64.nsc)
- (4) invalidate the DSTT entry.

To write an instruction subpage as data,
system software must:

- 5 (1) construct a DSTT entry which describes
 the SVA
- (2) ensure the subpage is not present in
 any instruction subcache (use mpsa to
 anchor the page, and mfsva.ex to
10 invalidate the subpage in all other
 caches and subcaches)
- (3) fetch the containing subblock (use
 ld64.nsc)
- (4) modify the subblock and write the
 instruction as data (use st64.nsc)
- 15 (5) release the page anchor.
- (6) invalidate the DSTT entry.

Instruction pages are typically pure and do
not need to be written from SVA space to an I/O
20 device as part of system SVA space management.
Before writing an instruction page, system software
must:

- 25 (1) ensure the subpage is not present in
 any instruction subcache (use mpsa to
 anchor the page, and mfsva.ex to
 invalidate the subpage in all other
 caches and subcaches). If this cell
 will never execute this instruction
30 page, this step is not required.
- (2) perform the I/O
- (3) ensure that the page is not present in
 the data subcache (use mflsp). If this
 cell will never execute this

instruction page, this step is not required.

(4) release the page anchor.

5 When reading an instruction page from an I/O device into SVA space, system software must:

- 10 (1) create the page (use mpsa to allocate and anchor the page, mpdw to complete page creation)
- (2) perform the I/O
- (3) ensure that the page is not present in the data subcache (use mflsp). If this cell will never execute this
- 15 instruction page, this step is not required.
- (4) clear the descriptor modified flag, set any other descriptor attributes such as subpage atomic state, and
- 20 descriptor.no_write.
- (5) release the page anchor.

RRC Organization and Data Structures

A preferred RRC constructed in accord with

25 the invention is organized into two independent banks, with each bank connected to a single subring. The banks of one to four RRCs are interleaved based on low order page address and the configured degree of interleaving. Interleaving is statically

30 configured during system configuration. An RRC connected to Ring:0 is referred to as RRC:0 and a RRC connected to Ring:1 is referred to as RRC:1.

For a single Ring:1 configuration both the RRC:0 and RRC:1 are interleaved based on SVA[14]. Figure 19a depicts one preferred such configuration.

For a dual Ring:1 configuration with one
5 RRC:0 (TwoRing:1/1), the single RRC:0 is interleaved based on SVA[14] and two Ring:1 RRCs are interleaved based on SVA[15:14]. Figure 19B depicts one preferred such configuration.

For a dual Ring:1 configuration with two
10 RRC:0s (TwoRing:1/2), the RRC:0s are interleaved based on SVA[15:14] and two RRC:1s are interleaved based on SVA [15:14]. Figure 19C depicts one preferred such configuration.

For a four Ring:1 configuration, the two
15 RRC:0s are interleaved based on SVA[15:14] and four Ring:1 RRCs are interleaved based on SVA[15:14]. Figure 19D depicts one preferred such configuration.

Multiple RRC:0s are on a single Ring:0 for TwoRing:1/2 and FourRing:1 configurations. Multiple
20 RRC:1s are configured on separate Ring:1s for TwoRing:1/1, TwoRing:1/2 and FourRing:1 configurations.

Figure 6 depicts a preferred Ring Routing Cell Pair (RRC Pair or RRCP) constructed in accord
25 with the invention. The pair comprises RRC:0 and RRC:1 which are connected by a single communication line, e.g., bus, optical transmission connection, referred to as the InterRingLink. The downlink RRC is the opposite RRC of an RRC Pair. RRC Pairs which
30 are connected to the same Ring:0 are referred to as an RRC Array (RRC Array). When discussing a cell or a Ring:0, the local RRC or RRC pair are the RRC(s) or RRC Pair(s) which are co-resident on the same

Ring:0. A remote RRC or RRC pair are all those RRCs which are not local.

Each RRC bank consists of three sections, Routing Interconnect, Routing Directory and
5 InterRingBuffer. The Routing Interconnect implements the RRC portion of ring interconnect, provides an interface to the Routing Directory and interfaces the ring to the RRC internal RoutingBus by means of insert and extract buffers. The RoutingBus consists
10 of two 32 bit paths. Packets from the extract buffer are transferred to the InterRingDram Buffer or IBU over the RoutingToLink Bus. Packets are transferred from the IBU to the RI insert buffer over the RoutingToRing Bus. The IBU also transfers packets to
15 the InterRingDramBuffer for error logging over the RoutingToLink Bus.

The Routing Directory contains an entry for all allocated pages within the Ring:0 which the RRC pair is connected. The Routing Directory consists of
20 an Routing Directory Unit (RDU) and the actual ram based directory. The RDU controls the directory access sequence and contains the routing rules. For each packet, the Routing Interconnect passes pkt.addr, pkt.command and fifo status to the Routing
25 directory. The Routing Directory returns a routing command, and instructions to update packet command to the Routing Interconnect. Ring related interleaving is specified by RduMasterConfig location.

The InterRingBuffer buffers incoming packets
30 from the Ring Interconnect extract buffer and interfaces to the InterRingLink. The InterRingBuffer consists of the InterRing BufferUnit (IBU), InterRing Dram Buffer (IDB) and InterRingLink Interface. The

IDB buffers all packets from the Routing Interconnect until the packet has been transmitted over InterRingLink to the downlink RRC and an acknowledge has been returned. The IDB also contains the RRC ErrorQueue. The IBU buffers and formats outgoing InterRingLink packets, provides error checking and buffer for InterRingLink incoming packets and controls the IDB. The InterRingLink interface converts the 32 bit wide IBU outgoing packet and incoming packet interfaces to or from the physical link format.

RRC banks can be configured to connect to separate InterRingLinks or share a single InterRingLink. In dual InterRingLink configuration, the InterRingLinks are completely independent. In single InterRingLink configuration, both RRC banks share a single InterRingLink. InterRingLinkIn is connected to the IBUs of both RRC banks. The RRC banks arbitrate to share InterRingLinkOut. The IBU interleave configurations and packet address determine which bank incoming packets are routed to.

RRC conforms to the UCS system fault detection strategy. All buses are parity or ecc protected. Parity or ECC is passed through chips, rather than regenerated whenever possible. All dram and large sram data structures are ecc protected. Single bit soft or hard errors are corrected without causing a fatal system error (requires system reboot) and in a software transparent manner if possible. Data links expecting high error rates (InterRingLink) provide software transparent frame level operation. Frame level operation ensures error free transmission of frames, including retransmission when errors are detected.

Routing Interconnect

RI Structure

Each Routing Interconnect (RI) interfaces a
5 RRC bank to a subring, RRC internal Routing Bus and
Routing Directory. The RI is implemented as two
Routing Interconnect Units (RIUs). A preferred RIU
constructed in accord with the invention is
illustrated in Figure 7. Each RIU interfaces to half
10 of 16 byte wide ring and to half of 8 byte wide
routing bus. The two RIUs pass the address and
command field of packet to Routing Directory and the
Routing Directory returns a packet routing command,
new packet command and error status.

15 Each packet consists of ten, 16 byte double
words (20 words) while on the ring. Each long packet
(contains data) is stored in 9 words in each RIU
extract buffer and 10 words in insert buffer. Each
short packet (no data) is stored in 1 word in each
20 RIU extract buffer and 2 words in insert buffer.
RIU0 and RIU1 interface to RingData[63:0]
RingData[127:64] respectively. RIU0 and RIU1 extract
buffers interface to RoutingToLink[15:0] and
RoutingToLink [31:16] respectively. RIU0 and RIU1
25 insert buffers interface to RoutingToRing[15:0] and
RoutingToRing[31:16] respectively.

The RingIn to RingOut portion of RIU
consists of a 6 stage shift register. Packets can be
copied or extract from the shift register into the
30 extract buffer and packets from the insert buffer can
be inserted into the shift register. The insert and
extract buffers are managed with a fifo discipline.

RI Operations

The Ring Interconnect routes from two sources to two destinations. The sources are previous Ring Interconnect and Insert Buffer.

5 Packets from previous Ring Interconnect are referred to as **Ring Packets** and packets from insert buffer are referred to as **Insert Packets**. The destinations are next Ring Interconnect and Extract Buffer. Packets going to next Ring Interconnect are referred to as
10 **Ring Packets** and packets going to extract buffer are referred to as **Extract Packets**. Insert packets are supplied to the insert buffer and extract packets are sourced from the extract buffer by the InterRing Buffer Unit (IBU).

15 The address and command fields of the packet are passed to the Routing Directory and the Routing Directory returns a packet routing command. The RI buffers all packets until the packet routing command is available. The assertion of RingEmpty field
20 indicates the Ring packet is empty. Non-empty ring packets are always processed instead of insert packets. If the Ring packet is Empty and the insert buffer is not empty, the next insert buffer packet is processed by the Ring Interconnect.

25 The command from Routing Directory consists of a packet routing, packet field , extract buffer, and insert buffer commands.

The packet routing commands for all packets are pass and empty. The packet pass command
30 specifies the Ring Interconnect should pass throughout the packet unaltered to the next Ring Interconnect. The packet empty command specifies that the Ring Interconnect empties the packet by setting RingEmpty signal and packet command CmdEmpty

field. All packets for which CmdRequestId field match Ring Interconnect CellNumber field are unconditionally emptied, the Routing Directory command is not required. This allows RingEmpty
5 signal to be set earlier than possible when waiting for command from Routing Directory.

The extract buffer commands are no_push and push. The no_push command specifies that the extract buffer does not store the current packet. The push
10 command specifies that the current ring packet is copied into the extract buffer.

The insert buffer commands are no_pop and pop. The no_pop command specifies that the top insert buffer packet has not been inserted. The pop
15 command specifies that the insert buffer top packet has been processed and is popped.

The packet field command specifies which outgoing packet command fields are modified. The packet command fields which can be modified are shown
20 in the following table. All fields, except pkt.Freeze are within pkt.command[31:0].

	<u>Packet Cmd Field</u>	<u>Command Description</u>
25	Freeze	Pass, clear or set
	Invalid	Pass, clear or set
	Pcopy	Pass, clear or set
	Timestamp	Pass or clear from rules set if RingMaster
30	InvalidatorBusy	Pass, clear or set
	ResponderBusy	Pass, clear or set
	InvalidateLookahead	Pass, clear or set
	rrc	Pass, clear or set
	RRCFault	Pass or set
35	SpaAddrError	Pass or set
	MarkedFault	Pass or set
	UnMarkedFault	Pass or set

A packet is extracted by an empty packet routing command and push extract buffer command. A packet is copied by a pass packet routing command and a push extract buffer command.

5

Extract Buffer Management

The extract buffer holds a combination of short (2 words) and long (18 words) packets which equals 512 words. This yields a maximum of 28 long, 10 or 256 short or a combination of long/short packets. An extract buffer management flow control algorithm is required since the peak subring packet rate exceeds the rate the IBU can empty the extract algorithm. The control algorithm is based on the 15 principle of maximizing forward progress in conjunction with reducing system loading. Prioritizing the allocation of system resources to completing requests in progress achieves both of these goals since completing a request simultaneously 20 reduces system loading and achieves most efficient forward progress. The control algorithm is based on categorizing packets and using the Ring flow control mechanisms.

Empty space in Ring Extract Buffer is 25 divided into four regions by two configurable pointers, ExtractAllThreshold and ExtractRspThreshold. The current region is selected by the number of empty words remaining in the Extract Buffer. The regions going from empty to full are: ExtBuf:All, 30 ExtBuf.NonOpt, ExtBuf:PendReq and ExtBuf:Full. Extract Buffer State is maintained on a per RRC bank basis. The boundaries of these regions are configurable by software. Ring packets are grouped

into three mutually exclusive categories for Extract Buffer management:

- 5
- 10
- 15
- 20
- 25
- 30
- 35
- 40
- 45
- 50
- 55
- 60
- 65
- 70
- 75
- 80
- 85
- 90
- 95
- 100
- 105
- 110
- 115
- 120
- 125
- 130
- 135
- 140
- 145
- 150
- 155
- 160
- 165
- 170
- 175
- 180
- 185
- 190
- 195
- 200
- 205
- 210
- 215
- 220
- 225
- 230
- 235
- 240
- 245
- 250
- 255
- 260
- 265
- 270
- 275
- 280
- 285
- 290
- 295
- 300
- 305
- 310
- 315
- 320
- 325
- 330
- 335
- 340
- 345
- 350
- 355
- 360
- 365
- 370
- 375
- 380
- 385
- 390
- 395
- 400
- 405
- 410
- 415
- 420
- 425
- 430
- 435
- 440
- 445
- 450
- 455
- 460
- 465
- 470
- 475
- 480
- 485
- 490
- 495
- 500
- 505
- 510
- 515
- 520
- 525
- 530
- 535
- 540
- 545
- 550
- 555
- 560
- 565
- 570
- 575
- 580
- 585
- 590
- 595
- 600
- 605
- 610
- 615
- 620
- 625
- 630
- 635
- 640
- 645
- 650
- 655
- 660
- 665
- 670
- 675
- 680
- 685
- 690
- 695
- 700
- 705
- 710
- 715
- 720
- 725
- 730
- 735
- 740
- 745
- 750
- 755
- 760
- 765
- 770
- 775
- 780
- 785
- 790
- 795
- 800
- 805
- 810
- 815
- 820
- 825
- 830
- 835
- 840
- 845
- 850
- 855
- 860
- 865
- 870
- 875
- 880
- 885
- 890
- 895
- 900
- 905
- 910
- 915
- 920
- 925
- 930
- 935
- 940
- 945
- 950
- 955
- 960
- 965
- 970
- 975
- 980
- 985
- 990
- 995
- 1000
- 1005
- 1010
- 1015
- 1020
- 1025
- 1030
- 1035
- 1040
- 1045
- 1050
- 1055
- 1060
- 1065
- 1070
- 1075
- 1080
- 1085
- 1090
- 1095
- 1100
- 1105
- 1110
- 1115
- 1120
- 1125
- 1130
- 1135
- 1140
- 1145
- 1150
- 1155
- 1160
- 1165
- 1170
- 1175
- 1180
- 1185
- 1190
- 1195
- 1200
- 1205
- 1210
- 1215
- 1220
- 1225
- 1230
- 1235
- 1240
- 1245
- 1250
- 1255
- 1260
- 1265
- 1270
- 1275
- 1280
- 1285
- 1290
- 1295
- 1300
- 1305
- 1310
- 1315
- 1320
- 1325
- 1330
- 1335
- 1340
- 1345
- 1350
- 1355
- 1360
- 1365
- 1370
- 1375
- 1380
- 1385
- 1390
- 1395
- 1400
- 1405
- 1410
- 1415
- 1420
- 1425
- 1430
- 1435
- 1440
- 1445
- 1450
- 1455
- 1460
- 1465
- 1470
- 1475
- 1480
- 1485
- 1490
- 1495
- 1500
- 1505
- 1510
- 1515
- 1520
- 1525
- 1530
- 1535
- 1540
- 1545
- 1550
- 1555
- 1560
- 1565
- 1570
- 1575
- 1580
- 1585
- 1590
- 1595
- 1600
- 1605
- 1610
- 1615
- 1620
- 1625
- 1630
- 1635
- 1640
- 1645
- 1650
- 1655
- 1660
- 1665
- 1670
- 1675
- 1680
- 1685
- 1690
- 1695
- 1700
- 1705
- 1710
- 1715
- 1720
- 1725
- 1730
- 1735
- 1740
- 1745
- 1750
- 1755
- 1760
- 1765
- 1770
- 1775
- 1780
- 1785
- 1790
- 1795
- 1800
- 1805
- 1810
- 1815
- 1820
- 1825
- 1830
- 1835
- 1840
- 1845
- 1850
- 1855
- 1860
- 1865
- 1870
- 1875
- 1880
- 1885
- 1890
- 1895
- 1900
- 1905
- 1910
- 1915
- 1920
- 1925
- 1930
- 1935
- 1940
- 1945
- 1950
- 1955
- 1960
- 1965
- 1970
- 1975
- 1980
- 1985
- 1990
- 1995
- 2000
- 2005
- 2010
- 2015
- 2020
- 2025
- 2030
- 2035
- 2040
- 2045
- 2050
- 2055
- 2060
- 2065
- 2070
- 2075
- 2080
- 2085
- 2090
- 2095
- 2100
- 2105
- 2110
- 2115
- 2120
- 2125
- 2130
- 2135
- 2140
- 2145
- 2150
- 2155
- 2160
- 2165
- 2170
- 2175
- 2180
- 2185
- 2190
- 2195
- 2200
- 2205
- 2210
- 2215
- 2220
- 2225
- 2230
- 2235
- 2240
- 2245
- 2250
- 2255
- 2260
- 2265
- 2270
- 2275
- 2280
- 2285
- 2290
- 2295
- 2300
- 2305
- 2310
- 2315
- 2320
- 2325
- 2330
- 2335
- 2340
- 2345
- 2350
- 2355
- 2360
- 2365
- 2370
- 2375
- 2380
- 2385
- 2390
- 2395
- 2400
- 2405
- 2410
- 2415
- 2420
- 2425
- 2430
- 2435
- 2440
- 2445
- 2450
- 2455
- 2460
- 2465
- 2470
- 2475
- 2480
- 2485
- 2490
- 2495
- 2500
- 2505
- 2510
- 2515
- 2520
- 2525
- 2530
- 2535
- 2540
- 2545
- 2550
- 2555
- 2560
- 2565
- 2570
- 2575
- 2580
- 2585
- 2590
- 2595
- 2600
- 2605
- 2610
- 2615
- 2620
- 2625
- 2630
- 2635
- 2640
- 2645
- 2650
- 2655
- 2660
- 2665
- 2670
- 2675
- 2680
- 2685
- 2690
- 2695
- 2700
- 2705
- 2710
- 2715
- 2720
- 2725
- 2730
- 2735
- 2740
- 2745
- 2750
- 2755
- 2760
- 2765
- 2770
- 2775
- 2780
- 2785
- 2790
- 2795
- 2800
- 2805
- 2810
- 2815
- 2820
- 2825
- 2830
- 2835
- 2840
- 2845
- 2850
- 2855
- 2860
- 2865
- 2870
- 2875
- 2880
- 2885
- 2890
- 2895
- 2900
- 2905
- 2910
- 2915
- 2920
- 2925
- 2930
- 2935
- 2940
- 2945
- 2950
- 2955
- 2960
- 2965
- 2970
- 2975
- 2980
- 2985
- 2990
- 2995
- 3000
- 3005
- 3010
- 3015
- 3020
- 3025
- 3030
- 3035
- 3040
- 3045
- 3050
- 3055
- 3060
- 3065
- 3070
- 3075
- 3080
- 3085
- 3090
- 3095
- 3100
- 3105
- 3110
- 3115
- 3120
- 3125
- 3130
- 3135
- 3140
- 3145
- 3150
- 3155
- 3160
- 3165
- 3170
- 3175
- 3180
- 3185
- 3190
- 3195
- 3200
- 3205
- 3210
- 3215
- 3220
- 3225
- 3230
- 3235
- 3240
- 3245
- 3250
- 3255
- 3260
- 3265
- 3270
- 3275
- 3280
- 3285
- 3290
- 3295
- 3300
- 3305
- 3310
- 3315
- 3320
- 3325
- 3330
- 3335
- 3340
- 3345
- 3350
- 3355
- 3360
- 3365
- 3370
- 3375
- 3380
- 3385
- 3390
- 3395
- 3400
- 3405
- 3410
- 3415
- 3420
- 3425
- 3430
- 3435
- 3440
- 3445
- 3450
- 3455
- 3460
- 3465
- 3470
- 3475
- 3480
- 3485
- 3490
- 3495
- 3500
- 3505
- 3510
- 3515
- 3520
- 3525
- 3530
- 3535
- 3540
- 3545
- 3550
- 3555
- 3560
- 3565
- 3570
- 3575
- 3580
- 3585
- 3590
- 3595
- 3600
- 3605
- 3610
- 3615
- 3620
- 3625
- 3630
- 3635
- 3640
- 3645
- 3650
- 3655
- 3660
- 3665
- 3670
- 3675
- 3680
- 3685
- 3690
- 3695
- 3700
- 3705
- 3710
- 3715
- 3720
- 3725
- 3730
- 3735
- 3740
- 3745
- 3750
- 3755
- 3760
- 3765
- 3770
- 3775
- 3780
- 3785
- 3790
- 3795
- 3800
- 3805
- 3810
- 3815
- 3820
- 3825
- 3830
- 3835
- 3840
- 3845
- 3850
- 3855
- 3860
- 3865
- 3870
- 3875
- 3880
- 3885
- 3890
- 3895
- 3900
- 3905
- 3910
- 3915
- 3920
- 3925
- 3930
- 3935
- 3940
- 3945
- 3950
- 3955
- 3960
- 3965
- 3970
- 3975
- 3980
- 3985
- 3990
- 3995
- 4000
- 4005
- 4010
- 4015
- 4020
- 4025
- 4030
- 4035
- 4040
- 4045
- 4050
- 4055
- 4060
- 4065
- 4070
- 4075
- 4080
- 4085
- 4090
- 4095
- 4100
- 4105
- 4110
- 4115
- 4120
- 4125
- 4130
- 4135
- 4140
- 4145
- 4150
- 4155
- 4160
- 4165
- 4170
- 4175
- 4180
- 4185
- 4190
- 4195
- 4200
- 4205
- 4210
- 4215
- 4220
- 4225
- 4230
- 4235
- 4240
- 4245
- 4250
- 4255
- 4260
- 4265
- 4270
- 4275
- 4280
- 4285
- 4290
- 4295
- 4300
- 4305
- 4310
- 4315
- 4320
- 4325
- 4330
- 4335
- 4340
- 4345
- 4350
- 4355
- 4360
- 4365
- 4370
- 4375
- 4380
- 4385
- 4390
- 4395
- 4400
- 4405
- 4410
- 4415
- 4420
- 4425
- 4430
- 4435
- 4440
- 4445
- 4450
- 4455
- 4460
- 4465
- 4470
- 4475
- 4480
- 4485
- 4490
- 4495
- 4500
- 4505
- 4510
- 4515
- 4520
- 4525
- 4530
- 4535
- 4540
- 4545
- 4550
- 4555
- 4560
- 4565
- 4570
- 4575
- 4580
- 4585
- 4590
- 4595
- 4600
- 4605
- 4610
- 4615
- 4620
- 4625
- 4630
- 4635
- 4640
- 4645
- 4650
- 4655
- 4660
- 4665
- 4670
- 4675
- 4680
- 4685
- 4690
- 4695
- 4700
- 4705
- 4710
- 4715
- 4720
- 4725
- 4730
- 4735
- 4740
- 4745
- 4750
- 4755
- 4760
- 4765
- 4770
- 4775
- 4780
- 4785
- 4790
- 4795
- 4800
- 4805
- 4810
- 4815
- 4820
- 4825
- 4830
- 4835
- 4840
- 4845
- 4850
- 4855
- 4860
- 4865
- 4870
- 4875
- 4880
- 4885
- 4890
- 4895
- 4900
- 4905
- 4910
- 4915
- 4920
- 4925
- 4930
- 4935
- 4940
- 4945
- 4950
- 4955
- 4960
- 4965
- 4970
- 4975
- 4980
- 4985
- 4990
- 4995
- 5000
- 5005
- 5010
- 5015
- 5020
- 5025
- 5030
- 5035
- 5040
- 5045
- 5050
- 5055
- 5060
- 5065
- 5070
- 5075
- 5080
- 5085
- 5090
- 5095
- 5100
- 5105
- 5110
- 5115
- 5120
- 5125
- 5130
- 5135
- 5140
- 5145
- 5150
- 5155
- 5160
- 5165
- 5170
- 5175
- 5180
- 5185
- 5190
- 5195
- 5200
- 5205
- 5210
- 5215
- 5220
- 5225
- 5230
- 5235
- 5240
- 5245
- 5250
- 5255
- 5260
- 5265
- 5270
- 5275
- 5280
- 5285
- 5290
- 5295
- 5300
- 5305
- 5310
- 5315
- 5320
- 5325
- 5330
- 5335
- 5340
- 5345
- 5350
- 5355
- 5360
- 5365
- 5370
- 5375
- 5380
- 5385
- 5390
- 5395
- 5400
- 5405
- 5410
- 5415
- 5420
- 5425
- 5430
- 5435
- 5440
- 5445
- 5450
- 5455
- 5460
- 5465
- 5470
- 5475
- 5480
- 5485
- 5490
- 5495
- 5500
- 5505
- 5510
- 5515
- 5520
- 5525
- 5530
- 5535
- 5540
- 5545
- 5550
- 5555
- 5560
- 5565
- 5570
- 5575
- 5580
- 5585
- 5590
- 5595
- 5600
- 5605
- 5610
- 5615
- 5620
- 5625
- 5630
- 5635
- 5640
- 5645
- 5650
- 5655
- 5660
- 5665
- 5670
- 5675
- 5680
- 5685
- 5690
- 5695
- 5700
- 5705
- 5710
- 5715
- 5720
- 5725
- 5730
- 5735
- 5740
- 5745
- 5750
- 5755
- 5760
- 5765
- 5770
- 5775
- 5780
- 5785
- 5790
- 5795
- 5800
- 5805
- 5810
- 5815
- 5820
- 5825
- 5830
- 5835
- 5840
- 5845
- 5850
- 5855
- 5860
- 5865
- 5870
- 5875
- 5880
- 5885
- 5890
- 5895
- 5900
- 5905
- 5910
- 5915
- 5920
- 5925
- 5930
- 5935
- 5940
- 5945
- 5950
- 5955
- 5960
- 5965
- 5970
- 5975
- 5980
- 5985
- 5990
- 5995
- 6000
- 6005
- 6010
- 6015
- 6020
- 6025
- 6030
- 6035
- 6040
- 6045
- 6050
- 6055
- 6060
- 6065
- 6070
- 6075
- 6080
- 6085
- 6090
- 6095
- 6100
- 6105
- 6110
- 6115
- 6120
- 6125
- 6130
- 6135
- 6140
- 6145
- 6150
- 6155
- 6160
- 6165
- 6170
- 6175
- 6180
- 6185
- 6190
- 6195
- 6200
- 6205
- 6210
- 6215
- 6220
- 6225
- 6230
- 6235
- 6240
- 6245
- 6250
- 6255
- 6260
- 6265
- 6270
- 6275
- 6280
- 6285
- 6290
- 6295
- 6300
- 6305
- 6310
- 6315
- 6320
- 6325
- 6330
- 6335
- 6340
- 6345
- 6350
- 6355
- 6360
- 6365
- 6370
- 6375
- 6380
- 6385
- 6390
- 6395
- 6400
- 6405
- 6410
- 6415
- 6420
- 6425
- 6430
- 6435
- 6440
- 6445
- 6450
- 6455
- 6460
- 6465
- 6470
- 6475
- 6480
- 6485
- 6490
- 6495
- 6500
- 6505
- 6510
- 6515
- 6520
- 6525
- 6530
- 6535
- 6540
- 6545
- 6550
- 6555
- 6560
- 6565
- 6570
- 6575
- 6580
- 6585
- 6590
- 6595
- 6600
- 6605
- 6610
- 6615
- 6620
- 6625
- 6630
- 6635
- 6640
- 6645
- 6650
- 6655
- 6660
- 6665
- 6670
- 6675
- 6680
- 6685
- 6690
- 6695
- 6700
- 6705
- 6710
- 6715
- 6720
- 6725
- 6730
- 6735
- 6740
- 6745
- 6750
- 6755
- 6760
- 6765
- 6770
- 6775
- 6780
- 6785
- 6790
- 6795
- 6800
- 6805
- 6810
- 6815
- 6820
- 6825
- 6830
- 6835
- 6840
- 6845
- 6850
- 6855
- 6860
- 6865
- 6870
- 6875
- 6880
- 6885
- 6890
- 6895
- 6900
- 6905
- 6910
- 6915
- 6920
- 6925
- 6930
- 6935
- 6940
- 6945

ExtractRspThreshold, all non-optional and response packets which would normally be copied into the extract buffer are copied. All optional packets are passed. The extract buffer and RRC sheds optional
5 load as it becomes more fully subscribed. Within the ExtBuf:PendReq region, between ExtractRspThreshold and full, only packets relating to RRC pending request are accepted. All optional packets and non-optional (includes new incoming requests) packets
10 are busied by assertion of responder busy (rsp_busy) or invalidator_busy (inv_busy). The RRC sheds both optional load and new requests to relieve the overload situation.

If the extract buffer is full, the RIU will
15 pass optional packets, assert responder busy or invalidator_busy to non-optional requests. If the packet is a response to a request originated by this RRC, then the packet is left on the ring and pkt.cmd.timestamp field is cleared. Since only
20 packets relating to RRC pending requests are accepted in the third region the packet will be copied on the subsequent traversals(s) around the Ring. The fact that the packet is on its second revolution is transparent to all cells within the Ring. If the
25 packet not the response to an request originated by this RRC, but a packet that will satisfy the pending request (snarf), the pending request state is transitioned to Pend4:Reissue state.

30 Routing Directory

Routing Directory Organization

The Routing Directory consists of an Routing Directory Unit (RDU) and the actual ram based

directory. The RDU controls the directory access sequence and contains the routing rules. The preferred implementation of the rules hardcodes SPA rules and allows SVA rules to be loaded during RRC
5 initialization.

Figure 8 depicts the organization of a Routing Directory constructed in accord with a preferred practice of the invention. The Routing Directory contains an entry for all allocated pages
10 within the Ring:0 which the RRC pair is connected. A single routing directory entry is required for SVA pages allocated in more than a single Local Cache. The Routing Directory is organized as a three level btree for each of 128 subdirectories. A subdirectory
15 contains an entry for all allocated pages within a single Local Cache set, as specified by SVA[20:14]. A subdirectory can describe a maximum of 2^{10} pages, which supports a maximum of 64 Ring:0 Cells. Each Cell's Local Cache contains 128 sets, each set
20 containing 16 pages.

Ring:0 subrings are normally configured to interleave subpage transfer based on cache set lsb, SVA[14]. Thus half of the Routing Directory can be associated with each Ring:0 subring (Ring:0_A,
25 Ring:0_B).

The L1Entry is indexed by SVA[39:25]. An L1Entry is statically assigned to each 2^4 pages within a set. Each valid L1Entry specifies a pointer to a dynamically allocated L2Entry. A maximum of 2^{10}
30 L2Entries are available. Unallocated L2Entries are organized into a singly linked free list. The Index; L2Free List Pointer points to the first L2 free list entry. Each L2Entry in free list contains a pointer to the next L2Free list entry. The last L2 free list

entry contains a null pointer. Each L2Entry contains 16 subentries. Each valid L2Subentry contains a pointer to a dynamically allocated L3Entry. A maximum of 2^{10} L3Entries are available. Each L3Entry
5 describes a single page. Unallocated L3Entries are organized into a singly link free list. The Index_i L3Free List Pointer points to the first L3 free list entry. Each L3Entry in free list contains a pointer to the next L3Free list entry. The last L3 free list
10 entry contains a null pointer. Each L3Entry consists of 16 L3Subentries, each of which describe subpage state for eight subpages.

Routing Directory Structure

15 Figure 9 depicts the structure of a Routing Directory constructed in accord with a preferred practice of the invention. The Routing Directory is physically organized into L1, L2 and L3 tables for each subring. The L1 Table is addressed by
20 concatenation of SVA[20:14] (set number) and SVA[39:25] (index1). SVA[39:25,20:14] is referred to as index1. Each L2Entry consists of a L2subentry and L2V. The L2subentry is addressed by a concatenation of SVA[20:14], L1Table[index1].l2ptr and SVA[24:21]
25 (index2). The L2V is addressed by concatenation of SVA[20:14] and L1Table[index1].l2ptr. The L2V field is interpreted as level2 subentry valid field if the L2Entry is valid and is interpreted as level2 NextFree list entry pointer if L2Entry is invalid.
30 Each L3Entry consists of a L3subentry and L3Valid. The L3subentry is addressed by a concatenation of SVA[20:14], L2Table.l2sub[index2].ptr and SVA[13:10] (index3). An individual subpage state is addressed by a concatenation of SVA[20:14],

L2Entry.l2sub[index2].ptr, SVA[13:10] and SVA[9:7] (indexsp). The L3V field is interpreted as level3 subentry valid field if the L3Entry is valid and is interpreted as levels NextFree list entry pointer if 5 L2Entry is invalid.

A c-like description of the table structure of one bank is shown below.

```

10      struct      L1Entry L1Table[212];
      struct      L1Entry {
                unsigned      *l2ptr      :10;
                unsigned      l1valid     :1;
15      struct      L2Entry      L2Table[216];
                unsigned      l1check     :5;
                };
      struct      L2Entry      L2Table[216];
20      struct      L2Entry {
                L2subentry      l2sub[16];
                L2V              l2v;
                };
      struct      L2subentry {
25      struct      L2subentry {
                unsigned      *l3ptr      :10;
                unsigned      check       :5;
                };
      struct      L2Valid {
30      struct      L2Valid {
                unsigned      l2valid     :16;
                unsigned      check       :5;
                };
      struct      L2Invalid {
35      struct      L2Invalid {
                unsigned      *NextFree   :11;
                unsigned      unused      :5;
                unsigned      check       :5;
                };
      union      {
40      union      {
                L2Valid      Valid;
                L2Invalid    Invalid;
                }L2V;
      struct      L3Entry      L3Table[216];
45      struct      L3Entry      L3Table[216];

```



```

5      struct    L3Entry {
          L3subentry    l3sub[16];
          L3V            l3v;
          };

10     struct    L3subentry {
          unsigned      SpState[8]    :8;
          unsigned      check         :8;
          };

15     struct    L3Valid {
          unsigned      l3valid       :16;
          unsigned      PagePend      :1;
          unsigned      created       :1;
          unsigned      OwnerLimit     :1;
          unsigned      DupLimit      :1;
          unsigned      check         :5;
          };

20     struct    L3Invalid {
          unsigned      *NextFree     :11;
          unsigned      unused        :7;
          unsigned      check         :5;
          };

25     union     {
          L3Valid      Valid;
          L3Invalid    Invalid;
          } L3V;

30

```

Figure 10 depicts formats of Routing Directory entries in a preferred ring routing cell constructed in accord with the invention.

35

Level1 Entry

L1Entry contains a L2Entry pointer (l2ptr), L2 pointer valid bit (l1 valid) and L1 check bits (l1check). The l2ptr is used to address L2Entry as described above. L1valid set indicates a valid l2ptr and clear indicates an invalid l2ptr. L1check consists of a 5 bit single bit correct and double bit detect ECC code.

Level2 Subentry

L2Subentry contains a L3Entry pointer (l3ptr) and L2 check bits (l2check). The l3ptr is used to address L3Entry as described above. L2check consists of a 5 bit single bit correct and double bit detect ECC code.

Level2 VEntry

The L2VEntry has two formats, one for use as
10 L2VEntry (L2V.Valid) and the second if part of L2
free list (L2V.Invalid). L2V.valid consists of an
array of 16 valid bits and 5 ecc check bits.
L2V.valid.l2valid[n] set indicates a valid
L2Subentry[n] and clear indicates an invalid
15 L2Subentry[n].

L2V.invalid consists of NextFree L2 entry,
unused field and ecc check field.
L2V.invalid.NextFree is used to address the next L2
free list entry. NextFree high order bit clear
20 indicates a null pointer.

Level3 Subentry

L3Subentry contains 8 subpage state fields
pointer (L3Subentry.SpState) and check bits
25 (L3Subentry.check). L3Subentry.check consists of a 8
bit single bit correct and double bit detect ECC code.

Level3 VEntry

The L3VEntry has two formats, one for use as
30 L3VEntry (L3V.Valid) and the second if part of L3
free list (L3V.Invalid). L3V.valid consists of an
array of 16 valid bits, created, OwnerLimit, PagePend
and ecc check field. L3V.valid.l3valid[n]

corresponds to L3Subentry[n] and is a function of L3V.created:

	<u>l3valid[n]</u>	<u>created</u>	<u>L3Subentry[n]</u>
5	0	0	all subpages invalid
	1	0	use actual subpage state
	0	1	all subpages exclusive
	1	1	use actual subpage state

10

L3V.valid.OwnerLimit is set or cleared by software to limit ownership of subpages within the page to the Local Ring:0. When clear, specifies that any Ring:0 can own the subpages. When set, specifies that only cells within the Local Ring:0 can own subpages in nonexclusive, exclusive, atomic or transcient atomic states. Remote Ring:0s can receive read only copies. A page fault is signalled by the RRC to ownership requests when OwnerLimit is set. Recombine requests will not propagate beyond Local Ring:0 when OwnerLimit is set. Cleared during RRC descriptor allocation.

L3V.valid.PagePend indicates that an RRC page deallocate sequence is in progress when set. L3V.invalid consists of NextFree L3 entry, unused field and ecc check field. L2Vinvalid.NextFree is used to address the next L3 free list entry. NextFree high order bit clear indicates a null pointer.

30

Routing Directory Operations

There are eight possible RRC routing directory operations. These operations occur within a single packet time. All packet types cause a single operation to execute.

- Level1 Miss lookup operation
- Level2 Miss lookup operation
- Normal lookup operation
- Deallocate Level3 Subentry
- 5 - Deallocate Level2 Subentry
- SPA Access
- Refresh

SVA Look-up Operations

10 All Lookup Operations proceed through five
steps, subdirectory selection, Level1 table lookup,
Level2 table lookup, Level3 table lookup and Level3
subpage state modification. The level2 and level3
entries might be required to be allocated in order to
15 complete lookup operation.

Normal Lookup Operation

 A normal lookup operation occurs when the
Level 1 table indicates that the level2 entry is
20 valid and the Level2 table indicates that the level3
entry is valid.

 The L1EntryAddr, L2EntryAddr and L3EntryAddr
are calculated as described under Routing Directory
Structure. L1EntryAddr-l1valid being set indicates
25 the level2 pointer is valid. The concatenation of
SVA[20:14] and L1Entry[index1].ptr specifies the
L2Entry, since the subdirectory is implied by the 10
bit pointer. The level2 subentry
(L2EntryAddr-l2sub[index2]) is indexed by
30 SVA[24:21]. L2Entry.l2v[index2] being set indicates
the level3 pointer (L2EntryAddr-l2sub[index2].ptr) is
valid. The concatenation of SVA[20:14] and
L2Entry[index2].ptr specifies the L3Entry, since the
subdirectory is implied by the 10 bit pointer. The

level3 subentry (L3EntryAddr→l3sub[index3]) is indexed SVA[13:10]. Each level3 subentry contains subpage state for eight subpages, indexed by SVA[9:7]. Index3 refers to SVA[13:10] and Indexs refers to SVA[9:7]. L3Entry.valid[index3] and L3Entry.created specifies whether the corresponding subentry is valid and default subpage state if invalid.

The subpage state is then presented to
10 Routing Directory Rules which generate the new subpage state to be stored in L3EntryAddr→l3sub[index3].SpState[indexs].

Level1 Miss Operation

15 A Level1 Miss lookup operation occurs when the Level 1 table lookup indicates that the level2 entry pointer is invalid. The level1 miss lookup operation allocates a level2 entry and level3 entry to allow the Routing Directory to describe the page
20 and subpage state of the request SVA. A level 1 miss is caused by the first reference within a Ring:0 to a page within an address range denoted by SVA[39:25]. If the packet command is ciu_ex_summary, the request is part of a page create operation.

25 The Level2 Free list head is pointed to by the L2 Free List Pointer. The free list head L2Entry.NextFree specifies the pointer to the next free Level2 Entry. The level2 free list head is allocated as a Level2 Entry by writing
30 LevelEntry.pointer with level2 free list pointer and setting L1Entry.valid. Level2 free list pointer is updated by writing Level2 free list pointer with level2 free list head, L2Entry.NextFree.

L2Entry.l2v.l2valid[SVA[24:21]] is set and the remaining valid bits are cleared.

<u>Step</u>	<u>Operation</u>
5	
1	L1EntryAddr-l2ptr = L2FreeListPointer
2	L1EntryAddr-valid = set
3	temp = L2FreeListPointer-l2v.Invalid.NextFree
4	L2FreeListPointer-l2v.Valid.l2valid[15:0] = clear
10 5	L2FreeListPointer-l2v.Valid.l2valid[(SVA[24:21])] = set
6	L2FreeListPointer = temp

If the L2FreeListPointer is null, a Routing
15 directory fault is signalled. This error can only occur if the corresponding Ring:0 has allocated greater than 2^{10} unique pages within a single subdirectory or the RRC has failed to allocate or deallocate level2 entries correctly. 2^{10} unique pages
20 corresponds to 64 processor cells (2^{10} pages/ 2^4 pages per cell).

Level2 Miss Operation

A Level2 Miss lookup operation occurs when
25 the Level2 table lookup indicates that the level2 entry pointer is invalid. The level2 miss lookup operation allocates a level3 entry to allow the Routing Directory to describe the page and subpage state of the request SVA. A level2 miss is caused by
30 the first reference within a Ring:0 to a page within an address range denoted by SVA[39:21]. If the packet command is ciu_ex_summary, the request is part of a page create operation.

The Level3 Free list head is pointed to by the L3FreeListPointer. The free list head, L3Entry.NextFree specifies the pointer to the next free L3 Entry. The level3 free list head is allocated as a L3Entry by writing L2Entry.l2sub[index2].l3ptr with level3 free list pointer. L3 free list pointer is updated by writing L3 free list pointer with levels free list head L3Entry.NextFree. L3Entry.l3v.l3valid[SVA[13:10]] is set and the remaining valid bits are cleared. L3Entry.l3sub(SVA[13:10]).SpState[SVA[9:7]] is set to proper subpage_state and the remaining seven subpage_states are set to invalid. If the packet command is ciu_ex_summary, L3Entry.l3v.Valid.created is set.

<u>Step</u>	<u>Operation</u>
7	L2EntryAddr-l2sub.l3ptr = L3FreeListPointer
20 8	temp = L3FreeListPointer-l3v.Invalid.NextFree
9	L3FreeListPointer-l3v.Valid.l3valid[15:0] = clear
10	L3FreeListPointer-l3v.Valid.l3valid[(SVA[13:10])] = set
25 11	L3FreeListPointer-l3sub[(SVA[13:10])].SpState[7:0] = invalid
12	L3FreeListPointer-l3sub[(SVA[13:10])].SpState[(SVA[9:7])] = new-sp-state
13	L3FreeListPointer-l3v.Valid.Created = clear
30 14	L3FreeListPointer = temp

SVA Deallocate Operations

SVA Deallocate Operations are initiated by the packet commands `ciu_inv_sum_alloc`, `ciu_pure_inv_sum_alloc` and `RrcDeallocDesc`. All deallocate operations
5 proceed through five steps, subdirectory selection, Level1 table lookup, Level2 table lookup, update Level2 entry to deallocate Level3 entry. If all level2 subentries are invalid, then the Level2Entry is also deallocated.

10

Deallocate Level3 Entry Operation

A deallocate Level3 Entry operation is executed if all level2 subentries are not invalid after the level3 entry is deallocated. The following
15 operation is executed if the packet command is `ciu_inv_summary`:

<u>Step</u>	<u>Operation</u>
20 1	<code>L2EntryAddr-12v.Valid.12valid[(SVA[24:21])] = clear</code>
2	<code>L3EntryAddr-13v.Invalid.NextFree = L3FreeListpointer</code>
3	<code>L3FreeListPointer =</code>
25	<code>L2EntryAddr-12sub[(SVA[24:21])].13ptr</code>

The following operation is executed if the packet command is `ciu_wr_descriptro.data` or `ciu_inv_sum_allocate.data`:

30

<u>Step</u>	<u>Operation</u>
1	L2EntryAddr→l2v.Valid.l2valid[(SVA[24:21]] = clear
5 2	Packet.DataWord2[9:0] = L2EntryAddr→l2sub[(SVA[24:21])].l3ptr

Deallocate Level2 Entry Operation

A deallocate Level2 Entry operation is
10 executed if all level2 subentries are invalid after the level3 entry is deallocated.

<u>Step</u>	<u>Operation</u>
15 1	Execute deallocate Level3 Entry Operation.
2	L1EntryAddr→l1valid = clear
3	L2EntryAddr→l2v.Invalid.NextFree = L2FreeListPointer
4	L2FreeListPointer = L1EntryAddr→l2ptr
20	

SPA Operation

SPA operation is initiated by two packet commands, spa_read_request and spa_write_request. There are two types of SPA operations, SpaDirect and
25 SpaLookup.

The SpaDirect operation preforms a spa read or write based of the directly on the spa location specified in the packet address. Spa_write_request data is supplied in Pkt.DataWord0. Spa_read_request
30 response data is supplied in Pkt.DataWord0.

The SpaLookup operation is used to emulate the RDU lookup and deallocate operations under

controlled conditions from spa address space. The
Pkt.Address location offset and
ctl\$RiuLookupCommand.SubAddr forms the emulated SVA
and the remainder of ctl\$RiuLookupCommand forms the
5 emulated packet command. The RDU uses the emulated
packet command and emulated to perform a normal
lookup or deallocate operation. The results of the
SpaLookup operation can be examined by access to
L2FreeListPointer, L3FreeListPointer and SpaDirect
10 locations.

Refresh Operation

Refresh is accomplished in a distributed
manner by each RRC independently controlling it
15 refresh. A single refresh requires one packet
processing time. Thus refresh can only be performed
during empty, or frozen packets, since these packet
require no processing. The RRC attempts perform
refresh in the background by using all empty, frozen
20 and invisible packet times for refresh.

The three dram based data structures,
L1Table, L2PTable and L3SubpageState require
refresh. During a refresh operation one of 512 rows
is refreshed in each of the data structures.

25

ECC Correct Operation

When the RRC detects a correctable ECC
error, additional packet times are required to
correct the data structure. Ring freezing is used to
30 allow multiple packet times for the correction and
re-lookup.

When a correctable ECC error is detected,
the RRC initiates Full-1 Ring Freeze method, since
the ECC error is not know early in packet

processing. The first packet is handled as described in Extract Buffer Full section and the remaining packets are marked frozen for one revolution. The corrected data structure re-written, saved in RDU and
5 then used to reprocess the packet which cause the ecc error after one Ring revolution. Only RRC originated packets will be processed, non-originated packets are extracted by another cell within the Ring. Saving the corrected packet allows RRC originated packets to
10 be correctly processed for single bit (correctable) hard errors.

Routing Directory Initialization

RRC hardware does not automatically
15 initialize the Routing Directory. All Level1Entries, Level2Subentries and Level3Subentries must be initialized to invalid. Level2 and Level3 FreeListPointers and free lists must be initialized to indicate all entries and subentries are free.

20

InterRing Buffer

The InterRingBuffer buffers incoming packets from the Ring Interconnect extract buffer and interfaces to the InterRingLink. A preferred
25 structure of the InterRingBuffer is shown in Figure 11A. The InterRingBuffer also controls the RRC Bank buffer management. The InterRingBuffer consists of the InterRing BufferUnit (IBU), InterRing Dram Buffer (IDB) and InterRingLink Interface.

30

InterRing Buffer Structure

IBU

The IBU buffers and formats outgoing
5 InterRingLink packets, provides error checking and
buffering for InterRingLink incoming packets and
controls the InterRing Dram Buffer. The Outgoing
Buffer contains 72 words, which holds a maximum of 4
long packets and the Incoming Buffer holds 72, words
10 for maximum of 4 long packets. The sizing of the
Incoming Buffer is dependent on the InterRingLink
rate and absolute delay (length and switching
delay). The 72 word Incoming Buffer is sized for a
10km maximum length, directly connected 1 gbit/sec
15 link. Short packets consume 2 words and long packets
consume 18 words of Outgoing or Incoming Buffers.
The Outgoing and Incoming buffers are managed with
fifo discipline.

20 InterRingDramBuffer

The InterRing Dram Buffer (IDB) contains the
Ring Packet Queue (RPQ) and RRC ErrorQueue. A
configuration for the InterRingDram queues in a
preferred ring routing cell constructed in accord the
25 invention is shown in Figure 11B. The RPQ buffers
all packets from Routing Interconnect unit until the
packet has been transmitted over InterRingLink to the
downlink RRC and an acknowledge has been returned.
When the Extract Buffer and the Outgoing Buffer are
30 full, the Extract Buffer-packets are emptied into the
IDB until the Outgoing Buffer can accept them. The
RPQ is organized as a circular queue. RPQPushPointer
points to the first free word for queueing the next
packet from the Extract Buffer. The RPQPopPointer

points to the first word of the next packet to be transmitted over the InterRingLink. The RPQCommitPointer points to the first word of oldest packet which has not be acknowledged by downlink
5 RRC. Packets between RPQPushPointer and RPQPopPointer are queued to be transmitted over the InterRingLink. Packets between RPQPopPointer and RPQCommitPointer are waiting for acknowledgement from the downlink RRC. If RPQPushPointer equals
10 RPQPopPointer, the there are no queued packets in the RPQ for transmission on InterRingLink. If RPQPopPointer equals RPQCommitPointer, all transmitted packets have been acknowledge. If Ring Packet Queue overflows, a fatal error is indicated by
15 IbuMasterConfig.RPQFull.

The IDB also contains the RRC ErrorQueue. The RRC places an entry in the ErrorQueue for each RRC detected error. Each ErrorQueue entry consists of the packet corresponding to the error (if any) and
20 one word of error information. EQPushPointer points to the first free words for queueing the next error entry. The EQPopPointer points to the first word of the least recent queued error. If EQPopPointer equals EQPushPointer, the RRC ErrorQueue is empty.
25 The EQPopPointer is not modified by the RRC, it is only modified by SPA access. IbuMasterConfig.EQFull is set when the RRC ErrorQueue overflows. The RRC continues normal operations but new errors are not logged.

30

InterRingLink Interface

The InterRingLink interface converts the 32 or 40 bit wide IBU outgoing packet and incoming packet interfaces to or from the physical link format.

35

RRC Buffer Management

Each RRC pair logically manages a pair of buffers with fifo discipline, one in each direction between the Ring:0 and Ring:1 (see Figure 11B). A
5 RRC bank manages 5 buffers, 3 buffers between the Ring and the outgoing InterRingLink and 2 buffers between the incoming InterRingLink and the Ring. A packet copied from the Ring of uplink RRC bank, passes through the Extract Buffer, optionally
10 throughout Ring Packet Buffer, through InterRingLink Outgoing Buffer, and is then transmitted over the InterRingLink. The packet then passes through the InterRingLink Incoming Buffer and Insert Buffer of downlink RRC bank. Packets traveling in the opposite
15 direction are copied from the Ring of downlink RRC bank, pass through the Extract Buffer, Ring Packet Buffer and InterRingLink Outgoing buffer. The packet then passes through the InterRingLink Incoming Buffer and Insert Buffer of uplink RRC bank.

20 The Extract and Insert buffer provide rate adjustment between the Ring and RRC Routing Bus. The Incoming and Outgoing buffers provide rate adjustment between the InterRingLink and the RRC Routing Bus. The Ring Packet Queue handles possible overflow of
25 these buffers by being able to store the maximum number of pending packets and prevents deadlock between the RRC pair. The buffers are sized such that the Ring Packet Queue is used infrequently for buffering between Extract and InterRingLink outgoing
30 buffers. The Ring Packet Queue is also used to log all packets transmitted over InterRingLink, in the event that re-transmission is required due to link errors.

The preferred RRC bank implementation provides an independent path between the Extract buffer and InterRingLink Outgoing Buffer and between the InterRingLink Incoming Buffer and the Insert
5 Buffer. When the InterRingLink Incoming buffer is not empty and the Insert Buffer is not full, the IBU transfers a packet from the Incoming to Insert buffer.

When the Extract buffer is not empty, InterRingDram Queue is empty and the InterRingLink
10 Outgoing buffer is not full, the IBU transfers a packet from the Extract to the Outgoing buffer. If the Outgoing buffer is full, the packet is transferred from the Extract buffer to the InterRingDram Queue. Packets contained in
15 InterRingDram Queue are always transferred to Outgoing buffer before any packets from Extract buffer.

InterRingLink Packet Protocol and Format

20 The InterRingLink packet protocol and format is required to be applicable across a range of link rates for multiple generations of UCS computers. The protocol and format should span interconnection distances from 25 meters within a single room to
25 10,000 km over tariffed telecommunication networks. The format should be compatible with emerging telecommunications standards, such as SONET. It is heavily based on simplified concepts borrowed from SDLC protocol; see Communications Architecture for
30 Distributed Systems, R.J. Cypser, Addison-Wesley Publishing Company, 1978 or Telecommunications Networks: Protocols, Modeling and Analysis, Mischa Schwartz, Addison Wesley, 1987. Figure 12 depicts a

preferred InterRingLink frame format used in a ring routing cell constructed in accord with the invention.

Acknowledgements and Retransmissions

- 5 The connection between an RRC pair is full duplex, consisting of two unidirectional InterRingLinks. Because of the uncertain reliability of the InterRingLinks, a very positive and definite acknowledgement is call for at the link level. The
- 10 transmitter RRC receives an positive acknowledgement over the incoming InterRingLink for each packet transmitted over outgoing InterRingLink. Waiting for an acknowledgement prior to transmitting the next packet leads to very low efficiency. The IBU
- 15 attaches a sequence number to each transmitted packet so that up to 2^{16} packets can be outstanding prior to the first acknowledge. Because of the duplex capability, two independent sets of sequence numbers are used, one for each direction of packet flow.
- 20 This permits the recovery of one InterRingLink to proceed independent of the sequence numbers of the other InterRingLink between the RRC pair and avoids possible timing ambiguities. Only data packets are sequenced checked.
- 25 Two sequence numbers are maintained by each IBU. The IbuXmitSeqNum, operated modulo 2^{16} , provides a count for each data transfer packet that is transmitted from the IBU. The IbuRcvSeqNum, also modulo 2^{16} , is incremented once for each valid
- 30 in-sequence, error-free data packet that is received by that IBU. When one IBU sends its IbuRcvSeqNum to the downlink IBU, it will served to indicate the next data packet that is expected and to acknowledge all data packets received up to, but not including the

value indicated by the packet receive sequence number. Each subsequent IbuRcvSeqNum transmitted reconfirms that all preceding messages have been accepted. Multiple data packets can be acknowledged 5 in a single packet.

Packet Format

InterRingLink packet format consists of five fields, Frame Start, Preamble, Ring Packet, CRC and 10 Frame End. The Frame Start and Frame End fields are link specific framing fields and are added and stripped off InterRingPacket by InterRingLink Interface.

The Frame Start and Frame End fields are 15 dependent on the type of physical link.

The PreAamble which specifies an link control command, a sequence acknowledge field, a Transmit frame sequence number (TFSN) and the received frame sequence number (RFSN). Control commands define the 20 number of words and content of Ring Packet subfield.

	<u>Control Cmd</u>	<u>Ring Pkt Size</u>	<u>Sequence</u>
25	idle	0 words	nonsequenced
	initialize request	0 words	nonsequenced
	initialize sequence	0 words	nonsequenced
	initialize acknowledge	0 words	nonsequenced
	short packet	2 words	sequenced
30	long packet	18 words	sequenced

Commands are divided into sequenced and nonsequenced commands. The sequence numbers for the 35 link to be incremented by the RRC pair for sequenced frames. Sequence numbers for the link are not incremented for nonsequenced frames.

Transmit Frame Sequence Number (TFSN) represents the sequence number of the frame. Each successive sequenced frame has its sequence number incremented by one. The InterRingLink used mod 2^{16} sequence numbering. When the transmitter reaches its maximum sequence number, it is forced to stop transmitting until a frame in the reverse direction is received, acknowledging an outstanding packet. The receive frame sequence number (RFSN) acknowledges the receipt of RFSN - 1 and any frames preceding that number not already acknowledged. RFSN indicates that the receiver is expecting sequenced frame number RFSN. The transmitting IBU buffers all frames not yet acked positively. Once acked positively the frame can be purged and its sequence number reused.

The InterRingLink, like HDLC, provides both a "piggy-back" feature, with the ack function embedded in a frame transmitted in the reverse direction. Multiple frames can be acknowledged in a single frame. The three possible acknowledges are tabulated as follows:

	<u>Acknowledge</u>	<u>Function</u>
25	Ready to receive (RR)	Acknowledges all frames received up to and including Received Frame Sequence Number (RFSN) minus one.
30	Not ready to receive (RNR)	Provides flow control for a temporary busy condition. Acknowledges all frames up to and including Received Frame Sequence Number (RFSN) minus one.
35	Reject (REJ)	Rejects all frames from Received Packet Sequence Number (RPSN) on. Positively

acknowledges all frames up to and including Received Frame Sequence Number (RFSN) minus one.

5

Idle command that specifies that the Ring packet field is null, but the remainder of preamble and checksum field are valid. The initialize request, initialize sequence and initialize acknowledge commands are used to initialize and synchronize sequence number protocol between IBUs at each end of InterRingLink. Initialize sequence indicates the transmit IBU has reset its sequence numbers. The receiving RRC should reset sequence number, Ring Packet Queue, clear IbuMasterConfig.RPQFull and transmit a reset acknowledge control command packet. Initialize request causes the downlink IBU to reset sequence number, Ring Packet Queue, clear IbuMasterConfig.RPQFull and transmit a initialize sequence packet. The short packet and long packet commands indicate that the Ring Packet field contains a short or long packet respectively.

25 The InterRingLink uses two mechanisms for error recovery.

1 - Reject recovery. This is assume to be always used where possible to speed up InterRingLink recovery. However, it can be used only once for a given frame. It cannot be invoked on repeats of that frame.

2 - Timeout recovery. This must always be used in addition to reject recovery. Without timeout recovery, an isolated sequenced frame or the last in a

35

sequence of frames could not be recovered if garbled. In addition, since a reject recovery may be used only once per sequence frame, multiple losses of a given frame must be handled through a timeout.

Link packet errors include:

- Incorrect checksum
- Out of sequence transmit packet number
- Clock recovery error
- Link framing error

The RNR acknowledge provides flow control when the incoming buffer is full. For optimal use of InterRingLink bandwidth, the minimum incoming buffer size is the number of frames received during a round trip InterRingLink delay.

The Cyclic Redundancy Field (CRC) is calculated from Preamble and RingPacket fields.

20 IB Operations

The IBU buffer management control executes the following IB operations to transfer packets between the 5 RRC buffers.

25 Ring Interconnect Packet

A packet from the Extract Buffer is transferred over the Routing Bus to both the Ring Packet Queue and stored starting at the location specified by RPQPushPointer. The pointer is incremented by 18 words for both short and long packets. If the Ring Packet is empty (RPQPushPointer equals RPQPopPointer) and Outgoing buffer is not full, the packet is also pushed into the Outgoing buffer and RPQPopPointer is incremented.

IbuXmitSeqNum is incremented by one for each data packet transmitted on InterRingLink.

Ring Interconnect Packet from Dram Buffer

- 5 A packet is transferred from the Ring Packet Queue to the Outgoing buffer when the Ring Packet Queue is not empty and the Outgoing Buffer is not full. The RPQPopPointer is incremented by 32 words.

10 InterRingLink Packet

A packet is transferred from the Incoming Buffer to the Insert Buffer. The count IbuRcvSeqNum is incremented by one.

15 InterRing Dram Buffer Refresh

A refresh sequence is performed on InterRing Dram when the IRBRefreshCount overflows. The refresh interval is set by IRBRefreshInterval. Refresh takes top priority over other Routing Bus operations.

20

InterRingLink Retry

- Each packet received over the InterRingLink specifies the last error free packet received and whether an error was detected. If the acknowledge
25 indicates no errors, the RPQCommitPointer and CommitPktNumber locations are updated in the following manner:

- 30 1 - $RPQCommitPointer = RPQCommitPointer + (RcvPktSeqNum - CommitPktNum) * 32$
 2 - $CommitPktNum = RcbPktSeqNum$

When a reject packet is received , the IBU must resend all packets with sequence numbers after

RPQCommitPointer through RPQPushPointer. Ibis is accomplished by adjusting the RPQPopPointer.

- 1 - RPQCommitPointer = RPQCommitPointer +
5 (RcvPktSeqNum - CommitPktNum) * 32
- 2 - CommitPktNum = PktRcvSeqNum
- 3 - RPQPopPointer = PktRcvSeqNum
- 4 - Store received error packet in ErrorQueue

10 Error Log

For any error detected by the RI, RDU and IB within RRC, the corresponding packet and 2 words of error status are recorded in the RRC ErrorQueue. The packet should also be queue in Ring Packet Queue if
15 it is being transmitted to downlink RRC.

SPA Read/Write InterRing Dram

Read or written through InterRingDram locations. Note that if Outgoing Buffer becomes full
20 or RRC detects errors, the InterRingDram is modified.

Ring:1 Operations

This section which follows describes preferred RRC operations in context of the
25 multi-level memory system hierarchy described above. The RRC can be viewed as a packet switch, which routes packets based on rules which are a function of packet type and RRC directory based subpage state. First the subpage, page and RRC based states and
30 pending states are described. Next, the an description of key rules concepts and sequences are given. Next complete subpage rules are provided. Page and special operations are then described.

Complete rules encompassing page operations, special operations and subpage operations

Originating Conventions

5 Figure 13 illustrates a typical request sequence and defines the location terminology used in this section. The Originating Ring:0 is the Ring:0 contains the requestor cell, which issued the original request packet (step 1). If the requested
10 subpage is not valid within the originating (or local) Ring:0, RRC:0 extracts the packet and forwards it to the local RRC:1. The locally originated request is then inserted by the local RRC:1 onto Ring:1 (step 3). An RRC:1, connected to a remote
15 Ring:0, which holds a valid copy of the requested subpage then extracts the remotely originated request (step 4). The remotely originated request is then inserted onto the remote Ring:0 by RRC:0 (step 5). The responding cell extracts the request and inserts
20 a response (step 6). The remotely originated request is extracted by the RRC:1 (step 7) and forwarded to the partner RRC:1. This RRC:1, called the responding RRC:1, inserts the response packet onto the Ring:1 (step 8). The packet travel around the Ring:1 to the
25 originating RRC:1 and it is extracted (step 9). The packet is forwarded to the originating Ring:0s RRC:0 and inserted onto Ring:0 (step 10). The request sequence is then completed by the originating (requestor) cell extracting the packet (step 11).

30 The packet is processed by RRCs eight times during this typical sequence. The two RRC:0s and two RRC:1s interpret the packet in eight unique ways based on packet source and packet origin as shown in the table below. The RRC operations and rules are

structured based on these eight combinations of RRC, packet source and packet origin.

For most part RRC:0 locally originated packet rules are analogous to RRC:1 remotely
5 originated packet rules. From the point of view of a local Ring:0, the local RRC:0 makes the remainder of the system logically appear as a single cell to locally originated requests. From the point of view
10 of the Ring:1 the local RRC:1 makes the local Ring:0 logically appear as a single cell to remotely originated requests.

The pkt.RequestorId and RiuCellAddress are compared to determine if the request is locally originated.

15

RRC Subpage State Description - Preferred Embodiment

Subpage and pending request state are stored in the RRC pair for all subpages which are present in any cache within the Ring:0 or for which a request is
20 pending to or from a nonlocal Ring:0. RRC:0 stores pending states, for requests from a remote Ring:0s. The RRC:1 stores pending state for local requests which are pending to a remote Ring:0. Both RRC:0 and RRC:1 store subpage state. RRC:0 and RRC:1 subpage
25 state is consistent, except during intervals when the subpage state is being updated. The rules and pending states properly handle all window cases. Subpage States for a preferred practice of the invention are set forth immediately below.

30

RDU:0 (RduState0) Enumeration

RDU:0 state consists of two independent fields and a reserved field. Subpage state field describes the collective state of all subpages

present within the local Ring:0. Pend field describes whether any remotely originated requests to this Ring:0 are pending. The individual states are described in the following sub-sections.

5

RDU:0 State

	<u>Field Name</u>	<u>States</u>	<u>Number bits</u>	<u>Description</u>
10	SubpageState	Invalid Ro Nex Ex	2	Ring:0 subpage state
15	Pend	None Ro Owner	2	Remote pending

SubpageState:Inv

20 Indicates there is no valid copy of subpage within the Ring:0.

SpState:Ro

25 SpState:Ro indicates that one or more caches within the Ring:0 might contain a read only copy of the subpage, but none contains owner state. Under certain circumstances, the RRC pair has recorded read only state (SpState:Ro) for a subpage, but the subpage is invalid (SpState:Inv) within the Ring:0,
30 since no read only copies exist in any of the cells.

SpState:Nex

35 SpState:Nex indicates the subpage owner, and zero or more read only copies of subpage exist within the local Ring:0. Zero or more read only copies might exist in remote Ring:0s. If no read copies exist outside the local Ring:0, the first invalidate, or recombine which indicates that no remote read only

copies exist causes a subpage state transition to exclusive. Remote RRCs did not signal when they deallocate descriptors.

5 SpState:Ex

SubpageState:Ex indicates that there are no valid copies of subpage outside of the local Ring:0. The local Ring:0 contains an exclusive owner and no read only copies or a nonexclusive owner with zero or
10 more read only copies. The RRC:0 records all local cell exclusive ownership states, exclusive, atomic, and transient-atomic as exclusive.

Pending:None

15 PendingNone request state indicates that there is no remotely originated pending activity to the local Ring:0 relating to the corresponding subpage.

20 Pendl:Ro

Pendl:Ro indicates that one remotely originated sva_read_ro request to the local Ring:0 is pending. Additional remotely originated sva_read_ro requests are returned to the partner RRC:1 with
25 RspBsy set, until the request is closed out.

Pendl:Owner

Pendl:Owner indicates that one or more remotely originated ownership or invalidation
30 requests to the local Ring:0 is pending. In addition one or more remotely originated sva_read_ro requests might also be pending. All remotely originated ownership, read only and invalidation requests are forwarded to the Ring:0, when in this state.

Ownership requests are sva_read_ex, sva_get and sva_getw. Invalidations requests are write_inv and lookahead invalidates.

5 RDU:1 (RduStatel) Enumeration

RDU:1 state consists of three independent fields and a reserved field. Subpage state field describes the collective state of all subpages present within the local Ring:0. Pend field describes whether any locally originated requests to Ring:1 are pending. Pend1 field describes whether the pending request might need to be reissued due to the RRC:1 missing a remotely originated ownership packet that wasn't extracted to satisfy a pending request and a missing page (snarf). Individual states are described in the following sub-sections.

RDU:1 State

20	<u>Field Name</u>	<u>States</u>	<u>Number bits</u>	<u>Description</u>
	SubpageState	Invalid Ro Owner	2	Ring:0 subpage state
25	Pend	None Ro Owner	2	Ring:0 pending
	Pend1	None ReIssue	1	Ring:0 pending
30				

SubpageState:Inv

Indicates there is no valid copy of subpage within the Ring:0.

35

SpState:Ro

SpState:Ro indicates that one or more caches within the Ring:0 might contain a read only copy of

the subpage, but none contains owner state. Under certain circumstances, the RRC pair has recorded read only state (SpState:Ro) for a subpage, but the subpage is invalid (SpState:Inv) within the Ring:0, 5 since no read only copies exist in any of the cells.

SubpageState:Owner

SubpageState:Owner indicates that the subpage is owned exclusively (Ex) or nonexclusively 10 (Nex) within the local Ring:0. The local Ring:0 contains an exclusive owner and no read only copies or a non-exclusive owner with zero or more read only copies. The RRC:1 records all local cell ownership states, non-exclusive, exclusive, atomic, and 15 transient-atomic as owner.

Pending:None

PendingNone request state indicates that there is no locally originated pending activity to 20 the Ring:1 relating to the corresponding subpage and pending field.

Pend:Ro

Pend:Ro indicates that one locally 25 originated sva_read_ro request to Ring:1 is pending. Additional locally originated sva_read_ro requests are returned to the partner RRC:0 with RspBsy set, until the request is closed out.

30 Pend:Owner

Pend:Owner indicates that one or more locally originated ownership or invalidation requests to Ring:1 are pending. In addition one or more locally originated sva_read_ro requests might also be

pending. All locally originated ownership, read only and invalidation requests are forwarded to the Ring:0, when in this state. Ownership requests are sva_read_ex, sva_get and sva_getw. Invalidations requests are write_inv and lookahead invalidates.

Pendl:ReIssue

Pendl:ReIssue indicates that the RRC:1 might have missed a potential non-RRC:1 originated response packet to an RRC:1 originated request. Pendl:Reissue state is used by RRC:1 to prevent spurious missing_page error (page fault) by differentiating between the case of a remotely originated ownership packet that wasn't extracted to satisfy a pending request and a missing page (snarf). If the RRC:1 originated request returns without a response and Pendl:ReIssue then the returning request is busied by the RRC causing the requesting processor cell or RRC to reissue the request, thereby preventing a spurious page fault. Pendl:Reissue is cleared when no locally originated requests are outstanding or during the first ownership (with data) response.

RRC Subpage State Description 25 - Alternative Embodiment I

Described below are subpage states used in a first alternate embodiment of the invention.

30 RDU:0 (RduState0) Enumeration

RDU:0 state consists of four independent fields. Subpage state field describes the collective state of all subpages present within the local Ring:0. Pending1 field describes whether any remotely originated sva_read_ro, sva_read_ex, or

write_invalidate requests to this Ring:0 are pending. Pending2 field describes whether a remotely originated sva_get request is pending. Pending3 field describes whether a remotely originated
5 sva_getw (get and wait) request is pending and whether additional getw requests were issued while this request was pending. Multiple sva_getw requests are recorded to properly signal transient atomic state to requesting cache.

10 The three pending fields independently recorded a maximum of one remotely originated pending request of each type per subpage. Multiple requests of the same type to the same subpage are returned with busy status to the partner RRC:1 until the
15 pending request has been closed out. The individual states are described in the following sub-sections.

	<u>Field Name</u>	<u>States</u>	<u>Number bits</u>	<u>Description</u>
20	SubpageState	Invalid Ro Nex Ex	2	Ring:0 subpage state
25	Pend1	None Ro Ex OwnerInv	2	Remote pending
	Pend2	None Get	1	Remote pending
30	Pend3	None Getw Atomic	2	Remote pending

35 SubpageState:Inv

Indicates there is no valid copy of subpage within the Ring:0.

SpState:Ro

SpState:Ro indicates that one or more caches within the Ring:0 might contain a read only copy of the subpage, but none contains owner state. Under 5 certain circumstances, the RRC pair has recorded read only state (SpState:Ro) for a subpage, but the subpage is invalid (SpState:Inv) within the Ring:0, since no read only copies exist in any of the cells. When no read only copies exist, read only subpage 10 state is adjusted to invalid subpage state, by the first locally originated request, remotely originated invalidate or remotely originated read only request.

SpState:Nex

15 SpState:Nex indicates the subpage owner, and zero or more read only copies of subpage exist within the local Ring:0. Zero or more read only copies might exist in remote Ring:0s. If no read copies exist outside the local Ring:0, the first invalidate, 20 or recombine which indicates that no remote read only copies exist causes a subpage state transition to exclusive. Remote RRCs did not signal when they deallocate descriptors.

SpState:Ex

25 SubpageState:Ex indicates that there are no valid copies of subpage outside of the local Ring:0. the Ring:0 contains an exclusive owner and no read only copies or a non-exclusive owner with zero or 30 more read only copies. The RRC:0 records all local cell exclusive ownership states, exclusive, atomic, and transient-atomic as exclusive.

Pending[1,2,3]:None

Pending None request state indicates that there is no remotely originated pending activity to the local Ring:0 relating to the corresponding 5 subpage and pending field.

Pendl:Ro

Pendl:Ro indicates that a remotely originated sva_read_ro request to the local Ring:0 is 10 pending. Additional remotely originated sva_read_ro request are returned to the partner RRC:1 with RspBsy set, until the request is closed out.

Pendl:Ex

15 Pendl:Ex indicates that a remotely originated sva_read_ex request to the local Ring:0 is pending. A preceding remotely originated sva_read_ro request might be pending. Additional remotely originated sva_read_ex requests are returned to the 20 partner RRC:1 with RspBsy set, until the request is closed out.

Pendl:OwnerInv

Pendl:OwnerInv indicates that a remotely 25 originated write_invalidate or lookahead invalidate (write_invalidate.lookahead, sva_read_ex.lookahead, sva_get.lookahead or sva_getw.lookahead) request to the local Ring:0 is pending. A preceding remotely originated sva_read_ro request might be pending. 30 Subsequent lookahead invalidates from Ring:1 are emptied. Subsequent remotely originated sva_read_ex or write_inv requests are busied. Since each subpage has a single owner and only the owner can issue a write_invalidate request, additional write

invalidates in this state indicate an error condition. Subsequent sva_read_ro requests are marked RspBsy and duplicate.data packets are emptied.

5 Pend2:Get

Pend2:Get indicates that a remotely originated sva_get request to the local Ring:0 is pending. Subsequent remotely originated sva_get requests are returned to the originating cell with
10 pkt.data set, which indicates subpage is already atomic, until the request is closed out. The subpage might not yet be atomic, but will be before the subsequent request is completed.

15 Pend3:Getw

Pend3:Getw indicates that one or more remotely originated sva_getw requests to the local Ring:0 are pending. Unlike all other read requests, all getw requests are forwarded when in this state,
20 rather than ensuing request being busied. When a getw response from Ring:0 indicates the page is already atomic (getw.data.RspBsy), Pend3:Atomic state is recorded.

25 Pend3:Atomic

Indicates that ownership packet is required to satisfy a pending getw request for atomic ownership. The owning cache is guaranteed to be in transient atomic state (tat). Subsequent remotely
30 originated getw or get requests can be returned indicating the subpage is already atomic. Subsequent Ring:0 ownership packets (release.data or sva_read ex.data) are extracted, forwarded to Ring:1, for

forwarding to other Ring:0s which are pending atomic (RRC:1 Pend3:Atomic).

Pend3:Atomic is set under two conditions. First, when a getw response indicates the page is already atomic. Second when a remotely initiated release.data enters a Ring:0, Pend3:Atomic is set, so that the RRC:0 can extract a subsequent release.data for forwarding to other Ring:0s which are pending atomic (RRC:1 Pend3:Atomic). The release.data could be locally originated if a cell within the Ring:0 acquired atomic ownership from the remotely originated release.

RDU:1 (RduStatel) Enumeration

15 RDU:1 state consists of six independent fields. Subpage state field describes the collective state of all subpages present within the local Ring:0 and whether an owner invalidation is in progress. Pend1a field describes whether a locally originated sva_read_ro request is pending on Ring:1. Pend1b field describes whether a locally originated sva_read ex request is pending on Ring:1. Pend2 field describes whether a locally originated sva_get request is pending. Pend3 field describes whether one locally originated sva_getw request is pending and whether additional getw requests were issued while this request was pending. Multiple sva_getw requests are recorded to properly signal transient atomic state to requesting cache. Pend5 field describes whether the pending request might need to be reissued due to the RRC:1 missing a remotely originated ownership packet that wasn't extracted to satisfy a pending request and a missing page (snarf).

The five pending fields independently recorded a maximum of one locally originated pending request of each type per subpage. Multiple requests of the same type are returned with busy status to the partner RRC until the pending request has been closed out. The individual states are described in the following sub-sections.

10	<u>Field Name</u>	<u>States</u>	<u>Number Bits</u>	<u>Description</u>
	SubpageState	Invalid Ro Owner OwnerInvalidtg	2	Ring:0 subpage state
15	Pend1a	None Ro	1	Ring:0 pending
	Pend1b	None Ex	1	Ring:0 pending
20	Pend2	None Get	1	Ring:0 pending
	Pend3	None Getw Atomic	2	Ring:0 pending
25	Pend4	None Reissue	1	Ring:0 pending

SubpageState:Inv

30 Indicates there is no valid copy of subpage within the Ring:0.

SpState:Ro

SpState:Ro indicates that one or more caches within the Ring:0 might contain a read only copy of the subpage, but none contains owner state. Under certain circumstances, the RRC pair has recorded read only state (SpState:Ro) for a subpage, but the subpage is invalid (SpState:Inv) within the Ring:0, since no read only copies exist in any of the cells.

When no read only copies exist, read only subpage state is adjusted to invalid subpage state, by the first locally originated request, remotely originated invalidate or remotely originated read only request.

5

SubpageState:Owner

SubpageState:Owner indicates that the subpage is owned exclusively (Ex) or non-exclusively (Nex) within the local Ring:0. The local Ring:0
10 contains an exclusive owner and no read only copies or a non-exclusive owner with zero or more read only copies. The RRC:1 records all local cell ownership states, non-exclusive, exclusive, atomic, and transient-atomic as owner.

15

SubpageState:OwnerInvalidating

Pend1:OwnerInv indicates that a locally originated write_invalidate request to the Ring:1 is pending. Since each subpage has a single owner and
20 only the owner can issue a write_invalidate request, additional write_invalidates in this state indicates an error condition. Invalidates implied by sva_read ex, get or getw are recorded as Pend1b:Ex, Pend2:Get, or Pend3:Getw states respectively. Subsequently
25 locally originated sva_read_ro requests are marked RspBsy and duplicate.data packets are emptied.

Pending[1a,1b,2,3]:None

Pending:None request state indicates that
30 there is no locally originated pending activity to the Ring:1 relating to the corresponding subpage and pending field.

Pendl:Ro

Pendl:Ro indicates that a locally originated sva_read_ro request to the Ring:1 is pending. Additional locally originated sva_read_ro requests 5 are returned to the partner RRC:0 with RspBsy set, until the request is closed out.

Pendlb:Ex

Pendl:Ex indicates that a locally originated 10 sva_read_ex request to the Ring:1 is pending. Additional remotely originated sva_read_ex requests are returned to the Partner RRC:0 with RspBsy set, until the request is closed out.

15 Pend2:Get

Pend2:Get indicates that a locally originated sva_get request to the Ring:1 is pending. Subsequent remotely originated sva_get requests are returned to the local originating cell with pkt.data 20 set, which indicates subpage is already atomic, until the request is closed out. The subpage might not yet be atomic, but will be before the subsequent request is completed.

25 Pend3:Getw

Pend3:Getw indicates that one or more locally originated sva_getw requests to the local Ring:0 are pending. Unlike all other read requests, all getw requests are forwarded when in this state, 30 rather than ensuing request being busied. When a getw response from Ring:1 indicates the page is already atomic (getw.data.RspBsy), Pend3:Atomic state is recorded.

Pend3:Atomic

Indicates that ownership packet is required to satisfy a pending getw request for atomic ownership. The owning cache is guaranteed to be in
5 transient atomic state (tat). Subsequent remotely originated getw or get requests can be returned indicating the subpage is already atomic. Subsequent Ring:0 ownership packets (release.data or sva_read ex.data) are extracted, forwarded to Ring:1, for
10 forwarding to other Ring:0s which are pending atomic (RRC:1 Pend3:Atomic). Pend3:Atomic does not guarantee that one or more getw requests are still pending.

Pend3:Atomic is set under two conditions.
15 First, when a getw response indicates the page is already atomic. Second when a remotely initiated release.data enters a Ring:0, Pend3:Atomic is set, so that the RRC:0 can extract a subsequent release.data for forwarding to other Ring:0s which are pending
20 atomic (RRC:1 Pend3:Atomic). The release.data could be locally originated if a cell within the Ring:0 acquired atomic ownership from the remotely originated release.

25 Pend4:ReIssue

Pend4:ReIssue indicates that the RRC:1 might have missed a potential non-RRC:1 originated response packet to an RRC:1 originated request. Pend4:ReIssue state is used by RRC:1 to prevent spurious missing
30 page_error (page fault) by differentiating between the case of a remotely originated ownership packet that wasn't extracted to satisfy a pending request and a missing page (snarf). If the RRC:1 originated request returns without a response and Pend4:ReIssue

then the returning request is busied by the RRC causing the requesting processor cell or RRC to reissue the request, thereby preventing a spurious page fault. Pend4:ReIssue is cleared when no locally
5 originated requests are outstanding or during the first ownership (with data) response.

RRC Subpage State Description
- Alternative Embodiment II

10

Described below are subpage states used in a second alternative embodiment of the invention.

Subpage State

15

RRC related information of subpage states is described here.

Invalid (Inv)

Invalid subpage state indicates that there
20 is no valid copy of subpage within the Ring:0.

Read Only (Ro)

Read only state indicates that one or more caches within the Ring:0 might contain a read only
25 copy of the subpage. It is possible that no local cache contains a read only copy of the subpage when the RRC has recorded read only state. If the last cache which held a particular subpage deallocated the descriptor, but another cache within the Ring:0 has a
30 descriptor for same SVA Page allocated, the RRC will not be immediately notified and will temporarily maintain read only state. Another case is when a recombine.nex.data leaves a Ring:0, the RRC under rare circumstances (pending read only request, PlRo
35 for instance), has to assume that a read only copy

might exist within the Ring:0. Normally the assertion of recombine.nex.inv_busy indicates a read only copy of the subpage exists. When no read only copies exist, read only subpage state is adjusted to 5 invalid subpage state, by the first invalidate or remote read only request.

NonExclusive (Nex)

Non-exclusive state indicates the subpage 10 owner, zero or more read only copies of subpage exist within the local Ring:0. Zero or more read only copies might exist in remote Ring:0s. If no read copies exist outside the local Ring:0, the first invalidate, or recombine which indicates that not 15 remote read only copies exist causes a subpage state transition to exclusive. Remote RRCs did not signal when they deallocate descriptors.

Exclusive (Ex)

20 Exclusive state indicates that there are no valid copies of subpage outside of the local Ring:0. The local Ring:0 contains an exclusive or non-exclusive owner and zero or more read only copies. The RRC records all local cell exclusive 25 ownership states, exclusive, atomic, and transient-atomic as exclusive.

Pending Request Subpage State

Pending request state indicates the type of 30 pending request, whether additional requests have been merged, whether the request needs to be reissued or whether the request has been completed and the RRC is wait for the original request to be returned. The

value of 0 or 1 for # indicates pending Ring:0 or Ring:1 request state respectively.

PendNone (P#None)

- 5 PendingNone request state indicates that there is no pending activity relating to the corresponding subpage.

PendWait (P#Wait)

- 10 PendingWait indicates that a read request (sva_read_{ro,nex,ex},get,getw) is still pending, but a response or merged response has already been issued. When the returning request packet is received, the packet is emptied and the pending
15 request is closed out. If the returning request packet contains data and ownership, an error should be signalled.

PendWaitMergeRo (P#WaitMrgRo)

- 20 PendingWaitMergeRo indicates that a read request is still pending and one or more sva_read ro.nodata packet(s) have been merged while waiting for the returning request. A second request cannot be issued until the first request is closed out.
25 When the returning request packet is received, the packet is emptied, the pending request is closed out and a sva_read_ro request is usually issued. If the returning request packet contains data and ownership, an error should be signalled.

30

PendWaitMergeEx (P#WaitMrgEx)

PendWaitMergeEx indicates that a read request is still pending and one or more sva_read ex.nodata packets(s) have been merged while waiting

for the returning request. Since the pending state always recognizes the strongest request merged, this state also includes zero or more merged sva_read_ro packets. A second request cannot be issued until the
5 first request is closed out. When the returning request packet is received, the packet is emptied, the pending request is closed out and a sva_read_ex request is usually issued. If the returning request packet contains data and ownership, an error should
10 be signalled.

PendWaitMergeAt (P#WaitMrqAt)

PendWaitMergeAt indicates that a read request is still pending and one or more
15 {get,getw}.nodata packet(s) have been merged while waiting for the returning request. Since the pending state always recognizes the strongest request merged, this state also includes zero or more merged sva_read {ro,ex} packets. A second request cannot be issued
20 until the first request is closed out. When the returning request packet is received, the packet is emptied, the pending request is closed out and a getw request is usually issued. If the returning request packet contains data and ownership, an error should
25 be signalled.

PendReissue (P#Rei)

PendReissue indicates that a read request is pending and the request should be reissued when it
30 returns. This state is usually entered when an event detect while the request is pending indicates that the request will be returned without a response or the response will return with stale data. For example, if a sva_read_ro is pending and an

invalidate is detected, the returned data might be stale. Sometime pendreissue state is entered to allow a local Ring:0 transaction to complete prior to moving ownership to a remote Ring:0.

5

PendReissueMergeRo (P#ReiMrgRo)

PendReissueMergeRo indicates that a read request is pending and the request should be reissued when it returns and one or more sva_read_ro.nodata
10 packet(s) have been merged while waiting for the returning request. Since the pending state always recognizes the strongest request merged, this state also includes zero or more merged sva_read_ro
15 packets. A second request cannot be issued until the first request is closed out. When the returning request packet is received, the packet is emptied, the stronger of the merged sva_read_ro or the pending request is reissued.

20

PendReissueMergeEx (P#ReiMrgEx)

PendReissueMergeEx indicates that a read request is pending and the request should be reissued when it returns and one or more sva_read_ex.nodata
25 packet(s) have been merged while waiting for the returning request. Since the pending state always recognizes the strongest request merged, this state also includes zero or more merged sva_read_ro
30 packets. A second request cannot be issued until the first request is closed out. When the returning request packet is received, the packet is emptied, the stronger of the merged sva_read_ex or the pending request is reissued.

PendReissueMergeEx (P#ReiMrgAt)

PendReissueMergeAt indicates that a read request is pending and the request should be reissued when it return and one or more {get,getw}.nodata
5 packet(s) have been merged while waiting for the returning request. Since the pending state always recognizes the strongest request merged, this state also includes zero or more merged sva_read_{ro,ex} packets. A second request cannot be issued until the
10 first request is closed out. When the returning request packet is received, the packet is emptied, the getw request is reissued.

PendRo (P#Ro)

15 PendRo indicates that a sva_read_ro request is pending. A sva_read_ex request received while in any pending ro state, causes a transition to ReissueMergeEx pending state.

20 PendRoMergeRo (P#RoMrgRo)

PendRoMergeRo indicates that a sva_read_ro request is pending and one or more sva_read_ro.nodata packet(s) have been merged while waiting for the returning request.

25 When the returning request is received, the packet is marked as merged. A sva_read_ex request received while in any pending ro state, causes a transition to ReissueMergeEx pending state.

30 PendEx (P#Ex)

PendEx indicates that a sva_read_ex request is pending.

PendExMerge (P#ExMrq)

PendExMerge indicates that a sva_read_ex request is pending and an sva_read_{ro,ex} request has been merged. When the returning request is 5 received, the packet is marked as merged.

PendExMergeAt (P#ExMrqAt)

PendExMerge indicates that a sva_read_ex request is pending and an get or getw request has 10 been merged. When the returning request is received, the packet is marked as merged.

PendAtomic (P#At)

PendAtomic indicates that a get or getw 15 request is pending or that a getw.nodata response has been returned.

PendAtomic (P#AtMrq)

PendAtomic indicates that a get or getw 20 request is pending or that a getw.nodata response has been returned and a get, getw or sva_read_{ro,ex} request has been merged. When the returning request is received, the packet is marked as merged.

25 PendRecombine (P#Rec)

PendRecombine indicates a recombine.data.{tat,at,ex,nex} is pending from the local Ring:0. PendRecombine is only used in RDU:1.

30 PendMrsp (PlMrsp)

PendMrsp indicates that a response which includes merged request(s) from this Ring:0 is currently on the Ring:0. Merged requests from this Ring:0 would be indicated by Rdu:1 pending states

PORoMrg*, POExMrg* or POAtMrg. While in PendMrsp additional sva_read requests are emptied to satisfy Cell transaction completion criteria.

5 PendMrspMrgRo (P#MrspMrgRo)

PendMrspMrgRo indicates that a response which includes merged request(s) from this Ring:0 is currently on the Ring:0 and a remote sva_read_ro request has been merged during this interval.

10

PendMrspMrgEx (P#MrspMrgEx)

PendMrspMrgEx indicates that a response which includes merged request(s) from this Ring:0 is currently on the Ring:0 and a remote sva_read_ex and
15 sva_read_ro requests has been merged during this interval.

PendMrspMrgAt (P#MrspMrgAt)

PendMrspMrgEx indicates that a response
20 which includes merged request(s) from this Ring:0 is currently on the Ring:0 and a remote get or getw and sva_read_ex or sva_read_ro requests has been merged during this interval.

25 PendInvalidate (pInv)

PendInvalidate indicates that a pending write_invalidate, write_invalidate.lookahead is pending. Invalidates implied by sva_read_{ro,ex}, get or getw are recorded as PendRo, PendEx, or PendAt
30 states.

RDU:0 State (RduState0) Enumeration

RduState0 stores Ring:0 subpage state and Ring:1 pending request (from remote Ring:0s) state.

Illegal combinations of Ring:0 subpage state and Ring:1 pending State are shown as dashes. Each RduState₁ is named by postfixing the Ring:1 pending state name to Ring:0 subpage state.

- 5 A table showing RDU:0 pending states used in this alternate embodiment is shown in Figure 17.

RDU:1 State (RduState₁) Enumeration

- RduState₁ stores Ring:1 subpage state and
10 Ring:0 pending request (from remote Ring:0s) state. Illegal combinations of Ring:0 subpage state and Ring:0 pending State are shown as dashes. Each RduState₁ is named by postfixing the Ring:1 pending state name to Ring:0 subpage state.

- 15 A table showing RDU:1 states used in this alternate embodiment is shown in Figure 18.

Page State

- The description in the sections which follow
20 assumes use of the preferred (as opposed to the alternate) subpage states.

OwnerLimit

- OwnerLimit is set or cleared by software to
25 limit ownership of subpages within the page to the Local Ring:0. It is stored in both RRC:0 and RRC:1. When clear, OwnerLimit specifies that any Ring:0 can own the subpages. When set, specifies that only cells within the Local Ring:0 can own subpages in
30 non-exclusive, exclusive, atomic or transcient atomic states. Remote Ring:0s can receive read only copies. When OwnerLimit is set, a page fault is signalled to ownership requests by the RRC returning the request with InvBsy and RspBsy clear. Recombine

requests will not propagate beyond Local Ring:0 when OwnerLimit is set. Default state for RRC descriptor allocation is clear.

5	<u>OwnerLimit</u>	<u>Mnemonic</u>	<u>Description</u>
	0	OL:All	Ownership not limited to local Ring:0
10	1	OL:Local	Ownership limited to local Ring:0

DupLimit

DupLimit specifies the conditions which causes a duplicate.data to be forwarded from the local Ring:0 to Ring:1. DupLimit is stored by RRC:0.

	<u>DupLimit</u>	<u>Mnemonic</u>	<u>Description</u>
20	0	Dup:CondForward	Conditional forwarding
	1	Dup:AllForward	Always forward

Dup:CondForward

Duplicate.data packets within the page are conditionally forwarded by RRC:0 to Ring:1, if the page is not exclusively owned by the local Ring:0. Dup:CondForward is default state for RRC descriptor allocation.

Dup:AllForward

Duplicate.data is always forwarded by RRC:0 to Ring:1.

35 Pending Request Page State

Some pending RRC operations effect how the RRC handles SVA requests to any subpage within a page. The Pending request page state field in

corresponding Level3VEntry specifies this pending information.

Page:None

- 5 No pending page level RRC operation is in progress for the corresponding descriptor. Page:None is default state for RRC descriptor allocation.

Page:PendDeallocate

- 10 Within the RRC:0 indicates that the RRC pair is in the process of determining whether the descriptor should be deallocated or that RRC:1 is being signalled to deallocate the descriptor. Within the RRC:1 indicates the RRC pair is in the process of
15 deallocating the descriptor.

 The RRC descriptor transitions into Page:PendDeallocate in response to a single cell indicating that it is deallocating this descriptor. If the descriptor is not present in any cell within
20 the Ring:0, RRC:0 and RRC:1 deallocate the descriptor. Other page level packets, such as ciu inv_sum.nodata, ciu_inv_sum_allocate.no_data and ciu ex_sum.nodata are busied while in this state. Remotely originated subpage requests are busied by
25 RRC:0 and RRC:1. Locally originated subpage requests are unaffected in RRC:0, since the requesting local cache has an allocated descriptor will prevent the RRC pair from deallocating its descriptor. Locally originated subpage requests are busied by RRC:1,
30 since the RRC pair has committed to deallocate the descriptor. The descriptor will be subsequently re-allocated on the next request after the page has been deallocated.

Other RRC State

The following state fields are stored on a per RRC or per RRC bank basis. These state fields globally apply to all subpages and pages, unless
5 otherwise stated.

Extract Buffer State

Empty space in Ring Extract Buffer is divided into four regions by two configurable
10 pointers, ExtractAllThreshold and ExtractRspThreshold. The current region is selected by the number of empty words remaining in the Extract Buffer. The regions going from empty to full are: ExtBuf.All,
ExtBuf:NonOpt, ExtBuf:PendReq and ExtBuf:Full.
15 Extract Buffer State is maintained on a per RRC bank basis. The boundaries of these regions are configurable by software. Ring packets are grouped into three mutually exclusive categories for Extract Buffer management:

20

1) Optional Packet The following remotely originated packets are optional from a coherency point of view:
25 - Recombine with SpState:(Inv|Ro)
 - Duplicate.data SpState:Inv
 - sva_read_ro.data with SpState:Inv

30

2) NonOptional Packet Must be extracted based on SpState to complete a remotely originated request. Does not include optional packets.

35

3) Response packet A locally originated packet or a packet which satisfies a RRC pending request (RRC:0 Pend(1,2,3), RRC:1 Pend(1,2,3,4,5)).
40

These packet categories are handled in the four extract buffer regions as follows:

5	<u>Region</u>	<u>Optional Packet</u>	<u>NonOptional Packet</u>	<u>Response Packet</u>
	ExtBuf:All	accepted	accepted	accepted
	ExtBuf:NonOpt	pass	accepted	accepted
	ExtBuf:PendReq	pass	pass and assert	accepted
10	ExtBuf:Full	pass	RspBsy,InvBsy pass and assert RspBsy,InvBsy	pass, see text

15 ExtBuf:All

ExtBuf:All region is between empty and ExtractAllThreshold. All packets are accepted in the ExBuf.All region.

20 ExtBuf:NonOpt

ExtBuf:NonOpt region is between ExtractAllThreshold and ExtractResponseThreshold. Nonoptional and Response packets are accepted in this region. No signalling within the packet is required
25 when optional packets are not accepted.

ExtBuf:PendReq

ExtBuf:PendReq region is between ExtractResponseThreshold and extract buffer full.
30 Only response packets are accepted in this region. Nonoptional packets are rejected by the assertion of RspBsy and InvBsy. No signalling within the packet is required when optional packets are not accepted.

35 ExtBuf:Full

ExtBuf:Full region is when the extract buffer is full. No packets can be accepted in this

region. Nonoptional packets are rejected by the
assertion of RspBsy and InvBsy. No signalling within
the packet is required when optional packets are not
accepted. A locally originated packet, is passed
5 with pkt.TimeStamp clear. The packet will be
extracted by this RRC on a subsequent Ring
revolution. For remotely originated response
packets, not carrying data, set both RspBsy and
InvBsy. Remotely originated response packets with
10 read only data are passed. Remotely originated
ownership response packets are passed and RRC:1
transitions to Pend1:ReIssue.

Pend1:Reissue state is used by RRC:1 to
prevent spurious missing_page_error (page fault) by
15 differentiating between the case of a remotely
originated ownership packet that wasn't extracted to
satisfy a pending request and a missing page. When
set Pend1:Reissue indicates that the RRC:1 might have
missed a potential non-RRC originated response packet
20 to an RRC originated request. If the RRC:1
originated request returns without a response and
Pend1:ReIssue then the returning request is busied by
the RRC causing the requesting processor cell or RRC
to reissue the request, thereby preventing a spurious
25 page fault. RRC:0 uses the SpState field to prevent
missing_page_error. The RRC:0 SpState field records
whether a read only copy or ownership is within the
Ring:0 to satisfy the request. If the RRC:0
originated request returns without a response and a
30 sufficient copy or ownership exists within the
Ring:0, the request is reissued.

Ring Freeze State

In addition to the following states, the RRC must keep track of the first and last packet marked frozen.

5

Fz:None

Fz:None indicates that the RRC is not in Ring Freeze sequence.

10 Fz:SingleEccErr

Indicates a correctable Ecc error has been detected during the lookup sequence corresponding to this packet time. Start of Ring:l Freeze sequence. Always transitions to Fz:Freeze during the next
15 packet time. The corresponding packet is handled as described under Extract Buffer full section.

Fz:Freeze

Indicates that the RRC is in a Ring Freeze
20 sequence. All packets not previously marked frozen, are marked frozen.

Fz:UnFreeze

Indicates that the RRC has ended a Ring
25 Freeze sequence. All packets previously marked frozen, are marked unfrozen and processed normally by the RRC. Any condition which initiates a Ring Freeze sequence, would transitions to Fz:SingleEccErr or Fz:Freeze states.

30

Rules Tables Basics

Ownership

The foundation of strong ordering is based
5 on the rules ownership concept. Each subpage in
memory system has a single owner. For the value of a
subpage to be modified, the owner must be exclusive
meaning all copies of the subpage has been
invalidated. When read only copies of a subpage
10 exist, the owner is non-exclusive and cannot modify
the subpage. The subpage ownership can be
transferred dynamically between caches, with the
owning cache being the dynamic home for the subpage.

The memory system and RRC specifically is
15 based on the property of inclusion. The RRC
directory keeps track of the state of all subpages
within the Ring:0 and the state of all pending
requests entering or leaving a Ring:0.

20 Other Concepts

Owner and Invalidation Requests

All owner requests and write invalidate
requests go all the way to the owning cell regardless
25 of what other requests are outstanding. All owner
requests and invalidates will cause an RRC to go to
Pend:Owner when the packet is inserted onto the
Ring. That is, the RRC:1 goes to Pend:Owner when the
packet is inserted onto Ring:1 and the RRC:0 goes to
30 Pend:Owner when it inserts the packet onto Ring:0.

An RRC will transition from Pend:Owner to
Pend:None when:

- 1 - An owner carrying packet is extracted
by the RRC.

- 2- An RRC:0 with SpState: (Inv|Ro)
extracts a remote owner requests or
invalidate with pkt.!InvBsy.
- 5 3- An RRC with SpState:Owner extracts a
locally originated invalidate with
pkt.!InvBsy.
- 10 4- Page destruction always sets
SpState:Inv and Pend:None for all
subpages within the page in the local
RRC:0 and RRC:1 and all remote RRC:1s
with the page allocated. All pending
state is cleared, so that page faults
can be signaled to other pending
requests (the requests would hang busy
15 or timeout otherwise).

Read Only Requests

Locally originated read only requests are
only inserted by the local RRC:1 if the RC:1 is not
20 Pend:Ro. This is an optimization to avoid flooding
too many identical requests onto Ring:1 and
eventually the responder. Instead of inserting a
such a request, the RRC:1 will extract it and set
pkt.RspBsy. The RRC:0 at the responder, however,
25 does not bounce back read only requests in this way.
Instead the read only request are inserted on the
Ring:0.

The RRC transitions from Pend:Ro to
Pend:None when an owner carrying packet or read only
30 data carrying packet is extracted by the RRC.

A packet carrying read only data is "killed"
when it passes an RRC in Pend:Owner state. An
sva_read_ro packet is "killed" by setting pkt.RspBsy
and a duplicate is emptied.

Clearing and Interpretation of Packet Busy Bits

RspBsy can always be cleared when responding with data to a request packet.

InvBsy is always clear when leaving the
5 originating Ring:0, since the RRC:0 will not extract
the packet if pkt.InvBsy. The RRC will pass the
packet and set pkt.InvBsy, RspBsy to cause the request
to be reissued. InvBsy is always clear, when a
remote ownership request or invalidate is enters a
10 remote Ring:0, since an RRC:1 will not extract a
request if pkt.InvBsy. The RRC:1 will pass the
packet and set pkt.InvBsy, RspBsy to cause the request
to be reissued. The lookahead will actually continue
and initiate invalidates on the remain Ring:0s with
15 SpState:Ro.

Use of Packet RRC Bit

Used by sva_read_ro, recombines and release
to indicate that the packet should be reissued by the
20 originating cell or RRC if the request is not
completed. Pkt.rrc is used in two manners.

Pkt.rrc is used to allow a
sva_read_ro.no_data or recombine.data request to
traverse the entire Ring:0, to allow the request to
25 completed locally, rather than the RRC:0 prematurely
extracting the packet. Pkt.rrc is set when a locally
originated sva_read_ro.no_data or recombine.data
passes the RRC:0. If the request has not been
completed (sva_read_ro.data or recombine.no_data)
30 prior to returning to the originating cell, the
originating cell passes the packet. The RRC:0
extracts the packet with rrc set and processes the
request normally.

When an RRC:0 or RRC:1 can't extract a non-originated recombine.data or release.data to satisfy a pending request (called snarf) it sets rrc bit so that the packet can be re-circulated. If the
5 release.data returns to the originating cell, it must be re-issued. If the release.data is extracted by another cell, no action is required, since another release will be issued since transient atomic subpage state is preserved (tat). The recombine.data should
10 not be extracted by any cell and is passed by the originating CIU, effecting a reissue.

Pkt.rrc is always initialized to clear and passed by CCU and CIU. It is cleared by RRC:0, when locally originated packet is extracted.

15

Invalidation Method

Several types of packets cause invalidation of read only copies. Write_invalidate causes an explicit invalidation and sva_read_ex, get, and getw
20 packets cause implicit invalidation in addition to a read request. All invalidates are guaranteed to remove all read only copies based on two methods. First is that the invalidate packet causes all read only copies it passes to be invalidated. Secondly,
25 when a sva_read_ro passes a request with a pending invalidate, the requestor asserts rsp_busy to indicate the data in packet has been invalidated and to prevent a spurious page fault. When a duplicate.data passes a request with a pending
30 invalidate, the packet is emptied. The first method invalidates all cells or RRCs and the second method prevents stale read only data from propagating to cells which have already been invalidated.

In order to maintain strong sequential consistency and strong ordering of events, all copies of subpage A must be invalidated prior to the invalidator updating and distributing subpage B. If the invalidate were processes serially by each Ring:0, the invalidation time would be the sum of the Ring transit times of all Ring:0s with copies of the subpage. If the invalidate were just copied by each RRC:1, invalidates of subpage A are not guaranteed to complete prior to the original invalidator updating subpage B and distributing it to processors which have not yet invalidated subpage A. Strong ordering of events is violated since all processors do not see the changes to subpages A and B in the same order.

A technique called invalidate lookahead is used to invalidate all Ring:0s in parallel, but maintain a strong ordering of events. When an invalidate packet reaches the first RRC:1, which would extract the packet, the packet is copied and forwarded to Ring:0 and a lookahead packet is spawned on on Ring:1 by setting lookahead mode in Ring:1 packet. The lookahead packet is copied by each RRC:1 in SpState:Ro to invalidate its read only copy prior to the real request packet reaching it. When the original request packet reaches subsequent RRC:1s, the packet is passed at Ring:1 speed, since the invalidate has previously completed due to the lookahead packet. The lookahead packet is always considered a short packet.

Figures 14A and 14B depict a preferred Ring:1 Invalidate Sequence. Write_invalidate is extracted from a Ring:0 (step 2) when RDU:0 Spstate is Nex and a sva_(read_ex,get,getw)) is extracted from a Ring:0 when the RDU:0 SpState is not Ex. The

packet is forwarded to RRC:1 and inserted onto Ring:1
(step 3). The first RRC:1 which owns or has a read
only copy of the subpage, copies the packet, it
changes lookahead mode to no_lookahead and forwards
5 the packet to the Ring:0 (step 5). The passed Ring:1
packet create lookahead mode is unchanged.
Subsequent RRC:1s with read only state, copy the
invalidate.lookahead packet (steps 6,7,8,9). The
invalidate_lookahead is passed to Ring:0, inserted
10 and PendlState transitions to OwnInv. When the
Ring:0 invalidation is complete, SpState transitions
to inv, Pend State transitions to none, the packet is
returned to RRC:1, RDU:1 SpState is updated to inv
and the lookahead packet is emptied (steps
15 13,14,15,16).

Meanwhile, when the first RRC:1 and
corresponding Ring:0 have completed the invalidation
or ownership response, the packet is inserted in
no_lookahead state back onto Ring:1 (step 12). When
20 this no_lookahead packet encounters the subsequent
RRC:1s (and corresponding Ring:0s) which have
previously completed the invalidation due to the
lookahead packet, the packet can be passed at full
Ring:1 speed (steps 17,18,19,20). If the actual
25 invalidate arrives before the invalidate lookahead is
complete, the actual invalidate also traverses the
Ring:0. When the lookahead packet reaches the
requesting RRC:1, the packet is emptied (step 11).
When the no_lookahead or create_lookahead packet
30 reaches the requesting RRC:1, it is extracted and
forwarded to the requesting Ring:0 (step 22) and
returned to the requestor (step 24).

False RRC SpState:Ro

Under certain circumstances, the RRC pair has recorded read only state (SpState:Ro) for a subpage, but the subpage is invalid (SpState:Inv) within the Ring:0, since no read only copies exist in any of the cells. The RRC pair deallocates the descriptor only when the last Cache within a Ring:0 deallocates the corresponding descriptor. When a Cache deallocates a descriptor, the state of the subpages within the Cache are invalid or read only. When the Cache deallocates the descriptor, read only copies are just deleted, but the RRC is not notified of the individual read only subpage involved. Thus when a cache within a Ring:0 deallocates the descriptor, with other Caches within the Ring:0 having the descriptor allocated and the only read only copy of the subpage within the Ring:0 is deleted, the RRC falsely still records read only state for the subpage. Another example is when a recombine with non-exclusive packet state leaves a Ring:0, the RRC:0 and RRC:1 must record SpState:Ro, in case a read only copy exists within the Ring:0.

False read only state is corrected by the RRC:1 picking up copy of read only data packets when in SpState:(Inv|Ro).

Duplicate distribution

Read only copies of subpages are distributed to multiple Caches in two manners. Copies are distributed implicitly by any RRC:1 or Cache copying a remotely originated packet with read only data when the descriptor is locally allocated and the SpState:(Inv|Ro). Thus the number of read only requests are reduced, since copies are distributed to

all Ring:0s connected to RRC:1s which the read only data carrying packet passes on Ring:1.

Duplicate.data packets are forwarded from the originating Ring:0 based on DupLimit. If
5 Dup:CondForward, duplicate packets are conditionally forwarded by RRC:0 to Ring:1, if the page is not exclusively owned by the local Ring:0. If
Dup:AllForward, duplicate packets are always forwarded by RRC:0 to Ring:1. Ring:1 duplicate.data
10 packets are copied by a remote RRC:1 when it has a descriptor allocated for the corresponding page within its local Ring:0 and the Extract Buffer is in the ExtBuf:All region.

Pkt.pcopy is set by RRC:1 to indicate the
15 packet was copied from Ring:1. Packets which are copied from Ring:1 with Pcopy set, are not required to be returned to RRC:1 or Ring:1 and are emptied by RRC:0 after successfully traversing Ring:0.

20 Recombine

Recombine.data packets are forwarded from the originating Ring:0 if no other cell within the Ring:0 accepts the recombine. The RRC:1 determines the packet must be forwarded out of the originating
25 Ring:0 in a similar manner to locally originated Ro request with false RRC SpState:Ro case. On the first pass of the recombine.data, the RRC:0 sets pkt.RRC. If the recombine.nodata returns to the originating cell, the recombine is extracted and closed out. If
30 the recombine.data.rrc returns to the originating cell, the cell passes the packet. The RRC:0 then extracts the recombine.data.rrc and forwards it the local RRC:1 for insertion onto Ring:1.

Ring:1 recombine.data packets are extracted
by a remote RRC:1 when it has a descriptor allocated
for the corresponding page within its local Ring:0
and the Extract Buffer is in the Extbuf:All region.
5 Recombine.data can also be used to satisfy any
pending request.

Release

When a locally originated release is
10 inserted onto a Ring:, Pend:Owner state is recorded
so that a subsequent release can be routed back to
the originating Ring:0 to satisfy possible pending
getw requests with prior already atomic response.
When a remotely originated release is inserted onto a
15 Ring:0, Pend:Owner state is recorded so that a
subsequent release can be routed back to the Ring:1.

Page Faults

Page faults are detected in a Ring:1 system
20 by a request completely traversing the Ring:1,
without any RRC detecting a descriptor match. The
packet.nodata indicates a page fault condition by
RspBsy clear and InvBsy clear. The packet is then
extracted by the originating RRC:1 and returned to
25 Ring:0 and then the requesting cell.

When a page fault occurs, the local RRC1
will be left in Pend:Owner or Pend:Ro. This
condition must be cleared by system software prior to
creation of the page. Execution of a mpdw.invsum
30 instruction from the any cell within the fault (after
the page has been anchored of course) will cause a
CiuInvSum packet to be issued.

RDU:1 Basic Subpage Rules

Packets From Connected Ring:0

5 Packet.nodata, Originated by Local Ring:0

```
    if page_not_allocated allocate page;
    if Page:PendDeallocate (extract.RspBsy.InvBsy,-) else
10  sva_read_ro:      if Pend:Inv (insert,Pend:Ro)
                        else if Pend:Owner (insert,-)
                        else if !ExtBuf:Full
                        (insert.extract.RspBsy,-);

    sva_read_ex,
    sva_get,
15  sva_getw,
    write_inv:      (insert,Pend:Owner);

    recombine,
    release,
20  duplicate:      if !ExtBuf:Full
    (insert.extract.error,-);
```

Packet.nodata, Originated by Remote Ring:0

```
25  if page_not_allocated error
    else if Page:PendDeallocate (insert.RspBsy.InvBsy,-)
    else
    sva_read_ro: if SpState:(Inv|ro) (insert,-)
30                  else (insert.RspBsy,-);

    sva_rea_ex,
    sva_get,
    sva_getw,
    write_inv:      if pkt.Lookahead
35                  if SpState:(Inv|Owner) (empty,-)
                      else if pkt. (!InvBsy&!RspBsy)
                      (empty,SpState:Inv) else if
                      !ExtBuf:Full
                      (insert.extract.!RspBsy.!InvBsy,-)
40                  else /* !Lookahead */
                      if SpState:Inv (insert,-)
                      else if SpState:Owner
                      (insert.RspBsy,-)
                      else if pkt.(!InvBsy&!RspBsy)
45                  (insert,SpState:Inv) else
                      (insert,-);

    recombine,
    release: if SpState:Owner (insert,-)
              else (insert.error,-)
50  duplicate:      (insert.error,-);
```


Packet.data

```
if page_not_allocated error
else if Page:PendDeallocate
5      if !OriginateLocal
        (insert.RspBsy.InvBsy,-)
        else if !ExtBuf:Full
          (insert.extract.RspBsy.InvBsy,-)
        else
10 sva_read_ro,
    duplicate:  if SpState: ( RoEx) (insert,-)
                else error,
    release: if OriginateLocal
              if SpState:Owner
15              (insert,SpState:Inv.Pend:Owner)
            else if !ExtBuf:Full
              (insert.extract.error,-)
              else/* !OriginateLocal */
                if SpState:Owner
20                (insert,SpState:Inv)
              else if !ExtBuf:Full
                (insert..error,-)
    recombine:  if SpState:Owner
                 if PktState.nex (insert,SpState:Ro)
25                 else (insert,SpState:Inv);
                 else if OriginateLocal
                   if ExBuf.Full
                     (insert.extract.error,-);
                   else (insert.error,-);
30 sva_read_ex: if OriginateLocal
                 if SpState:Owner
                   if PktState.nex
                     (insert,SpState:Ro)
                     else (insert,SpState:Inv)
35                 else if
                   !ExtBuf:Full(insert.extract.error,-)
                   else/* !OriginateLocal */
                     if pkt.InvBsy (insert,SpState:Ro)
                     else (insert,SpState:Inv)
40                 else (insert.error,-);
```

Packets From Ring:1

45 Packet.nodata, Originated by Local Ring:0

```
if page_not_allocated error
else if Page:PendDeallocate error
else
```



```

    sva_read_ro:  if ExtBuf:Full (pass.!Timestamp,-)
                  else if Pendl:Reissue
                  (extract.RspBsy,-) else
                  (extract,-);
5  sva_read_ex,
   sva_get,
   sva_getw,
   write_inv:    if pkt.Lookahead (empty,-)
                  else if ExtBuf:Full
10                  (pass.!Timestamp,-) else if
                  Pendl:Reissue (extract.RspBsy,-)
                  else if SpState:Owner
                  (extract,Pend:None) else
                  (extract,-);
15 recombine,
   release: if ExtBuf:Full (pass.!Timestamp,-)
            else (extract,-);
   duplicate: (empty.error,-)

20 Packet.nodata, Originated by Remote Ring:0

   if page_allocated
       if Page:PendDeallocate
25       (pass.RspBsy.InvBsy,-)
       else (

   sva_read_ro:  if SpState:Inv (pass,-)
                  else if ExtBuf: (All|Opt)
30                  (extract,-)
   sva_read_ex,
   sva_get,
   sva_getw,
   write_inv:    if pkt.Lookahead
35                  if SpState:Ro &
                  ExtBuf{All|NonOpt}
                  (copy.!InvBsy,-) else (pass,-)
                  else if pkt.(!InvBsy&!RspBsy)
                  if SpState:(Ro|Owner)
40                  if ExtBuf: (All|NonOpt)
                  (pass.Lookahead.,copy,-)
                  else (pass.RspBsy.InvBsy,-)
                  else (pass,-)
                  else if SpState:Owner
45                  (pass.RspBsy,-)
                  else (pass,-);

   recombine,
   release: (pass,-)
   duplicate: (pass.error,-);
50
```


Packet.data

```
if page_not_allocated error
else if Page:PendDeallocate error
5 else (
  sva_read_ro: if OriginateLocal
                if pkt.pcopy (empty,-)
                if ExtBuf:Full
                (pass.!Timestamp,-)
10             else if Pend:Ro
                if SpState:Inv
                (extract,SpState:Ro.Pend:None.
                Pendl:None)
                else
15             (extract,Pend:None.Pendl:None)
                ;
                else
                if SpState:Inv
                (extract,SpState:Ro);
20             else (extract,-);
            else /* !OriginateLocal */
            if Pend:Owner (pass.RspBsy,-)
            else if Pend:Ro
            if ExtBuf:All
25             (copy,SpState:Ro,Pend:None.Pen
            d1:None)
            else if !SpState:Owner
            (copy,SpState:Ro)
            else (pass,-);
30 duplicate:  if OriginateLocal (empty,-);
                else if Pend:Ro
                if ExtBuf:All
                If SpState:Inv
                (copy,Spstate:Ro,Pend:None.Pen
35             d1:None)
                else
                (copy,Pend:None,Pendl:None);
                else if Pend:Owner (empty,-)
                else (copy,-);
40             /*dependent on false ro
                policy*/
            sva_read_ex: if OriginateLocal
                if ExtBuf-Full (pass.!Timestamp,-)
                else if Pend:None
45             (extract,SpState:Owner);
                else
                (extract,SpState:Owner.Pend:None.Pe
                ndl:None);
            else /* !OriginateLocal */
```



```

                                if Pend:None
                                (extract,SpState:Owner);
                                else (pass,Pendl:Reissue);
recombine:    if OriginateLocal
5              if ExtBuf:Full
                (pass.!Timestamp,-)
                else if Pend:None
                (extract,SpState:Owner);
                else
10              (extract,SpState:Owner.Pend:None.Pendl:None);
              else/* !OriginateLocal */
                if Pend: (Ro|Owner)
                if !Extbuf:Full
15              (extract,SpState:Owner.Pend:None.Pendl:None)
                else (pass,Pendl:Reissue);
                else if page_allocated
                if ExtBuf:All
20              (extract,SpState:Owner)
                else (pass,-);
release:    if OriginateLocal
              if Extbuf:Full
                (pass.!Timestamp,-)
25              else if Pend:None
                (extract,SpState:Owner);
                else
                (extract,SpState:Owner.Pend:None.Pendl:None);
30              else/* ! OriginateLocal */
                if Pend: (Ro|Owner)
                if !Extbuf:Full
                (extract,SpState:Owner.Pend:None.Pendl:None)
35              else (pass.rrc,-);
                else (pass,-);

    sva_get,
    sva_getw:    if Originate Local
40              if ExtBuf:Full
                (pass.!Timestamp,-)
                else (extract,-);
              else (pass,-);

45
```


RDU:0 Basic Subpage Rules

Packets From Connected Ring:1

5 Packet.nodata, Originated by Remote Ring:0

```
    if page_not_allocated error
    else if Page:PendDeallocate
        if !ExtBuf:Full
10         (insert.extract.RspBsy.InvBsy,-)
        else
            sva_read_ro:  if Pend:None (insert,Pend:Ro)
                        else (insert,-);
            sva_read_ex,
15 sva_get,
            sva_getw,
            write_inv:  if pkt.Lookahead
                        if SpState:Inv|Pend:Owner (empty,-)
                        else (insert,Pend:Owner);
20         else /* !Lookahead */
            if SpState:Inv & !ExtBuf:Full
                (insert.extract,-)
            else (insert,Pend:Owner);
            recombine,
25 release,
            duplicate:  if !ExtBuf:Full (insert.extract.error,-)
```

Packet.nodata, Originated by Local Ring:0

```
30 /* not effected by Page:PendDeallocate state */
    if page_not_allocated error else
        sva_read_ro:  (insert,-);
        sva_read_ex,
35 sva_get,
        sva_getw,
        write_inv:  if pkt.(InvBsy|RspBsy) & !ExtBuf:Full
                    (insert.extract.!RspBsy.!InvBsy,-)
                    else if SpState:Nex (insert,SpState:Ex)
40         else (insert,-);
        recombine,
        release:  (insert,-)
        duplicate:  (pass.error,-)
```

45

Packet.data

```
    if page_not_allocated error
    else Page:PendDeallocate
50         if !OriginateLocal
```



```

                                if !ExtBuf:Full
                                (insert.extract.RspBsy.InvBsy,-)
else (
  sva_read_ro,
5 duplicate:  if SpState:Inv (insert,SpState:Ro)
               else if SpState: (Ro|Nex) (insert,-)
               else if OriginateLocal (insert.error,-)
               else if !ExtBuf:Full
               (insert.extract.error,-)
10 sva_read_ex:  if SpState:(Inv|Ro)
                  if PktState.nex
                    if pkt.!InvBsy (insert,SpState:Ex);
                    else (insert,SpState:Nex)
                    else (insert,SpState:Ex);
15               else if OriginateLocal (insert.error,-)
               else if !ExtBuf:Full
               (insert.extract.error,-);
  recombine:  if SpState:Inv
                if PktState.nex (insert,SpState:Nex)
20               else (insert,SpState:Ex)
               else if SpState:Ro & PktState:Nex
               (insert,SpState:Nex)
               else if OriginateLocal (insert.error,-)
               else if !ExtBuf:Full
25               (insert.extract.error,-);
  release:    if OriginateLocal
                if SpState:Inv (insert,SpState:Ex)
                else (insert.error,-);
               else/* !OriginateLocal */
30               if SpState:Inv
               (insert,SpState:Ex,Pend:Owner)
               else if !ExtBuf:Full
               (insert.extract.error,-);
               /* Same as
35 sva_get,
   sva_getw:  (insert,-)
   )

```

40 Packets From Ring:0

Packet.nodata, Originated by Remote Ring:0

```

45 if page_not_allocated error
   else if Page:PendDeallocate
       if !ExtBuf:Full (extract.RspBsy.InvBsy,-)
       else (pass.!Timestamp.RspBsy.InvBsy,-)
   else (

```



```
sva_read_ro:    if !ExtBuf:Full (extract,-)
                else (pass.!Timestamp,-);
sva_read_ex:    if pkt.Lookahead
                if SpState:(Nex|Ex) (empty,-)
5               else if ExtBuf:Full (pass.!Timestamp,-)
                else if SpState:Inv
                  (extract,Pend:None);
                else if pkt.!InvBsy
                  (extract,SpState:Inv.Pend:None)
10              else (extract,-)
                else /* !Lookahead */
                  if ExtBuf:Full (pass.!Timestamp,-)
                  else if SpState: (Nex|Ex ) (extract,-)
                  else if SpState:Inv (extract,Pend:None)
15              else if pkt.  [!InvBsy&!RspBsy]
                  (extract,SpState:Inv.Pend:None)
                  else (extract,-);
    release,
    recombine:  if ExtBuf:Full (pass.!Timestamp,-)
20              else (extract,-);
    duplicate:  if !ExtBuf:Full (extract.error,-)
                )

25 Packet.nodata, Originated by Local Ring,0

    /* not effected by Page:PendDeallocate state */
    if page_not_allocated & pkt.!Invis allocate_page;
    (
30 sva_read_ro:  if pkt.rrc
                  if SpState:(Nex|Ex) (pass,-)
                  else if ExtBuf:(All|NonOpt) (extract,-)
                    , else (pass,-)
                  else /* !rrc */
35                  if SpState:(Ro|Nex|Ex) (pass.rrc,-)
                  else if ExtBuf(All|NonOpt) (extract,-)
                  else (pass.RspBsy,-);
    sva_read_ex,
    sva_get,
40 sva_getw:     if SpState:Ex (pass,-)
                  else if pkt.InvBsy (pass.RspBsy,-)
                  else if ExtBuf:(All|NonOpt)(extract,-)
                    else (pass.RspBsy,-);
    write_inv:   if SpState:Ex (pass,-)
45              else if pkt.InvBsy (pass.RspBsy,-)
                  else if SpState:Nex
                    if ExtBuf:(All|NonOpt) (extract,-)
                    else (pass.InvBsy,-);
                  else (pass.error,-)
```



```
release,
recombine: (pass,-);
duplicate: (pass.error,-)
)
5

Packet.data

if page_not_allocated error
10 else Page:Penddeallocate
    if !OriginateLocal
        if !ExtBuf:Full
            (extract.RspBsy.InvBsy,-)
        else (pass.!TimeStamp,-);
15 else
    sva_read_ro: if OriginateLocal
        if Pend:Owner (pass.RspBsy,-)
        else if Pend:None (pass,-)
        else if !ExtBuf:Full
20         if SpState:Ex
            (pass,copy.pcopy,SpState:Nex.Pend:N
            one)
        else
            (pass,copy.pcopy,Pend:None)
25         else pass;
        else /* !OriginateLocal */
            if pkt.pcopy (empty,-)
            else if ExtBuf:Full
                (pass.!Timestamp,-)
30         else if Pend:Ro
            if SpState:Ex
                (extract,SpState:Nex.Pend:None
                )
            else (extract,Pend:None);
35         else
            if SpState:Ex
                (extract,SpState:Nex)
            else (extract,-);
duplicate: if OriginateLocal
40         if Pend:Owner (empty,-)
        else/* !OriginateLocal */
            if Pend:Ro & ExtBuf: (All|NonOpt)
            if SpState:Ex
                (pass,copy.pcopy,SpState:Nex.Pend:N
25         one)
            else (pass,copy.pcopy,Pend:None);
        else if (!page_all_ex&Dup:ConForward)
            |Dup:AllForward
                if ExtBuf:All
```



```

                                if SpState:Ex
                                (pass,copy.pcopy,SpState:Nex)
                                else (pass,copy.pcopy,-);
                                else (pass,-);
5 sva_read_ex:  if OriginateLocal
                  if Pend:None (pass,-)
                  else if !ExtBuf:Full
                  if PktState:Nex
10                  (extract,SpState:Ro.Pend:None)
                  else
                  (extract,SpState:Inv.Pend:None);
                  else (pass,-)
                  else/* !OriginateLocal */
15                  if ExtBuf:Full (pass.!Timestamp,-)
                  else if pkt.(!RspBsy&!InvBsy)
                  (extract,SpState:Inv.Pend:None)
                  else (extract,SpState:Ro);
release:  if OriginateLocal
20          if Pend:None (pass,-)
          else if !ExtBuf:Full
          (extract,SpState:Inv.Pend:None)
          else (pass.rrc,-);
          else/* ! OriginateLocal */
25          if ExtBuf:Full (pass.!Timestamp,-)
          else (extract,SpState:Inv.Pend:None);
recombine:  if OriginateLocal
              if Pend: (Ro|Owner) & !ExtBuf:Full
              if PktState:Nex
30              (extract,SpState:Ro.Pend:None)
              else
              (extract,SpState:Inv.Pend:None);
              else if pkt.!rrc (pass.rrc,-)
35              else if !ExtBuf:Full
              if PktState:Nex
              (extract,SpState:Ro)
              else (extract,SpState:Inv);
              else (pass,-);
40              else/* !OriginateLocal */
              if ExtBuf:Full (pass.!Timestamp,-)
              else if Pend:(Ro|Owner)
              if PktState:Nex
              (extract,SpState:Ro.Pend:None)
45              else
              (extract,SpState:inv.pend:None);
              else
              if PktState:Nex
              (extract,SpState:Ro)
50              else (extract,SpState:Inv);
```



```

sva_get,
sva_getw:  if OriginateLocal (pass,-)
            else if ExtBuf:Full (pass.!Timestamp,-)
            else (extract,-);
5

```

Page Operations

Page Allocate

10 A descriptor is allocated by the local cell
prior to any cell SVA request. The RRC pair
allocates a descriptor when the first locally
originated request packet (always nodata) is received
from the local Ring:0. If a page fault was signalled
15 to a request, the descriptor remains allocated in the
cell and RRC pair. One or more descriptors can be
allocated for pages not present in the memory
system. A descriptor must be allocated for the page
to be created. The following table specifies the
20 default state of an allocated descriptor:

	<u>Descriptor Field</u>	<u>Default State</u>
25	Descriptor Tag	Pkt page address
	Subpage State	Invalid
	OwnerLimit	OL:All
	DupLimit	Dup:CondForward
	Pending Page State	Page:None

30 Page Deallocate

Local cache and local RRC pair descriptors
are only deallocated by the CiuInvSumAlloc.nodata,
CiuPureInvSumAlloc.nodata or RrcDeallocDesc.nodata
35 packets. These three packets will be referred to as
deallocate packets. If the deallocate packet
indicates that no other caches within the Ring:0 have
the page allocated, the local RRC pair deallocates
the descriptor corresponding to the page. RRC:0

transitions to Page:PendDealloc while it is determining if the descriptor should be deallocated to prevent race conditions between multiple deallocate packets and other page oriented packets.

- 5 If no other cells have allocated descriptors for the page, the RRC pair deallocates the descriptor. Both RRCs must be in Page:PendDeallocate state prior to either RRC deallocating the descriptor. The local RRC pair does not signal remote RRCs that descriptors
10 have been deallocated. First the deallocation sequence is described, then interactions with subpage and page operations are discussed.

The deallocate sequence is initiated by the CCU sending a CiuLoadDropReg packet to the CIU. The
15 packet address and CCU prt entry correspond to the page to be deallocated. The CIU processes the packet by loading the address into its dropping register, extracting the packet to return it to the CCU. When loaded with a valid address, the CIU dropping
20 register causes the CIU to assert RspBsy to requests for owned subpages and ignore all other page or subpage address which matches the dropping register. Dropping register match overrides cache group match.

- When the CCU receives the returned
25 CiuLoadDropReg packet, it deallocates the old descriptor by writing the new descriptor in its place, deallocates the prt entry and issues a CiuInvSumAlloc packet if the deallocated page is not pure. If the deallocated page is pure,
30 CiuPureInvSumAlloc packet is issued. The deallocate packet address is the page address to be deallocated (pkt.address[5:0] contains the CCU descriptor group) and dataword0 contains the page address being allocated. A pending CEU request can be processed,

even at the newly allocated address, since no prt is required for the deallocate packet. The mrdd (RRC deallocate descriptor) instruction causes the CCU to issue a RrcDeallocDesc packet without allocating a
5 prt.

Figure 15 depicts a preferred deallocate sequence used in a system constructed in accord with the invention. The CIU processes the CiuInvSumAlloc or CiuPureInvSumAlloc packet by deallocating the old
10 descriptor by writing the newly allocated descriptor in its place, setting all subpages of new allocated page to invalid (using summaries), setting dropping register to invalid and inserting the packet (step 1). The CIU processes the RrcDeallocDesc packet by
15 inserting the packet without any modification to CIU directory state.

As the deallocate packet traverses the Ring:0, RspBsy is asserted by any CIU which has the descriptor allocated (steps 2,4,6).

20 The RRC:0 empties the CiuInvSumAlloc or RrcDeallocDesc request if RspBsy is asserted, indicating that the descriptor is allocated in other caches within the Ring:0 (step 3). If RspBsy is not asserted, the RRC:0 the deallocate packet must now
25 completely traverse the Ring:0 to check all local caches for this page being allocated, since the packet has only traversed the portion of the Ring:0 between the deallocating cache and the RRC:0. The RRC:0 passes the packet, setting the RRC bit,
30 clearing RspBsy and InvBsy and transitions to Page:PendDeallocate. The RRC bit is used to indicate that the packet has passed the RRC:0. If the RRC cannot process the packet due to refresh or ecc error it sets InvBsy and processes the deallocate packet on

the next revolution. Subsequent deallocate.!RRC
packets received, while PagePendDeallocate are
emptied, since the cells which issued these
deallocate packets will not assert RspBsy on the
5 current revolution of the first deallocate packet.
The reason that the RRC:0 negates RspBsy on the first
deallocate packet is handle the window where a cell
which issued the subsequent deallocate packet
asserted RspBsy on the first partial revolution of
10 the first deallocate packet.

If the deallocate packet is
CiuPureInvSumAlloc, the RRC:0 extracts the packet,
sets pkt.cmd_pend_dealloc and forwards the packet to
the local RRC:1 skipping to step 8 (RRC:1
15 processing). CiuPureInvSumAlloc can only be issued
by a cell which owns all subpages within the page,
guaranteeing that any other allocated descriptors
have all subpages invalid. Thus the RRC:0 skips
sending the CiuPureInvSumAlloc packet around the Ring
20 to determine if any cells have the page allocated.
The page will be allocated by the RRC is one of the
cells within the Ring:0 subsequently requests a
subpage within the page.

All remotely originated packets received
25 from Ring:1 when in Page:PendDealloc state, are
returned unmodified to the Ring:1. If a sva_read_ro
packet is returned (Ring:0 state was SPState:Ro),
another Ring:0 will supply a copy of the subpage. A
recombine.data is returned since the cell and RRC
30 pair descriptors are about to be deallocated. Page
destroy sequences always set all subpages to invalid
prior to deallocation, so that SpState:{Nex,Ex} will
not be present.

If the configuration does not contain an RRC, both InvBsy and RRC will be negated. The original requestor empties the deallocate packet if packet.request_id matches cell_address, !InvBsy and
5 !RRC (step 5).

The RRC:0 empties the deallocate.RRC packet after it full Ring:0 traversal (step 7), if caches have indicated the page is allocated by asserting RspBsy. Negation of RspBsy indicates that no caches
10 within the Ring:0 have the page allocated. The RRC:0 extracts the deallocate packet, sets pkt.cmd_pend_dealloc and forwards the packet to the local RRC:1. Both RRC:0 and RRC:1 must be in Page:PendDealloc prior to the descriptor being
15 deallocated to prevent race conditions.

The RRC:1 processes the packet by transitioning to Page:PendDealloc, sets pkt.cmd_dealloc_page, extracts the packet and forwards the packet back to RRC:0 (step 8). All
20 remotely originated packets received from Ring:1 when in Page:PendDealloc state, are returned unmodified to the Ring:1.

The RRC:0 processes the packet by deallocating the descriptor, transitioning to
25 Page:None, extracts the packet and forwards the packet back to RRC:1 (step 9). The RRC:1 processes the returning packet by also deallocating the descriptor, transitioning to Page:None, and emptying the packet.

30 Page and subpage oriented packets are handled specially when in Page:PendDeallocate state. Other page level packets, such as CiuInvSum.nodata, CiuInvSumAllocate.no_data and CiuExSum.nodata are busied. Remotely originated subpage requests are

busied by RRC:0 and RRC:1. Locally originated subpage requests are unaffected in RRC:0, since the requesting local cache has an allocated descriptor will prevent the RRC pair from deallocating its
5 descriptor. Locally originated subpage requests are busied by RRC:1, since the RRC pair has committed to deallocate the descriptor. The descriptor will be subsequently re-allocated on the next request after the page has been deallocated. Locally originated
10 packets which are forwarded by RRC:0 between steps 7 through 9, will also be busied by the RRC:1. Locally originated requests processed after step 9 by RRC:0, are guaranteed to be processed after step 10 by RRC:1, and the descriptor will be re-allocated.

15

Page Create

The memory system assumes that system software guarantees that an existing page in memory is not created. System software must also guarantee
20 that a page is not simultaneously created by more than one cache. Memory System hardware does not immediately detect these conditions.

Mpsa instruction provides an anchored descriptor in the local cache for an SVA page. If
25 the descriptor already existed prior to the mpsa, its anchor flag is set. Otherwise, the local cache allocates a descriptor as described under page deallocate/reallocate section above.

The local cache then issues a
30 ciu_ex_sum.nodata packet in as a result of a mpdw.allex instruction and transitions all subpages within the page to owned. The RRC:0 extracts the ciu_ex_sum.nodata packet, allocates the descriptor if necessary and forwards the request to RRC:1. If

Page:PendDealloc, the RRC:0 passes the request and sets pkt.InvBsy. RRC:1 extracts the packet, allocates the descriptor if necessary, transitions all subpages to SpState:Owner and forwards the
5 request back to RRC:0. RRC:0 inserts the packet, and transitions all subpages within the page to SpState:Ex. The originating CIU extracts the packet, returns the packet to the CCU and the CCU returns completion status to the local processor.
10 A mpdw instruction is then issued to clear the anchor bit.

Page Destruction Issues

Pages can be destroyed in three manners.
15 Any time a page is destroyed, all subpages within the page must be exclusively owned by the local Cache. First, when pure pages are deallocated, they are destroyed. The CiuPureInvSumAlloc packet indicates that a pure page is being deallocated. Secondly,
20 when a descriptor's address is changed, the old page is destroyed, but the contents of the old page are present in the new page. The RrcDeallocDesc packet indicates that the local RRCs should dealloc the page. Thirdly, the page and its contents are
25 explicitly destroyed. The CiuInvSum packet indicates that the local RRCs should set all subpages to invalid state.

When a page is destroyed, the CiuPureInvSumAlloc, RrcDeallocDesc and CiuInvSum
30 packets traverse the Ring:1 to clear out any remote pending or Reissue state for the page. Three separate packet types are required, since the local Cache and RRC actions for the three cases are different.

Page Destroy

Mpsa instruction provides an anchored descriptor in the local cache for an SVA page. If the descriptor already existed prior to the mpsa, its anchor flag is set. Otherwise, the local cache allocates a descriptor as described under page deallocate/reallocate section above.

The mfsva instruction is used to cause the local cache to obtain exclusive ownership of all subpages.

The local cache then issues a ciu_inv_sum.nodata packet in as a result of a mpdw.allinv instruction and transitions all subpages within the page to invalid. The RRC:0 extracts the ciu_inv_sum.nodata packet, and forwards the request to RRC:1. If Page:PendDealloc, the RRC:0 passes the request and sets pkt.InvBsy. RRC:1 extracts the packet, transitions all subpages to SpState:Inv and forwards the request back to RRC:0. RRC:0 inserts the packet, and transitions all subpages within the page to SpState:Inv. The originating CIU extracts the packet, returns the packet to the CCU and the CCU returns completion status to the local processor.

The page destroy sequence and ciu_inv_sum.nodata in particular do not deallocate the local cache or local RRC pair descriptors. The local cache and local RRC pair descriptors are not deallocated until the descriptor is reused.

30 Change Descriptor

Change descriptor page address is logically a destroying the old page followed by creating a new page which happens to have the same data. The memory system assumes that system software guarantees that

an existing page in memory is not created. System software must also guarantee that a page is not simultaneously created by more than one cache. Memory System hardware does not immediately detect
5 these conditions.

Change descriptor used three instructions, mpdw, mrdd and mrwd. Mpdw instruction side effects the local cell (CCU and CIU) only and is not inserted on the Ring:0.

10 The RRC Deallocate Descriptor (mrdd) instruction from the CCU reserved memory system opcodes. The CCU issues a RrcDeallocDesc.nodata packet when decoding this instruction from the CEU. !m0 supplies the page address to be deallocated, in
15 an identical manner to other memory system instructions. The CIU inserts this packet onto the Ring:0. The RRC copies the packet and initiates a deallocation sequence. If the RRC is unable to copy or process the packet, it sets packet.RspBsy
20 (responder busy). The CIU extracts the returning RrcDeallocDesc.nodata packet and passes it to the CCU. If RspBsy is set the request is reissued by the CCU. If RspBsy is clear, the CCU returns a successful completion status to the CEU. If an RRC:0
25 is not configured, the returning packet, RrcDeallocDesc.nodata.!RspBsy, is identical to the original request, and causes the CCU to return a successful completion status to CEU. This instruction has no effect on CCU or CIU data
30 structures. The addition of this instruction is transparent to the CEU.

The RRC Write Descriptor (mrwd) instruction from the CCU reserved memory system opcodes. The CCU issues a RrcWrDesc.nodata packet when decoding this

instruction from the CEU. !m0 supplies the page address to be allocated, in an identical manner to other memory system instructions. The CIU inserts this packet onto the Ring:0. The RRC:0 extracts the
5 packet. The local RRC:1 allocates the descriptor if required, sets the descriptor all exclusive, writes RRC resident descriptor bits (ownerlimit) and extracts the packet, sending it back to the Ring:0. The RRC:0 allocates the descriptor if required, sets
10 the descriptor all exclusive, writes RRC resident descriptor bits (ownerlimit) and inserts the packet onto Ring:0. If the RRCs are unable to process the packet, packet.RspBsy is set. The CIU extracts the returning RrcWrDesc.nodata packet and passes it to
15 the CCU. If RspBsy is set the request is reissued by the CCU. If RspBsy is clear, the CCU returns a successful completion status to the CEU. If an RRC:0 is not configured, the returning packet, RrcWrDesc.nodata.!RspBsy, is identical to the
20 original request, and causes the CCU to return a successful completion status to CEU. This instruction has no effect on CCU or CIU data structures. The addition of this instruction is transparent to the CEU.

- 25 The change descriptor sequence is:
- 1 - Anchor old page (mpsa).
 - 2 - Gather all subpages of old page with exclusive ownership.
 - 3 - Deallocate RRC descriptor (mrdd) (old
30 page).
 - 4 - Write new RRC descriptor (mrwd) (new page).
 - 5 - Write new descriptor in cell and clear anchor bit (mpdw).

35

Mrdd and mrwd instructions are not required if the tag portion of descriptor is not being changed. Additionally, gathering all subpages with exclusive ownership might not be required to change
5 some descriptor bits as discussed above. Atomic and transient atomic state are preserved in the Cell. The RRC subpage state only records non-exclusive and exclusive ownership with atomic and transient atomic falling in the later category. The RRC does maintain
10 pending getw state which is lost by the change sequence. The getw requestor(s) will timeout.

RDU:0 Page Operations Rules

Allocation rules are enumerated under RDU:0
15 Basic Rules.

Packets From Connected Ring:1

Packet.nodata, Originated by Remote Ring:0

20

CiuInvSumAlloc,
CiuPureInvSumAlloc,
RrcDeallocDesc,
RrcWrDesc,
25 CiuUpdateState,
CiuWrDesc
CiuInvSum,
CiuExSum: error,

30

Packet.nodata, Originated by Local Ring:0

CiuInvSumAlloc,
CiuPureInvSumAlloc,
35 RrcDeallocDesc: if page_not_allocated error
else if pkt.cmd_dealloc_page
(extract,Page:None.deallocate page)
else error;


```
RrcWrDesc,  
CiuInvSum,  
CiuExSum:      (insert,-);  
CiuUpdateState,  
5 CiuWrDesc:    error;
```

Packet.data

```
10 CiuPureInvSumAlloc,  
   CiuInvSumAlloc,  
   RrcDeallocDesc,  
   RrcWrDesc,  
   CiuUpdateState,  
15 CiuWrDesc  
   CiuInvSum,  
   CiuExSum:      error;
```

20 Packets From Ring:0

Packet.nodata, Originated by Remote Ring:0

```
   CiuInvSumAlloc,  
25 CiuPureInvSumAlloc,  
   RrrDeallocDesc,  
   RrcWrDesc,  
   CiuUpdateState,  
   CiuWrDesc,  
30 CiuInvSum,  
   CiuExSum:      error;
```

Packet.nodata, Originated by Local Ring:0

```
35 CiuInvSumAlloc,  
   RrcDeallocDesc:  if page_not_allocated error  
                    else if Page:PendDealloc  
                      if pkt.rrc  
40                      if pkt.RspBsy (empty,-)  
                        else (extract.cmd_pend_  
                             dealloc,-)  
                        else (empty,-)  
                      else if pkt.RspBsy (empty,-)  
45                      else  
                        (pass.!RspBsy.!InvBsy.rrc,Page:PendDe  
                        alloc);  
   CiuPureInvSumAlloc:  
                        if page_not_allocated error
```



```

else if !Page:PendDealloc
  (extract.cmd_pend_dealloc,-)
else (pass.InvBsy,-);
5  RrcWrDesc:  if page_not_allocated allocate_page;
               if Page:PendDealloc (pass.InvBsy,-)
               else (extract,write desc);
               CiuInvSum:  if page_not_allocated allocate_page;
                           if Page:PendDealloc (pass.InvBsy,-)
10              else (extract, set all subpages
                           invalid);
                           CiuExSum:  if page_not_allocated allocate_page;
                                       if Page:PendDealloc (pass.InvBsy,-)
                                       else (extract, set all subpages
                                       exclusive);
15  CiuUpdateState,
     CiuWrDesc:  error,
```

```

Packet.data
20  CiuInvSumAlloc,
     CiuPureInvSumAlloc,
     RrcDeallocDesc,
     RrcWrDesc,
25  CiuUpdateState,
     CiuWrDesc,
     CiuInvSum,
     CiuExSum:  error,
```

```

30  RDU:1 Page Operations Rules
      Allocation rules are enumerated under RDU:1
      Basic Rules.
```

```

35  Packets From Connected Ring:0
```

```

Packet.nodata, Originated by Local Ring:0
40  CiuInvSumAlloc,
     RrcDeallocDesc:  if page_not_allocated error
                       else if pkt.cmd_pend_dealloc
```



```

                                if Page:None
                                (extract.cmd_dealloc_page,Page:PendDealloc)
                                else error,
5                                else if pkt.cmd_dealloc
                                if Page:PendDealloc
                                (empty,Page:None.deallocate page)
                                else error;
                                else error,
10 CiuPureInvSumAlloc,
   RrcDeallocDesc: if page_not_allocated error
                   else if pkt.cmd_pend_dealloc
                       if Page:None (insert,-)
                       else error,
15                   else if pkt.cmd_dealloc
                       if Page:PendDealloc
                       (empty,Page:None.deallocate page)
                       else error,
                   else error;
20 RrcWrDesc:      if page_not_allocated allocate_page;
                   if Page:None (extract,write desc)
                   else error;
                   CiuInvSum: if page_not_allocated allocate_page;
                               if Page:None (insert,-)
25                               else (extract.error,-);
                               CiuExSum: if page_not_allocated allocate_page;
                                         if Page:None (extract, set all
                                         subpages exclusive)
                                         else error-,
30 CiuUpdateState,
   CiuWrDesc:      error,

Packet.nodata, Originated by Remote Ring:0
35 CiuPureInvSumAlloc,
   CiuInvSumAlloc,
   RrcDeallocDesc,
   RrcWrDesc,
40 CiuUpdateState,
   CiuWrDesc,
   CiuInvSum,
   CiuExSum:      error,

45 Packet.data

   CiuPureInvSumAlloc,
   CiuInvSumAlloc,
50 RrcDeallocDesc,
```



```
RrcWrDesc,
CiuUpdateState,
CiuWrDesc:      error,

5
  Packets From Ring:1

    Packet.nodata, Originated by Local Ring:0

10  CiuPureInvSumAlloc,
    RrcDeallocDesc:  if page_not_allocated error
                     else if pkt.cmd_pend_dealloc
                       (extract.cmd_dealloc_page,
                        Page:PendDealloc)
15  CiuInvSum:        else (extract.error,-);
    CiuInvSumAlloc,  (extract, set all subpages invalid)
    RrcWrDesc
    CiuExSum,
20  CiuUpdateState,
    CiuWrDesc:      error,

    Packet.nodata, Originated by Remote Ring:0

25  CiuPureInvSumAlloc,
    RrcDeallocDesc,
    CiuInvSum:        (pass,set all subpages invalid)
    CiuInvSumAlloc,
30  RrcWrDesc
    CiuExSum,
    CiuUpdateState,
    CiuWrDesc:        (pass.error,-);

35
    Packet.data

    CiuPureInvSumAlloc,
    CiuInvSumAlloc,
40  RrcDeallocDesc,
    RrcWrDesc,
    CiuUpdateState,
    CiuWrDesc
    CiuInvSum,
45  CiuExSum:        error;
```


Special Operations

RRC Extract Buffer Full

An Extract Buffer flow control algorithm is required since the peak packet rate exceeds the rate the RRC can empty the buffer. The control algorithm is based on the principle of maximizing forward progress in conjunction with reducing system loading. Prioritizing the allocation of system resources to completing requests in progress achieves both of these goals since completing a request simultaneously reduces system loading and achieves most efficient forward progress.

A packet which cannot be extracted or copied into the Extract Buffer is handled based on whether the packet is Originated by this RRC. If the packet is not originated by this RRC, it is handled based on whether the packet does not carry data, carries read only data or carries ownership. The packet busy bits, RspBsy and InvBsy (conveniently situated late in the packet), and SpState and Pend1 subpage state fields are used to cause the standard protocol to do the right thing.

For all packets which were originated by this RRC, pkt.TimeStamp is cleared so that the packet can circle the Ring again.

For all packets which were not originated by this RRC, pkt.TimeStamp is also cleared so that the packet can circle the Ring if necessary. All page operations are considered as no data request packets. For packets not carrying data, both RspBsy and InvBsy are set. RspBsy and no data are set for sva_read_ro.data responses and duplicates are emptied. Ownership packets are passed and RRC:1

transitions to Pend1:ReIssue. Pend1:Reissue state is used by RRC:1 to prevent spurious missing_page_error (page fault) by differentiating between the case of a remotely originated ownership packet that wasn't
5 extracted to satisfy a pending request and a missing page. When set Pend1:Reissue indicates that the RRC:1 might have missed a potential non-RRC originated response packet to an RRC originated request. If the RRC:1 originated request returns
10 without a response and Pend1:ReIssue then the returning request is busied by the RRC causing the requesting processor cell or RRC to reissue the request, thereby preventing a spurious page fault. RRC:0 uses the SpState field to prevent
15 missing_page_error. The RRC:0 SpState field records whether a read only copy or ownership is within the Ring:0 to satisfy the request. If the RRC:0 originated request returns without a response and a sufficient copy or ownership exists within the
20 Ring:0, the request is reissued.

The following paragraphs analyze the behavior of all cases.

RRC:0 From Ring:1; and RRC:1 From Ring:0

25 Packets received by RRC:1(0) from Ring:0(1) are usually inserted onto the Ring:1(0). Under certain circumstance the packet is re-routed back to the other RRC. This is accomplished by immediately extracting the inserted packet. An example is a
30 locally originated request of the same type as a pending request (sva_read_ex while in Pend1:Ro) processed in RRC:1. RspBsy is asserted and the request is returned to the local Ring:0.

If the Extract Buffer cannot accept the re-routed packet, it is held in the insert buffer until the extract buffer is not full and an empty packet is available. The rules notation
5 insert.extract, indicates a re-routed packet which follows this behavior.

RRC:0 From Ring:0, Packets Originated by Local Ring:0

10	Request packet	Will return to originator (maybe getting satisfied along the way) and get re-issued if necessary.
	ReadOnly Data	Will return to originator (maybe getting satisfied along the way) and get reissued if necessary.
15	Carrying packet	
	Ownership Data	Will return to originator to satisfy request. But if there was a pending request from another Ring:0 this packet should have been snarfed. When that packet returns to RRC:1, the RRC:1 will notice that ownership is still
20	Carrying packet	
		contained within the Ring:0 and return the request back to Ring:0.
25		

RRC:0. From Ring:0.
30 Packets Originated by Remote Ring:0

35	Request packet	Will circle the Ring:0 (maybe getting satisfied along the way), and will be extracted on the next encounter with the RRC:0.

	ReadOnly Data Carrying packet	Will circle the Ring:0 (maybe getting satisfied along the way), and will be extracted on the next encounter with the RRC:0.
5	Ownership Data Carrying packet	Will circle the Ring:0 (maybe getting snarfed along the way), and will be extracted on the next encounter with the RRC:0.
10		
	<u>RRC:1. From Ring:1. Packets Originated by Local Ring:0</u>	
15	Request packet	Will circle the Ring:0 (maybe getting satisfied along the way), and will be extracted on the next encounter with the RRC:0.
20	ReadOnly Data Carrying packet	Will circle the Ring:0 (maybe getting satisfied along the way), and will be extracted on the next encounter with the RRC:0.
25	Ownership Data Carrying packet	Will circle the Ring:0 (maybe getting snarfed along the way), and will be extracted on the next encounter with the RRC:0.

RRC:1. From Ring:1. Packets Originated by Remote
Ring:0

5	Request packet	Will return to originator (maybe getting satisfied along the way) and get re-issued if necessary.
10	ReadOnly Data Carrying packet	Will return to originator (maybe getting satisfied along the way) and get reissued if necessary.
15	Ownership Data Carrying packet	Will return to originator to satisfy request. But if there was a pending request from the Ring:0 connected to this RRC, this packet should have been snarfed. When that packet returns to RRC:1, the RRC:1 will notice that ownership is still contained within the
20		Ring:0 and return the request back to Ring:0.

Ring Freezing

Freeze bit is used by a cell to freeze all
25 packets on the Ring for a single revolution. Freezing a Ring means that every packet on a Ring is in a state where no member on a Ring can change or copy any information in the packets. A cell freezes a Ring by systematically asserting Freeze bit in every
30 packet that passed by it until every packet on the Ring was frozen. During the freeze time, each cell is free to perform any operation without having to process to any Ring packets. Once each Ring packet has been frozen, the cell unfreezes each packet it

had frozen, starting at the first packet it froze,
until all packets it froze were unfrozen. If a
condition which initiates a Ring Freeze sequence is
detected during unfreezing, Ring Freeze sequence is
5 immediately started. The RRC must remember the
packets which it previously froze, that are being
refrozen. Figure 16 depicts a full ring freeze. The
lightly shaded region depicts packet slots not marked
as frozen, while the darkly shaded region depicts
10 those marked as frozen.

To handle multiple cells attempting to
freeze a Ring, each cell behaves almost as it would
if there were only one cell trying to freeze the
Ring. However, the cell must only unfreeze packets
15 which it marked frozen. Therefore each cell must
keep track of the packets it marked frozen.

A cell keeps track of which packets it
marked frozen by knowing the first and the last
packet it marked as frozen. This simplification is
20 possible since the freezing mechanism guarantees that
any given cell freezing sequence freezes one
contiguous string of packets. Thus the begin and end
freeze packets can be maintained using counters
rather than a shift register.

25 Ring freeze can be initiated by two methods,
Full Ring Freeze and Full-1 Ring Freeze. Since
cmd.Freeze is in the early portion of the packet
command, status leading to a Ring freeze must be
known early in packet processing. Full ring freeze
30 is when the Freeze bit is used for the first packet
to be frozen. When the status required for Ring
Freeze is not available early enough in packet
processing, Full-1 Ring Freeze method can be used.

Full:1 Ring Freeze method does not freeze the first packet to be frozen. Instead the packet busy bits, RspBsy and InvBsy (conveniently situated late in the packet), and SpState and Pendl subpage
5 state fields are used to cause the standard protocol to do the right thing (Fz:SingleEccErr state). The remaining packets are marked frozen and unfrozen in the same manner as Full Ring Freeze method.

The first packet is handled in a manner
10 identical to that described in Extract Buffer Full section for packets which cannot be extracted or copied into the extract buffer.

Refresh

15 Refresh is accomplished in a distributed manner by each RRC independently controlling it refresh when enabled by RiuMasterConfig.RefreshEnable. Each 8 msec., 512 refreshes must be performed. The interval between refreshes is 15.6 usec and is
20 specified in Ring clocks by RiuRefreshInterval location.

A single refresh requires one packet processing time. Thus refresh can only be performed during empty, or frozen packets, since these packet
25 require no processing. The RRC attempts perform refresh in the background by using all empty, frozen and invisible packet times for refresh. The RRC keeps track of the number of refreshes it has performed up to a limit of 31 by incrementing
30 RiuRefreshCount.NumRef for each refresh. At the end of each refresh interval NumRef is decremented by one. If NumRef was zero, then a refresh is immediately forced by the RRC using Full Ring Freeze method (NumRef is not incremented for the first

frozen packet). All RRCs can store up multiple refreshes equal to the number of formatted packets in the Ring. Full Ring Freeze refresh method is only required if the Ring was full utilized for extended
5 periods of time. Otherwise, refresh adds no additional overhead. The refresh interval must be adjusted .94 usec ($15.6 \times 31 / 512$) to 14.7 usec to account for the effect of saving up, up to 31 refreshes.

10

ECC Correctable Error

When the RRC detects a correctable ECC error, additional packet times are required to correct the data structure. Ring freezing is used to
15 allow multiple packet times for the correction and re-lookup.

When a correctable ECC error is detected, the RRC initiates Full-1 Ring Freeze method, since the ECC error is not known early in packet
20 processing. The first packet is handled as described in Extract Buffer Full section and the remaining packets are marked frozen for one revolution. The corrected data structure re-written, saved in RDU and then used to reprocess the packet which caused the ecc
25 error after one Ring revolution. Only RRC originated packets will be processed, non-originated packets are extracted by another cell within the Ring. Saving the corrected packet allows RRC originated packets to be correctly processed for single bit (correctable)
30 hard errors.

The RRC sets `pkt.RrcFault` in the packet corresponding to the correctable ECC error. The RRC also asserts `pkt.UnMarkedFault` on RRC originated packets when uncorrectable ECC errors are detected.

The requesting processor cell interprets the RrcFault bit in the following manner. The RrcFault bit indicates that an RRC fault was detected and logged by an RRC. If the remaining fault bits
5 (MarkedFault, LocalFault, DescFault, and UnMarkedFault) are clear then the packet request was successfully completed. If the returning packet was busied, the originating cell should signal a fault with RrcAtt status and not retry the request. Other
10 CIUs and CCUs handle this packet normally. If additional fault bits are set, the actions specified by these bits override the actions of RrcFault bit.

RDU:0 Rules

15 Previously defined RDU:0 Subpage and Page rules are referenced by these rules.

```
if pkt.frozen (pass,cond_refresh)
else if pkt_empty /* from insert buffer */
20     if pkt.nodata
        if !OriginateLocal
            if Fz:Freeze (pkt.freeze,-)
            else if Fz:SingleEccErr
                (insert.RspBsy.InvBsy.RrcFault,-)
25             else use Rules(Packets From
                Connected Ring:1,
                Packet.nodata,Originated by
                Remote Ring:0)
        else /* OriginateLocal */
30             if Fz:Freeze (pkt.freeze,-)
            else if Fz:SingleEccErr
                (insert.RspBsy.InvBsy.RrcFault,-)
            else use Rules(Packets From
25             Connected Ring:1,
            Packet.nodata,!Originated by
            Remote Ring:0)
35     else /* pkt.data */
        if Fz:None use Rules(Packets From
            Connected Ring:1, Packet.data,-)
40         else if Fz:Freeze (pkt.freeze,-)
            else /* Fz:SingleEccErr */
                if pkt.duplicate (empty,-)
```



```

                    else if pkt.sva_read_ro
                    (insert.nodata.RspBsy.InvBsy.RrcFa
                    ult,-)
                    else (insert,-); /* ownership
5          packets */
          else (pass,cond_refresh);
else/* !pkt_empty */
    if pkt.nodata
    if !OriginateLocal
10      if Fz:Freeze (pkt.freeze,-)
      else if Fz:SingleEccErr
      (pass.RspBsy.InvBsy.RrcFault,-)
      else use Rules(Packets From
15      Ring:0, Packet.nodata,Originated
      by Remote Ring:0)
      else /* OriginateLocal */
      if Fz:Freeze (pkt.freeze,-)
      else if Fz:SingleEccErr
      (insert.RspBsy.InvBsy.RrcFault,-)
20      else use Rules(Packets From
      Ring:0, Packet.nodata,!Originated
      by Remote Ring:0)
    else /* pkt.data */
    if Fz:None use Rules(Packets From
25      Ring:0,Packet.data,-)
    else if Fz:Freeze (pkt.freeze,-)
    else /* Fz:SingleEccErr */
    if pkt.duplicate (empty,-)
    else if pkt.sva_read_ro
30      (insert.nodata.RspBsy.InvBsy.RrcFa
      ult,-)
      else (insert,-); /* ownership
      packets */

35
RDU:1 Rules

Previously defined RDU:1 Subpage and Page rules are
referenced by these rules.
40
if pkt.frozen (pass,cond_refresh)
else if pkt_empty
    if pkt.nodata
    if OriginateLocal
45      if Fz:Freeze (pkt.freeze,-)
      else if Fz:SingleEccErr
      (insert.RspBsy.InvBsy.RrcFault,-)
```



```

else use Rules(Packets From
Connected Ring:0,
Packet.nodata,Originated by
Remote Ring:0)
5 else/* !OriginateLocal */
  if Fz:Freeze (pkt.freeze,-)
  else if
    Fz:SingleEccErr(insert.RspBsy.InvB
10 sy.RrcFault,-)
    else use Rules(Packets From
      Connected Ring:0,
      Packet.nodata,!Originated by
      Remote Ring:0)
  else /* pkt.data */
15   if Fz:None use Rules(Packets From
     Connected Ring:0, Packet.data,-)
   else if Fz:Freeze (pkt.freeze,-)
   else /* Fz:SingleEccErr */
20     if pkt.duplicate (empty,-)
     else if pkt.sva_read_ro
       (insert.nodata.RspBsy.InvBsy.RrcFa
       ult,-)
     else (insert,-); /* ownership
       packets */
25   else (pass,cond_refresh);
else /* pkt_empty */
  if pkt.nodata
  if OriginateLocal
    if Fz:Freeze (pkt.freeze,-)
30    else if Fz:SingleEccErr
      (pass.RspBsy.InvBsy.RrcFault,-)
    else use Rules(Packets From
      Ring:1, Packet.nodata,Originated
      by Remote Ring:0)
35    else/* !OriginateLocal */
      if Fz:Freeze (pkt.freeze,-)
      else if Fz:SingleEccErr
        (insert.RspBsy.InvBsy.RrcFault,-)
      else use Rules(Packets From
40      Ring:1, Packet.nodata,!Originated
      by Remote Ring:0)
    else /* pkt.data */
      if Fz:None use Rules(Packets From
        Ring:1,Packet.data,-)
45      else if Fz:Freeze (pkt.freeze,-)
      else /* Fz:SingleEccErr */
        if pkt.duplicate (empty,-)
        else if pkt.sva_read_ro
          (insert.nodata.RspBsy.InvBsy.RrcFa
50          ult,-)
```



```
else (insert,-); /* ownership  
packets */
```

5 Ring:0 Rules Assumptions

The following changes are required to packet specification, Ring specification, CCU rules and CIU rules for Ring:N operation.

10 Cell Descriptor Deallocation

Cell descriptor deallocation has changed.

Write Descriptor

As currently defined, the RRC must perform
15 two operations to implement the write descriptor operation. The RRC deallocates the descriptor at the old SVA and allocates the same descriptor at the new address. These two operations require two packet times with the deallocate address in the packet
20 address field and allocate address in packet dataword0. The new sequence described under the section "Change Descriptor" describes how the two RRC operations required for write descriptor (mpdw) are divided into two separate instructions and CCU
25 requests.

CiuUpdateState

Ciu_update_state is not required to be sent on Ring, assuming its only use is to change ciu
30 pending state to invalid.

Packet Operations

Three RRC related operations have been added, RrcDeallocDesc, RrcWrDesc and
35 CiuPureInvSumAlloc. The page_fault operation has

been deleted, since it is not required for Rrc operation. The following table specifies the revised encodings.

5	<u>Operation[4:0]</u>	<u>Name</u>	<u>Comment</u>
	00000	sva_read_ro	
	00001	duplicate	
	00010	release	
10	00011	recombine	
	00100	sva_read_ex	
	00101	write_invalidate	
	00110	get	
	00111	get_and_wait	
15	<u>Operation[4:0]</u>	<u>Name</u>	<u>Comment</u>
	01000	ciu_wr_descr	
	01001	ciu_update_state	
	01010	ciu_inv_sum	
20	01011	ciu_ex_sum	
	01100	reserved	page_fault deleted
	01101	ciu_id_drop_req	
	01110	spa_read_request	
25	01111	spa_write_request	
	10xxx	reserved	
	11000	RrcDeallocDesc	added
	11001	RrcWrDesc	added
	11010	Ciu_inv_sum_alloc	
30	11011	ciu_pure_inv_sum_alloc	added
	111xx	reserved	

Invalidator and Responder Busy

35 The definition of invalidator busy and responder busy as a function of each packet operation is specified in the table below. On certain operations these bits are also set to indicate that the a cell or Ring:0 has the page allocated

40 (descriptor match). This information is used by RRC:0 to determine whether a page is allocated in any additional cells within the Ring:0 during a cell deallocate sequence.

	<u>Operation</u>	<u>ResponderBusy</u>	<u>InvalidatorBusy</u>
5	sva_read_ro.no_data	responder busy or pending invalidate	set if descriptor match
10	sva_read_ro.data	responder busy or pending invalidate	
	release.data		
	release.no_data		
15	recombine.data		
20	duplicate.data	responder busy or pending invalidate	
	recombine.nodata	set if descriptor match	
25	sva_read_exc.nodata	responder busy	invalidator busy
30	sva_read_exc.data	responder busy	invalidator busy
	write_invalidate. nodata		invalidator busy
35	get.no_data	responder busy	invalidator busy
	get.data		invalidator busy
40	get_and_wait.nodata	responder busy	invalidator busy
45	get_and_wait.data		invalidator busy
	ciu_wr_descr.data	responder busy	
50	ciu_updata_state. nodata	responder busy	

	<u>Operation</u>	<u>ResponderBusy</u>	<u>InvalidatorBusy</u>
	ciu_inv_sum.nodata	responder busy	
5	ciu_ex_sum.nodata	responder busy	
	ciu_ld_drop_req. nodata		
10	ciu_inv_sum_allocate. nodata	set if descriptor match	invalidator busy
	RrcDeallocDesc	responder busy	
15	RrcWrDesc	responder busy	
	spa_read_request	responder busy	
20	spa_write_request	responder busy	

Packet Subpage State Field

The following table lists valid subpage
25 states as a function of packet operation.

	<u>Operation</u>	<u>Subpage States</u>
	sva_read_ro.no_data	
30	sva_read_ro.data	ro
	duplicate.data	ro
	release.data	ex
	release.no_data	
	recombine.data	nex,ex,at,tat
35	recombine.no_data	
	sva_read_ex.no_data	
	sva_read_ex.no_data	
	sva_read_ex.data	nex,ex,tex,at,tat
	write_invalidate.no_data	
40	get.no_data	
	get.data	
	get_and_wait.no_data	
	get_and_wait.data	
	ciu_wr_descr.data	
45	ciu_update_state.no_data	inv,ro,ex
	ciu_inv_sum.no_data	
	page_fault.data	
	ciu_ld_drop_req.no_data	

	<u>Operation</u>	<u>Subpage States</u>
	ciu_inv_sum_allocate.no_data	
	RrcDeallocDesc	
5	RrcWrDesc	
	spa_read_request	
	spa_write_request	
10	<u>Ring Freezing</u>	
	Ring freezing is used to handle ECC single bit errors and foreground refresh.	
	<u>Packet Pcopy</u>	
15	Assign pkt.cmd[33] as Pcopy bit for CIU, RIU and RRC usage. Always initialized to clear by inserting CIU. Always set by RRC:1 to indicate the packet was copied from Ring:1. Packets which are copied from Ring:1 with Pcopy set, are not required	
20	to be returned to RRC:1 or Ring:1 and are emptied by RRC:0 after successfully traversing Ring:0.	
	<u>Packet Freeze</u>	
	Assign pkt.cmd[32] as Freeze bit for CIU,	
25	RIU usage. Always initialized to clear by inserting CIU or RIU. When set indicates that the contents of the packet are ignored by all cells. Must be cleared by the same cell which sets it. When clear, the packet is interpreted normally. Valid for empty and	
30	non-empty packets.	
	<u>Packet Invalidate Lookahead</u>	
	Change packet.cmd[13] to reserved. Name packet.cmd[28] as cmd_lookahead[0] with the following	
35	encoding:	

	<u>packet.cmd[28]</u>	<u>Description</u>
	0	NoLookahead
5	1	Lookahead

Pkt.Lookahead always overrides data to indicate the packet is short.

10 Packet RRC Bit

Assign packet.cmd[8] as RRC bit. Always initialized to clear and passed by CCU and CIU. Used by sva_read_ro, recombines and release to indicate that the packet should be reissued if no cell within

15 Ring:0 picks up the subpage.

Packet RrcFault Bit

The definition of RrcFault bit has been changed RCC attention. The RrcFault bit indicates
20 that an RRC fault was detected and logged by an RRC. If the remaining fault bits (MarkedFault, LocalFault, DescFault, and UnMarkedFault) are clear then the packet request was successfully completed. If the returning packet was busied, the originating cell
25 should signal a fault with RrcAtt status and not retry the request. Other CIUs and CCUs handle this packet normally. If additional fault bits are set, the actions specified by these bits override the actions of RrcFault bit.

30

Packet RrcCmd

Rename pkt.cmd[23:16] to RrcCmd for RRC and RIU usage. RrcCmd is currently used to pass state between RRC:0 and RRC:1 for deallocate sequences.

35

	<u>packet.cmd[23:16]</u>	<u>Description</u>
	00000000	none
	00000001	cmd_pend_dealloc
5	00000010	cmd_dealloc_page
	00000011	reserved
	000001xx	reserved
	00001xxx	reserved
	0001xxxx	reserved
10	001xxxxx	reserved
	01xxxxxx	reserved
	1xxxxxxx	reserved

15 Cell Treatment of CEU Aborted Requests

Packets which are returned busy, corresponding to CEU aborted CCU requests should not be reissued by the CCU. System error conditions signalled to the local cache, which in turn, signal the local processor, causing an abort are guaranteed to get the local processors and local caches attention without hanging up.

Control Locations

25

Addressing

SPA space addresses are 64 bits, and are divided into two parts; a cell address and a cell address offset. The cell address specifies a particular cell in the hierarchy of Rings. The offset specifies a particular location within a cell.

Cell Address Structure

A cell address specifies the location of a cell within the three defined levels of the Ring interconnection hierarchy, Ring:2, Ring:1 and Ring:0. The leaf cell address is constructed hierarchically from Ring:2Address, Ring:1Address and

Ring:0Address. The three Ring:nAddress fields are used to construct absolute or relative hierarchical addresses to leaf or non-leaf cells.

Referring to the table below, there are
 5 three types of Ring:nAddresses; normal address, local address and routing address. A normal address is an absolute reference to one of 128 cells within the level of ring hierarchy. The cccccc field specifies the absolute cell number. A local address is a
 10 relative reference within level n-1 of ring the hierarchy. A routing address is used to address non-leaf cells within the hierarchy. The rrrrrr field within the routing address specifies the ring number or routing to Ring:n+1 for non-leaf cell
 15 access. Local address and Routing encodings are synonyms.

	<u>Ring:nAddr</u>	<u>Description</u>
20	0ccccccc lxxxxxxx lxrrrrrr	Normal Address Local Address Routing Ring:n+1
25	Cell addressing modes are categorized into leaf cell addressing and non-leaf (RRC) cell addressing. Cell addressing mode is determined by the combination of address types within the Cell Address.	
30	Leaf cell relative mode allows a cell to address itself without knowledge of its actual full Cell Address. Ring:2Addr, Ring:1Addr and Ring:0Addr specify local address. Leaf cell relative to local Ring:0 mode allows a cell to address other cells	
35	within its local Ring:0, without knowledge of Ring:1Addr or Ring:2Addr. Ring:0Addr field specifies	

the leaf cell addressed. Leaf cell relative to local
Ring:1 mode allows a cell to address other cells
within its local Ring:1, without knowledge of its
actual Ring:2Addr. Ring:1Address specifies Ring:0
5 number within the local Ring:2 and Ring:0Address
specifies the cell number within the Ring:0. Leaf
cell within a Ring:2 mode allows a specific leaf cell
within the system to be accessed with the full cell
address. Ring:2Address, Ring:1Address and
10 Ring:0Address fields hierarchically specify the
Ring:1, Ring:0 and cell number of the accessed leaf
cell.

RRC:1 (non-leaf) pair relative to local
Ring:1 mode, allows a non-local RRC:1 pair to be
15 addressed without knowledge of Ring:2Address.
Ring:2Addr type is local, Ring:1Addr type is normal
to specify a non-leaf cell within Ring:1 and
Ring:0Addr type is routing to specify routing. RRC:0
interleave configuration and Ring:0Addr field are
20 used to route the packet to the specified Ring:1.
RRC:1 pair mode allows a RRC:1 pair to be addressed
by specifying Ring:2Addr and Ring:1Addr. RRC:2 pair
mode allows an RRC:2 pair to be addressed by
specifying Ring:2Addr, Ring:1Addr routing and
25 Ring:0Addr routing.

	<u>Ring:2Addr</u>	<u>Ring:1Addr</u>	<u>Ring:0Addr</u>	<u>Cell Addressing Mode</u>
30	local	local	local	Leaf cell relative
	local	local	normal	Leaf cell relative
	local	normal	normal	to local Ring:0
35	normal	normal	normal	Leaf cell relative
	local	normal	normal	to local Ring:1
	local	normal	normal	to Ring:2
	local	normal	routing	RRC:1 pair relative
				to local Ring:1

	<u>Ring:2Addr</u>	<u>Ring:1Addr</u>	<u>Ring:0Addr</u>	<u>Cell Addressing Mode</u>
	normal	normal	routing	RRC:1 pair
	normal	routing	routing	RRC:2 pair
5	normal	routing	normal	reserved, detected by RRC:0

SPA[39] controls whether the addressed RRC or its
10 partner is accessed.

Routing of Interconnect Packets (SVA,SPA)

SVA request packets are routed based on RRC
directory lookup, packet command, match of
15 pkt.RequestorId and CellAddress, RRC rules and
interleave configuration.

SPA request packets are routed based on
packet command, match of pkt.Addr[63:40] and cell
address, match of pkt.RequestorId and CellAddress and
20 interleave configuration. If pkt.RequestorId and
CellAddress match, the packet is a returning response
and is forwarded back to the requesting leaf cell.
Otherwise if pkt.RequestorId and CellAddress match,
the packet is a request for the non-leaf or child
25 leaf cell depending on CellAddress mode.

Packet RequestorId Match and SPA address
decode are a function of the Ring level the RRC is
connected to.

30 SVA and SPA Packet RequestId Match

		<u>RRC Cell Address</u>		
	<u>RRC Ring</u>			
35	<u>Level</u>	<u>Ring:2Addr</u>	<u>Ring:1Addr</u>	<u>Ring:0Addr</u>
	0	=pkt.RequestId [23:16]	=pkt.RequestId [15:8]	=pkt.RequestId [7:0]
	1	=pkt.RequestId [23:16]	=pkt.RequestId [15:8]	not compared
40				

	<u>RRC Ring</u> <u>Level</u>	<u>Ring:2Addr</u>	<u>Ring:1Addr</u>	<u>Ring:0Addr</u>
5	2	=pkt.RequestId [23:16]	not compared	not compared

Packet SPA Address Decode

10

RRC Cell Address

	<u>RRC Ring</u> <u>Level</u>	<u>Ring:2Addr</u>	<u>Ring:1Addr</u>	<u>Ring:0Addr</u>
15	0	=pkt.Address [63:56]	=pkt.Address [55:48]	=pkt.Address [47:40]
	1	=pkt.Address [63:56]	=pkt.Address [55:48]	not compared
20	2	=pkt.Address [63:56]	not compared	not compared

Interleave configuration parameters are

25 RduMasterConfig.RingInterleave,
RduMasterConfig.RingValue, and
IbuMasterConfig.LinkInterleave.

SPA packet requests to leaf cells and
non-leaf cells are routed identically. SPA packet
30 requests are routed by an RRC bank based on
Interleave Value matching Ring:n Address field of
cell address, where n corresponds to the level of the
routing RRC. RingInterleave configuration determines
which bits from Ring:n Address[3:0] are matched. For
35 non-leaf cell access, RingAddr[3:0] specifies the
exact routing path. For leaf cell access, an exact
routing path between the requesting and accessed
Ring:0s is not required. The routing is based on the
cell addressed. Access to non-leaf Ring:2 devices
40 requires that both Ring:0 and Ring:1 fields of cell
address be specified.

	<u>RingInterleave</u> <u>Config</u>	<u>Interleave</u> <u>Value</u>	<u>SVA</u> <u>Interleave</u>	<u>SPA</u> <u>Interleave</u>
	One way	xxxx	--	--
5	Two way	xxxv	SVA[14]	RingAddr[0]
	Four way	xxvv	SVA[15:14]	RingAddr[1:0]
	Eight way	xvvv	SVA[16:14]	RingAddr[2:0]
	Sixteen way	vvvv	SVA[17:14]	RingAddr[3:0]

10

In configurations where both RRC banks connect to a single InterRingLink, incoming packets are routed to one bank based on IbuMasterConfig.LinkInterleave and a single address 15 bit.

	<u>Link</u> <u>Interleave</u>	<u>SVA</u> <u>Interleave</u>	<u>SPA Interleave</u>	
			<u>Leaf Cell</u>	<u>Non-Leaf Cell</u>
20	00	SVA[14]	SPA[14]	RingAddr[0]
	01	SVA[15]	SPA[15]	RingAddr[1]
	10	SVA[16]	SPA[16]	RingAddr[2]
	11	SVA[17]	SPA[17]	RingAddr[3]
25				

Offset Structure

The offset portion of an SPA selects a particular control register in the RRC pair, SPA[39] specifies whether the addressed RRC or its partner is accessed.

	<u>SPA[39]</u>	<u>Description</u>
35	0	addressed RRC
	1	partner of addressed RRC

When in each unit, control locations are 40 grouped by type and location offset. The range of

locations offsets for each type are specified with its individual description.

cell_address[38:32] - Unit Number

5 Bank A and bank B units are associated with subring A (odd ring number) and B (even ring number) respectively. Units are only accessible when accessed over the subring with which they are associated as specified by the Cell Address Ring:n
10 Addr field.

	<u>Unit</u>	<u>Name</u>
	00	RDUA (bank A)
15	02	RDUB (bank B)
	01	IBUA (bank A)
	03	IBUB (bank B)
	10	RIUA0 (bank A)
	11	RIUA1 (bank A)
20	12	RIUB2 (bank B)
	13	RIUB3 (bank B)
	04-0f	reserved
	14-7f	reserved

25 cell_address[31:28] - Location Type

The locations types for unit numbers 0x0 and 0x2, which address RDUs are:

	<u>Type</u>	<u>Name</u>
	0	RduMasterConfig
	1	L2FreeListPrtTable
	2	L3FreeListPrtTable
35	3	L1TableDirect
	4	L2PTableDirect
	5	L2VTableDirect
	6	L3PTableDirect
	7	L3VTableDirect
40	8	SvaRulesTable
	9-a	reserved
	b	L1TableLookup
	c	L2PTableLookup

	<u>Type</u>	<u>Name</u>
	d	L2VTableLookup
	e	L3PTableLookup
5	f	L3VTableLookup

The locations types for unit numbers 0x10 through 0x13, which address RIUs are:

10	<u>Type</u>	<u>Name</u>
	0	RiuMasterConfig
	1	RiuCellAddress
	2	LookupCommand
15	3	RiuRefreshInterval
	4	RiuRefreshCount
	5-f	reserved

20 The locations types for unit numbers 0x1 and 0x3, which address IBUs are:

	<u>Type</u>	<u>Name</u>
	0	IbuMasterConfig
25	1	IbuRefreshCount
	2	IbuRefreshInterval
	3	IRBPushPointer
	4	IRBPopPointer
	5	IRBCommitPointer
30	6	CommitPktNumber
	7	IbuXmitSeqNum
	8	IbuRcvSeqNum
	9	EQPushPointer
	10	EQPopPointer
35	11	EQPopData
	12	RPQPopData
	13-15	reserved

40 Summary

Those skilled in the art will appreciate that the embodiments described above are exemplary only, and that other apparatuses and methods -- including modifications, additions and deletions --
45 fall within the scope and spirit of the invention.

By way of example, it will be appreciated that the central processing unit and subcache included within the processing cells of the illustrated embodiment may be replaced by other logic
5 circuitry capable of operating on, or in conjunction with, data allocated to those cells.

In view of the foregoing, what we claim is:

-207-

1. A switch for selectively routing digital information packets received from at least first and second external sources to at least first and second external destinations,

5

at least said first external destination having associated therewith zero, one or more data or copies thereof, each said data or copy thereof corresponding to a unique descriptor by which that
10 datum or copy is referenced,

at least one of said first and second external sources including means for generating information packets comprising one or more signals
15 representative of a datum or copy thereof or request therefor and further comprising at least a signal representative of a descriptor,

said routing switch comprising:

20

A. first and second routing interconnect elements, each including

i) a first input means for receiving
25 an information packet from a respective one of said first and second external sources,

ii) a first output means for
transmitting an information packet to a
30 respective one of said first and second external destinations,

iii) second input means for receiving an information packet from the other of said routing interconnect elements, and

5 iv) second output means for transmitting an information packet to the other of said routing interconnect element,

B. each of said routing interconnect
10 elements including directory means having allocated therein zero, one or more entry signals respectively representative of descriptors corresponding to at least one of

15 i) said zero, one or more data or copies thereof associated with said first external destination,

20 ii) zero, one or more data or copies thereof requested in information packets received from by the second input means of that routing interconnect element and routed for transmission by the first output means of that routing interconnect element, and

25 C. each of said routing interconnect elements further including control means for selectively transferring an information packet received by at least one of the first and second
30 input means of the associated routing interconnect element to at least one of the first and second output means of that routing interconnect element

based on correspondence, or lack thereof, between the descriptor in that information packet and an entry signal allocated to that directory means.

5

2. A switch according to claim 1, wherein said control means of the first routing interconnect element includes at least one of

10

A. means responsive to receipt, at the second input means of that routing interconnect element, of a packet including a request for a datum or copy thereof referenced by a descriptor that corresponds to an entry signal allocated to the
15 directory means of that routing interconnect element for a datum or copy thereof associated with said first external destination for routing that packet from that second input means to the first output means of that routing interconnect element, and

20

B. means responsive to receipt, at the second input means of that routing interconnect element, of a packet including a request for a datum or copy thereof referenced by a descriptor that does
25 not correspond to an valid entry signal allocated to the directory means of that routing interconnect element for a datum or copy thereof associated with said first external destination for routing that packet from that second input means to the second
30 output means of that routing interconnect element.

3. A switch according to claim 2, wherein said control means includes means selectively responsive

to transfer of packet including a datum or copy thereof from at least one of

5 i) the first input means to the second
output means of the associated routing
interconnect element, and

10 ii) the second input means to the first
output means of the associated routing
interconnect element

for at least one of allocating, deallocating or
modifying an entry signal in the directory means of
that element.
15

4. A switch according to claim 3, wherein said
control means includes means selectively responsive
to transfer of a packet including a request for a
20 datum or copy thereof from the second input means to
the first output means of the associated routing
interconnect element for at least one of allocating,
deallocating and modifying an entry signal in the
directory means of that element.
25

5. A switch according to claim 4, wherein said
first and second routing interconnect elements are
disposed physically remotely from one another.
30

6. A switch according to claim 5, wherein at
least one information packet generated by said first
and second external sources includes a requester-ID

-211-

signal associated with an external destination to which that packet is directed, at least one of said routing interconnect elements comprising at least one of

5

A. means responsive to receipt on any of the first and second input means of that routing interconnect element of at least a selected packet including a requester-ID signal associated with the
10 external destination to which the first output means transmits information packets for routing that packet from that first input means to that first output means, and

15

B. means responsive to receipt on any of the first and second input means of that routing interconnect element of at least a selected packet including a requester-ID signal that is not associated with the external destination to which the
20 first output means transmits information packets for routing that packet from that first input means to the second output means of that routing interconnect element.

25

7. A switch for selectively routing digital information packets received from at least first and second external sources to at least first and second external destinations,

30

at least said first external destination having associated therewith zero, one or more data or copies thereof, each said data or copy thereof

corresponding to a unique descriptor by which that datum or copy is referenced,

at least one of said first and second
5 external sources including means for generating information packets comprising one or more signals representative of a datum or copy thereof or request therefor and further comprising at least a signal representative of a descriptor,

10

said routing switch comprising:

A. first and second routing interconnect elements, each including

15

i) a first input means for receiving an information packet from a respective one of said first and second external sources,

20

ii) a first output means for transmitting an information packet to a respective one of said first and second external destinations,

25

iii) second input means for receiving an information packet from the other of said routing interconnect elements, and

30

iv) second output means for transmitting an information packet to the other of said routing interconnect element,

B. each of said routing interconnect elements including directory means having allocated

therein zero, one or more entry signals respectively representative of descriptors corresponding to at least one of

5 i) said zero, one or more data or copies thereof associated with said first external destination,

10 ii) zero, one or more data or copies thereof requested in information packets received from by the second input means of that routing interconnect element and routed for transmission by the first output means of that routing interconnect element,

15 C. each of said routing interconnect elements further including control means for selectively transferring an information packet received by at least one of the first and second
20 input means of the associated routing interconnect element to at least one of the first and second output means of that routing interconnect element based on correspondence, or lack thereof, between the descriptor in that information packet and an entry
25 signal allocated to that directory means, and

 D. said control means of the second routing interconnect element includes at least one of

30 i) means responsive to receipt, at the second input means of that routing interconnect element, of a packet including a request for a datum or copy thereof referenced by a descriptor that does not

- 5 correspond to an entry signal allocated to the directory means of that routing interconnect element for a datum or copy thereof associated with said first external destination for routing that packet from that second input means to the first output means of that routing interconnect element, and
- 10 ii) means responsive to receipt, at the second input means of that routing interconnect element, of a packet including a request for a datum or copy thereof
- 15 referenced by a descriptor that corresponds to a valid entry signal allocated to the directory means of that routing interconnect element for a datum or copy thereof associated with said first external destination for routing that packet from
- 20 that second input means to the second output means of that routing interconnect element.
8. A switch according to claim 7, wherein said
- 25 control means includes means selectively responsive to transfer of packet including a datum or copy thereof from at least one of
- 30 i) the first input means to the second output means of the associated routing interconnect element, and

ii) the second input means to the first output means of the associated routing interconnect element

5 for at least one of allocating, deallocating or modifying an entry signal in the directory means of that element.

10 9. A switch according to claim 8, wherein said control means includes means selectively responsive to transfer of a packet including a request for a datum or copy thereof from the second input means to the first output means of the associated routing
15 interconnect element for at least one of allocating, deallocating and modifying an entry signal in the directory means of that element.

20 10. A switch according to claim 9, wherein said first and second routing interconnect elements are disposed physically remotely from one another.

25 11. A switch according to claim 10, wherein at least one information packet generated by said first and second external sources includes a requester-ID signal associated with an external destination to which that packet is directed, at least one of said
30 routing interconnect elements comprising at least one of

:

A. means responsive to receipt on any of the first and second input means of that routing

-216-

interconnect element of at least a selected packet including a requester-ID signal associated with the external destination to which the first output means transmits information packets for routing that packet
5 from that first input means to that first output means, and

B. means responsive to receipt on any of the first and second input means of that routing
10 interconnect element of at least a selected packet including a requester-ID signal that is not associated with the external destination to which the first output means transmits information packets for routing that packet from that first input means to
15 the second output means of that routing interconnect element.

12. A switch according to any one of claims 1-11,
20 wherein said control means of the first routing interconnect element includes at least one of

A. means responsive to receipt, at the first input means of that routing interconnect element, of
25 a packet including a request for a datum or copy thereof referenced by a descriptor that corresponds to an entry signal allocated to the directory means of that routing interconnect element for a datum or copy thereof associated with said first external
30 destination for routing that packet from that first input means to the first output means of that routing interconnect element, and

-217-

B. means responsive to receipt, at the first input means of that routing interconnect element, of a packet including a request for a datum or copy thereof referenced by a descriptor that does not
5 correspond to an entry signal allocated to the directory means of that routing interconnect element for a datum or copy thereof associated with said first external destination for routing that packet from that first input means to the second output
10 means of that routing interconnect element.

13. A switch according to any one of claims 1-11, wherein said control means of the second routing
15 interconnect element includes at least one of

A. means responsive to receipt, at the first input means of that routing interconnect element, of a packet including a request for a datum or copy
20 thereof referenced by a descriptor that does not correspond to an entry signal allocated to the directory means of that routing interconnect element for a datum or copy thereof associated with said first external destination for routing that packet
25 from that first input means to the first output means of that routing interconnect element, and

B. means responsive to receipt, at the first input means of that routing interconnect element, of
30 a packet including a request for a datum or copy thereof referenced by a descriptor that corresponds to a entry signal allocated to the directory means of that routing interconnect element for a datum or copy thereof associated with said first external

-218-

destination for routing that packet from that first input means to the second output means of that routing interconnect element.

5

14. A switch according to any one of claims 5, 6, 10, or 11, wherein at least one of said first and second routing interconnect elements includes remote interface means for at least one of

10

A. transmitting information packets from the second output means of that routing interconnect element to the remotely disposed second input means of the other of said routing interconnect elements,
15 and

B. receiving information packets at the second input means of that routing interconnect element transmitted from the remotely disposed second
20 output means of the other of said routing interconnect elements.

15. A switch according to claim 14, wherein said
25 remote interface means includes fiber optic transmission media for carrying signals representative of said information packets between said at least one remotely disposed second input means and second output means of said routing
30 interconnect elements.

16. A switch according to any one of claims 1-11, wherein at least one of said routing interconnect

elements includes means selectively responsive to receipt at said first input means of a packet including a response to a request packet previously routed from the second input means to the first
5 output means of that routing interconnect element for routing that packet from that first input means to the second output means of that routing interconnect element.

10

17. A switch according to claim 16, wherein at least said second routing interconnect element includes means selectively responsive to receipt at said first input means of a packet including a datum
15 or copy thereof for which at least one of

i) a read-only request has previously been transferred from the second input means of that routing interconnect element to the
20 first output means of that routing interconnect element, and

ii) an invalid copy is associated to said first external destination,
25

for generating an information packet including a copy of that datum or copy thereof and for transferring that second duplicate data packet to said second output means.

30

18. A switch according to claim 16, wherein at least said second routing interconnect element includes means selectively responsive to receipt at

-220-

said first input means of a packet including a datum or copy thereof for which any of an ownership request and an atomic request packet has been transferred from the second input means to the first output means of that routing interconnect element for transferring that packet received at said first input means to the second output means of that routing interconnect element.

10

19. A switch according to any one of claims 1-11, wherein the directory means of at least one of said routing interconnect elements includes means for storing association state-defining information for at least selected data, or copies thereof, associated with said first external destination, said association state-defining information including at least one of

20

i) an invalid association state for an invalid datum associated with said first external destination,

25

ii) a read-only association state for a copy of a datum associated with and subject to read-only access within said first external destination, no copy of which datum is modifiable within said first external destination,

30

iii) a nonexclusive owner association state for a datum associated with and modifiable within said first external destination, one or more read-only copies of

-221-

which datum are associated with at least one of said first and second external destinations, and

5 iv) an exclusive owner association state for a datum exclusively associated with and modifiable within said first external destination, no read-only copies of which datum are associated with either of
10 said first and second external destinations.

20. A switch according to any one of claims 1-11, wherein the directory means of at least one of said
15 routing interconnect elements includes means for storing pending state-defining information for at least selected data for which requests are pending to the associated external destination, said pending state-defining information including at least one of
20

 i) a none-pending pending state for a datum for which no requests are pending,

 ii) a read-only pending state for a
25 datum for which a read-only association request is pending,

 iii) an ownership pending state for a
30 datum for which an ownership association request is pending, or for which an invalidation request is pending, and

 iv) a reissue pending state for a datum for which a request is pending and for which

-222-

a response other than one generated in
response to that request has issued, which
response was not routed from the first input
means to the second output means of that
5 routing interconnect element.

21. A switch according to any one of claims 1-11,
wherein at least one of said routing interconnect
10 elements includes multiple request-handling means
selectively responsive to receipt, at the second
input means of that routing interconnect element, of
multiple information packets requesting association
of a datum for transferring an initial one of those
15 request packets to the first output means of that
routing interconnect element and for modifying later
received ones of such request packets to indicate
that a responder is busy and transferring those
modified packets to the second output means of that
20 routing interconnect element.

22. A switch according to claim 21, wherein said
multiple request-handling means includes at least one
25 of

A. means selectively responsive to
receipt, at the second input means of the associated
routing interconnect element, of multiple information
30 packets requesting read-only allocation of a datum
for transferring an initial one of those request
packets to the first output means of that associated
routing interconnect element and for modifying later
received ones of such request packets to indicate

that a responder is busy and transferring those modified packets to the second output means of that associated routing interconnect element, and

5 B. means selectively responsive to receipt, at the second input means of the associated routing interconnect element, of one or more information packets requesting ownership allocation or invalidation of a datum for transferring those
10 packets to the to the first output means of that associated routing interconnect element.

23. An apparatus according to claim 21, wherein
15 said multiple request-handling means includes at least one of

 A. means selectively responsive to receipt, at the second input means of the associated
20 routing interconnect element, of multiple information packets requesting read-only allocation of a datum for transferring an initial one of those request packets to the first output means of that associated routing interconnect element and for modifying later
25 received ones of such request packets to indicate that a responder is busy and transferring those modified packets to the second output means of that associated routing interconnect element,

30 B. means selectively responsive to receipt, at the second input means of the associated routing interconnect element, of multiple information packets requesting exclusive ownership of a datum for transferring an initial one of those request packets

to the first output means of that associated routing
interconnect element and for modifying later received
ones of such request packets to indicate that a
responder is busy and transferring those modified
5 packets to the second output means of that associated
routing interconnect element,

C. means selectively responsive to
receipt, at the second input means of the associated
10 routing interconnect element, of multiple information
packets requesting, without wait, atomic ownership of
a datum for transferring an initial one of those
request packets to the first output means of that
associated routing interconnect element and for
15 modifying later received ones of such request packets
to indicate that the requested datum is atomically
owned and transferring those modified packets to the
second output means of that associated routing
interconnect element,

20

D. means selectively responsive to
receipt, at the second input means of the associated
routing interconnect element, of multiple information
packets requesting, with wait, atomic ownership of a
25 datum for transferring an initial one of those
request packets to the first output means of that
associated routing interconnect element and for
modifying subsequently received ones of such request
packets, as well as to subsequently received packets
30 requesting atomic ownership without wait, to indicate
that the requested datum is atomically owned and
transferring those modified packets to the second
output means of that associated routing interconnect
element, and

35

-225-

E. reissue means selectively responsive to receipt, at the second input means of the associated routing interconnect element, of an unresolved information packet requesting a datum, which
5 information packet was previously issued by said multiple request-handling means, for modifying that packet to indicate that a responder is busy and transferring that modified packets to the second output means of that associated routing interconnect
10 element.

24. A switch according to claim 21,
wherein the directory means of at least one of said
15 routing interconnect elements includes means for storing pending state-defining information for at least selected data for which requests are pending to the associated external destination, said pending state-defining information including at least one of
20

i) a none-pending pending state for a datum for which no requests are pending,

25 ii) a read-only pending state for a datum for which a read-only association request is pending,

30 iii) an exclusive ownership pending state for a datum for which an exclusive ownership association request is pending,

iv) an ownership invalidate pending state for a datum for which an invalidate request is pending;

35

-226-

v) an GET pending state for a datum for which an atomic ownership request, without wait, is pending,

5 vi) a GETW pending state for a datum for which an atomic ownership request, with wait, is pending, and

10 vii) an atomic pending state for a datum for which an atomic ownership request, with wait, is pending, which data is previously associated with said atomic state in one of said first and second external destinations, and

15 iv) a reissue pending state for a datum for which a request is pending and for which a response other than one generated in response to that request has issued, which
20 response was not routed from the first input means to the second output means of that routing interconnect element.

25 25. A switch according to any one of claims 1-11, wherein at least one of said first and second routing interconnect elements include

A. multiple request-handling means
30 selectively responsive to receipt, at the second input means of the associated routing interconnect element, of multiple information packets requesting allocation of a datum for transferring an initial one of those request packets to the first output means of

that associated routing interconnect elements and for deleting later issued ones of said request packets, and

5 B. said multiple request-handling means including merge means for associating a merged-state signal with a datum for which such later issued ones of said requests have been deleted.

10

26. A switch according to claim 25, wherein said multiple request means includes means responsive to a packet including a datum associated with said merged-state signals for modifying that packet to
15 include a merged-response signal indicating that a requestor to which that response packet is directed should not delete that packet.

20 27. A switch according to claim 25, wherein

 A. multiple request-handling means selectively responsive to receipt, at the second input means of the associated routing interconnect
25 element, of multiple information packets requesting like allocation of a datum for transferring an initial one of those request packets to the first output means of that associated routing interconnect elements and for deleting later issued ones of said
30 like request packets, and

 B. said multiple request-handling means including merge means for associating a merged-state

signal with a datum for which such later issued ones of such like requests have been deleted.

5 28. A switch according to claim 25, wherein the directory means of at least one of said routing interconnect elements includes means for storing pending state-defining information for at least selected data for which requests are pending to the
10 associated external destination, said pending state-defining information including at least one of

15 i) a none-pending pending state for a datum for which no requests are pending,

ii) a wait pending state for a datum for which a request packet is pending and for which a response has previously issued, which request packet is to be emptied upon
20 receipt by the routing interconnect cell,

iii) a wait-merge-read-only pending state for a datum for which multiple information packets requesting read-only
25 association have been received and for which a single read-only request packet is pending, which request packet is to be emptied upon receipt by the routing interconnect element,

30 iv) a wait-merge-exclusive pending state for a datum for which multiple information packets requesting exclusive ownership association have been received and

for which a single exclusive ownership request packet is pending, which request packet is to be emptied upon receipt by the routing interconnect element,

5

v) a wait-merge-atomic pending state for a datum for which multiple information packets requesting atomic ownership association have been received and for which a single atomic ownership request packet is pending, which request packet is to be emptied upon receipt by the routing interconnect element,

10

vi) a reissue pending state for a datum for which a request is pending, which request is to be reissued upon receipt by the routing interconnect element,

15

vii) a reissue-merge-read-only pending state for a datum for which multiple information packets requesting read-only association have been received by the routing interconnect element, and for which a single read-only request is pending, which request is to be reissued upon receipt by the routing interconnect element,

20

25

viii) a reissue-merge-exclusive pending state for a datum for which multiple information packets requesting exclusive ownership association have been received by the routing interconnect element, and for which a single exclusive ownership request

30

is pending, which request is to be reissued upon receipt by the routing interconnect element,

5 ix) a reissue-merge-atomic pending
state for a datum for which multiple
information packets requesting atomic
ownership association have been received by
the routing interconnect element, and for
10 which a single atomic ownership request is
pending, which request is to be reissued
upon receipt by the routing interconnect
element,

15 x) a read-only pending state for a
datum for which a read-only association
request is pending,

20 xi) a read-only-merge pending state for
a datum for which multiple information
packets requesting read-only association
have been received and for which a single
read-only request packet is pending, which
request packet is to be marked as merged
25 upon receipt by the routing interconnect
element,

30 xii) an exclusive ownership pending
state for a datum for which an exclusive
ownership association request is pending,

 xiii) an exclusive-merge pending state
for a datum for which an exclusive ownership
request packet was pending upon receipt by

-231-

the routing interconnect unit of a further
request for association with that datum,
which request packet is to be marked as
merged upon receipt by the routing
5 interconnect element,

xiv) an atomic pending state for which
an atomic ownership request is pending,

10 xv) an atomic-merge pending state for a
datum for which an atomic ownership request
packet was pending upon receipt by the
routing interconnect unit of a further
request for association with that datum,
15 which request packet is to be marked as
merged upon receipt by the routing
interconnect element,

20 xvi) a recombine pending state for a
packet including that datum, routed for
recombine association with an external
destination, is pending, and

25 xvii) a merged response pending state
for a datum for which a response packet
marked as merged is pending and for which a
further request has been received by the
routing interconnect element.

30 29. A switch according to any one of claims 1-11,
wherein the control means of at least one of said
routing interconnect elements includes routing
directory unit means connected with the directory

-232-

means of that routing interconnect element for
generating an association-based routing signal
indicative of a packet routing based on
correspondence, or lack thereof, between the
5 descriptor in an information packet received at any
of said first and second input means of that routing
interconnect element and an entry signal allocated to
that directory means.

10

30. A switch according to any one of claims 1-11,
wherein

A. the first input means and the first
15 output means of said first routing interconnect
element are arranged together to form a first
terminal pair, and

B. said first terminal pair comprises a
20 bus connecting said first external source to said
first external destination.

31. A switch according to any one of claims 1-11,
25 wherein

A. the first input means and the first
output means of said second routing interconnect
element are arranged together to form a second
30 terminal pair, and

B. said second terminal pair comprises a
bus connecting said second external source to said
second external destination.

35

-233-

32. A switch according to claim 31, wherein at least one of said first and second terminal pairs comprise a plurality of shift register stages each for sequentially and successively transferring at
5 least a portion of an information packet.

33. A switch according to any one of claims 1-11, wherein at least one of said interconnect elements
10 comprise

A. extract buffer means connected to said first input means and to said second output means of that routing interconnect element for buffering
15 information packets transferred therebetween, and

B. insert buffer means connected to said second input means and to said first output means of that routing interconnect element for buffering
20 information packets transferred therebetween.

34. A switch according to claim 33, wherein

25 A. the control means of at least one of said routing interconnect elements includes routing command means connected to said routing directory unit means and responsive to said association-based routing signal for selectively generating a push
30 signal indicating that an information packet received at the first input means of that routing interconnect element is to be copied to the extract buffer means of that routing interconnect element, and

B. said first input means of at least one of said routing interconnect elements includes means connected to said routing command means of that routing interconnect element for responding to said
5 push signal for generating a signal representing a copy of an information packet received at said first input means of that routing interconnect element and for transferring that signal to the extract buffer of that routing interconnect element.

10

35. A switch according to claim 33, wherein

A. the control means of at least one of
15 said interconnect elements includes routing command means connected to said routing directory unit means and responsive to said association-based routing signal for selectively generating a pop signal indicating that an information packet buffered in
20 said insert buffer means of that routing connect element is to be transferred to said first output means of that routing connect element, and

B. said insert buffer of at least one of
25 said routing interconnect elements includes means connected to said routing command means of that routing connect element for responding to said pop signal for transferring an information packet buffered therein to said first output means of that
30 routing connect element.

36. A switch according to claim 33, wherein

A. the control means of at least one of
said routing interconnect elements includes routing
5 command means connected to said routing directory
unit means of that routing connect element and
responsive to said association-based routing signal
for selectively generating a pass signal indicating
that an information packet received at said first
10 input means of that routing connect element is to be
passed to said first output means of that routing
connect element,

B. said first input means of at least one
15 of said routing interconnect elements includes means
connected to said routing command means of that
routing connect element for responding to said pass
signal for transferring an information packet
received at said first input means of that routing
20 connect element to said first output means of that
routing connect element.

37. A switch according to claim 33, wherein
25

A. the control means of at least one of
said routing interconnect elements includes routing
command means connected to said routing directory
unit means of that routing connect element and
30 responsive to said associated-based routing signals
for selectively generating an empty signal indicating
that an information packet received at said first
input means of that routing connect element is to be
emptied, and

35

B. said first input means of at least one of said routing interconnect elements includes means connected to said routing command means of that routing connect element for responding to said empty
5 signal for modifying an information packet received at said first input means of that routing connect element to indicate that it is empty.

10 38. A switch according to claim 33, wherein

A. the control means of at least one of said routing interconnect elements includes extract
15 buffer management means connected to said extract buffer means of that routing connect element for generating a buffer-region signal representative of a number of packets buffered therein,

B. said routing directory unit means of
20 that control means includes means for selectively generating an extract signal indicating that an information packet received at said first input of that routing connect element is to be transferred to said extract buffer means of that routing connect
25 element, and

C. said control means further includes routing command means connected to said routing
directory unit means of that routing connect element
30 and to said extract buffer management means of that routing connect element for generating a routing signal indicative of a packet routing based upon said extract signal and said buffer-region signal.

39. A switch according to claim 38, wherein said routing command means of at least one of said routing interconnect elements includes extract-all means responsive to assertion of said extract signal and to
5 a buffer-region signal representative of a number less than a first threshold value for transferring the corresponding packet to said extract buffer means of that routing connect element.

10

40. A switch according to claim 38, wherein said routing command means of at least one of said routing interconnect elements includes extract-noptional means responsive to assertion of said extract signal
15 and to a buffer-region signal representative of a number greater than said first threshold value and less than a second, higher threshold value for selectively transferring the corresponding packet to said extract buffer means of that routing connect
20 element or, alternatively, transferring that packet to said first output means of that routing connect element.

25 41. A switch according to claim 40, wherein said extract-noptional means of at least one of said routing interconnect elements includes means for selecting for transfer to said first output means of that routing connect element an information packet
30 including any of

i) a read-only copy of a datum received at the first input means of that routing connect element of the second

routing interconnect element, no read-only copy of which datum is associated with said first external destination; and

5 ii) a datum subject to ownership access received at the first input means of the second routing interconnect element, no copy of which data subject to ownership access is associated with said first
10 external destination.

42. A switch according to claim 38, wherein said routing command means of at least one of said routing
15 interconnect elements includes extract-pending means responsive to assertion of said extract signal and to a buffer-region signal representative of a number greater than said second threshold value and less than a third, highest threshold for selectively
20 transferring the corresponding packet to said extract buffer means of that routing connect element or, alternatively, transferring that packet to said first output means of that routing connect element.

25
43. A switch according to claim 38, wherein said extract-pending means of at least one of said routing interconnect elements includes means for selecting for transfer to said first output means of that
30 routing connect element an information packet including any of

:

 i) a read-only copy of a datum received at the first input means of the

-239-

second routing interconnect element, no read-only copy of which datum is associated with said first external destination,

5 ii) a datum subject to ownership access received at the first input means of the second routing interconnect element, no copy of which data subject to ownership access is associated with said first
10 external destination, and

 iii) a response packet received at the first input means of the first routing interconnect element to a request packet
15 previously received at the first input means of the second routing interconnect element.

44. A switch according to claim 38, wherein said
20 routing command means of at least one of said routing interconnect elements includes extract-full means responsive to a buffer-region signal representative of a number greater than said third threshold for preventing the transfer of an information packet to
25 said extract buffer means of that routing connect element.

45. A switch according to any one of claims 1-11,
30 wherein the second output of at least one of said routing interconnect elements includes outgoing buffer means for buffering information packets transferred from the corresponding extract buffer

-240-

means to the second output means of that routing connect element.

5 46. A switch according to claim 45, wherein the second output of said at least one interconnect elements includes ring packet queue means connected in parallel between the extract buffer means of that routing connect element and outgoing buffer means of
10 that routing connect element for at least one of

i) buffering overflow packets transferred therebetween,

15 ii) buffering packets for which acknowledgement of transfer is pending transferred between the corresponding outgoing buffer means and the second input means of the other interconnect element.

20

47. A switch according to claim 45, wherein said ring packet queue means comprises a circular queue.

25

48. A switch according to any one of claims 1-11, wherein the second output means of at least one of said routing interconnect elements includes error queue means for storing one or more packets in which
30 an error is detected.

49. A switch according to claim 48, wherein said error queue means includes means for storing, along

-241-

with each packet in which an error is detected, a signal representative of an error associated with that packet.

5

50. A switch according to any one of claims 1-11, wherein the second input means of at least one of said routing interconnect element includes incoming buffer means for buffering packets received at the
10 second input means of that routing connect element.

51. A method for routing digital information packets received from at least first and second
15 external sources to at least first and second external destinations,

at least said first external destination having associated therewith zero, one or more data or
20 copies thereof, each said datum or copy thereof corresponding to a unique descriptor by which that datum or copy is referenced,

at least one of said first and second
25 external sources including means for generating information packets comprising one or more signals representative of a datum or copy thereof or request therefor and further comprising at least a signal representative of a descriptor,

30

said method comprising:

A. providing first and second routing interconnect elements, each including

35

i) a first input for receiving an information packet from a respective one of said first and second external sources,

5 ii) a first output for transmitting an information packet to a respective one of said first and second external destinations,

10 iii) second input for receiving an information packet from the other of said routing interconnect elements, and

15 iv) second output for transmitting an information packet to the other of said routing interconnect element,

B. allocating in a directory of each of said routing interconnect elements zero, one or more entry signals respectively representative of
20 descriptors corresponding to at least one of

25 i) said zero, one or more data or copies thereof associated with said first external destination,

30 ii) zero, one or more data or copies thereof requested in information packets received from by the second input of that routing interconnect element and routed for transmission by the first output of that routing interconnect element, and

C. selectively transferring an information packet received by at least one of the first and

second input of each of said routing interconnect elements to at least one of the first and second output of that routing interconnect element based on correspondence, or lack thereof, between the
5 descriptor in that information packet and an entry signal allocated to that directory.

52. A method according to claim 51, comprising
10

A. responding within the first interconnect element to receipt, at the second input of that routing interconnect element, of a packet including a request for a datum or copy thereof referenced by a
15 descriptor that corresponds to an entry signal allocated to the directory of that routing interconnect element for a datum or copy thereof associated with said first external destination for routing that packet from that second input to the
20 first output of that routing interconnect element, and

B. responding within the first interconnect element to receipt, at the second input of that routing interconnect element, of a packet including a
25 request for a datum or copy thereof referenced by a descriptor that does not correspond to an valid entry signal allocated to the directory of that routing interconnect element for a datum or copy thereof associated with said first external destination for
30 routing that packet from that second input to the second output of that routing interconnect element.

:

53. A method according to claim 52, comprising selectively responding within at least one of said first and second routing interconnect elements to at least one of

5

i) the transfer of a packet including a datum or copy thereof from the first input to the second output of the associated routing interconnect element, and

10

ii) the transfer of a packet including a datum or copy thereof from the second input to the first output of the associated routing interconnect element

15

for at least one of allocating, deallocating or modifying an entry signal in the directory of that routing interconnect element.

20

54. A method according to claim 53, comprising selectively responding within at least one of said first and second routing interconnect elements to transfer of a packet including a request for a datum
25 or copy thereof from the second input to the first output of that routing interconnect element for at least one of allocating, deallocating and modifying an entry signal in the directory of that element.

30

55. A method according to claim 54, comprising disposing said first and second routing interconnect elements physically remotely from one another.

35

56. A method according to claim 55, wherein at least one information packet generated by said first and second external sources includes a requester-ID signal associated with an external destination to which that packet is directed, said method comprising

A. responding within at least one of said first and second routing interconnect elements to receipt on any of the first and second inputs of that routing interconnect element of at least a selected packet including a requester-ID signal associated with the external destination to which the first output transmits information packets for routing that packet from that first input to that first output, and

15

B. responding within at least one of said first and second routing interconnect elements to receipt on any of the first and second inputs of that routing interconnect element of at least a selected packet including a requester-ID signal that is not associated with the external destination to which the first output transmits information packets for routing that packet from that first input to the second output of that routing interconnect element.

25

57. A method for routing digital information packets received from at least first and second external sources to at least first and second external destinations,

at least said first external destination having associated therewith zero, one or more data or copies thereof, each said datum or copy thereof

corresponding to a unique descriptor by which that datum or copy is referenced,

at least one of said first and second
5 external sources including means for generating information packets comprising one or more signals representative of a datum or copy thereof or request therefor and further comprising at least a signal representative of a descriptor,

10

said method comprising:

A. providing first and second routing
interconnect elements, each including

15

i) a first input for receiving an information packet from a respective one of said first and second external sources,

20

ii) a first output for transmitting an information packet to a respective one of said first and second external destinations,

25

iii) second input for receiving an information packet from the other of said routing interconnect elements, and

30

iv) second output for transmitting an information packet to the other of said routing interconnect element,

B. allocating in a directory of each of said routing interconnect elements zero, one or more

entry signals respectively representative of
descriptors corresponding to at least one of

5 i) said zero, one or more data or
 copies thereof associated with said first
 external destination,

10 ii) zero, one or more data or copies
 thereof requested in information packets
 received from by the second input of that
 routing interconnect element and routed for
 transmission by the first output of that
 routing interconnect element, and

15 C. selectively transferring an information
 packet received by at least one of the first and
 second input of each of said routing interconnect
 elements to at least one of the first and second
 output of that routing interconnect element based on
20 correspondence, or lack thereof, between the
 descriptor in that information packet and an entry
 signal allocated to that directory,

 D. said selectively transferring step
25 including the steps of

30 i) responding to receipt, at the second
 input of that routing interconnect element,
 of a packet including a request for a datum
 or copy thereof referenced by a descriptor
 that does not correspond to an entry signal
 allocated to the directory of that routing
 interconnect element for a datum or copy
 thereof associated with said first external

destination for routing that packet from that second input to the first output of that routing interconnect element, and

5 ii) responding to receipt, at the
second input of that routing interconnect
element, of a packet including a request for
a datum or copy thereof referenced by a
10 descriptor that corresponds to a valid entry
signal allocated to the directory of that
routing interconnect element for a datum or
copy thereof associated with said first
external destination for routing that packet
from that second input to the second output
15 of that routing interconnect element.

58. A method according to claim 57, comprising
selectively responding within at least one of said
first and second routing interconnect elements to at
20 least one of

 i) the transfer of a packet including a
datum or copy thereof from the first input
to the second output of the associated
25 routing interconnect element, and

 ii) the transfer of a packet including
a datum or copy thereof from the second
input to the first output of the associated
30 routing interconnect element

for at least one of allocating, deallocating or
modifying an entry signal in the directory of that
routing interconnect element.

35

59. A method according to claim 58, comprising selectively responding within at least one of said first and second routing interconnect elements to transfer of a packet including a request for a datum
5 or copy thereof from the second input to the first output of that routing interconnect element for at least one of allocating, deallocating and modifying an entry signal in the directory of that element.

10

60. A method according to claim 59, comprising disposing said first and second routing interconnect elements physically remotely from one another.

15

61. A method according to claim 60, wherein at least one information packet generated by said first and second external sources includes a requester-ID signal associated with an external destination to
20 which that packet is directed, said method comprising

A. responding within at least one of said first and second routing interconnect elements to receipt on any of the first and second inputs of that
25 routing interconnect element of at least a selected packet including a requester-ID signal associated with the external destination to which the first output transmits information packets for routing that packet from that first input to that first output, and

30

B. responding within at least one of said first and second routing interconnect elements to receipt on any of the first and second inputs of that routing interconnect element of at least a selected

-250-

packet including a requester-ID signal that is not associated with the external destination to which the first output transmits information packets for routing that packet from that first input to the
5 second output of that routing interconnect element.

62. A method according to any one of claims 51-61, comprising:

10

A. responding within the first routing interconnect element to receipt, at the first input of that routing interconnect element, of a packet including a request for a datum or copy thereof
15 referenced by a descriptor that corresponds to an entry signal allocated to the directory of that routing interconnect element for a datum or copy thereof associated with said first external destination for routing that packet from that first
20 input to the first output of that routing interconnect element, and

B. responding within the first routing interconnect element to receipt, at the first input
25 of that routing interconnect element, of a packet including a request for a datum or copy thereof referenced by a descriptor that does not correspond to an entry signal allocated to the directory of that routing interconnect element for a datum or copy
30 thereof associated with said first external destination for routing that packet from that first input to the second output of that routing interconnect element.

35

-251-

63. A method according to any one of claims 51-61, comprising:

A. responding within the second routing
5 interconnect element to receipt, at the first input
of that routing interconnect element, of a packet
including a request for a datum or copy thereof
referenced by a descriptor that does not correspond
to an entry signal allocated to the directory of that
10 routing interconnect element for a datum or copy
thereof associated with said first external
destination for routing that packet from that first
input to the first output of that routing
interconnect element, and

15

B. responding within the second routing
interconnect element to receipt, at the first input
of that routing interconnect element, of a packet
including a request for a datum or copy thereof
20 referenced by a descriptor that corresponds to a
entry signal allocated to the directory of that
routing interconnect element for a datum or copy
thereof associated with said first external
destination for routing that packet from that first
25 input to the second output of that routing
interconnect element.

64. A method according to any one of claims 55,
30 56, 60 or 61, comprising:

A. transmitting within at least one of said
first and second routing interconnect elements to
information packets from the second output of that

-252-

routing interconnect element to the remotely disposed second input of the other of said routing interconnect elements, and

5 B. receiving information packets at the second input of that routing interconnect element transmitted from the remotely disposed second output of the other of said routing interconnect elements.

10

65. A method according to claim 64, comprising carrying on a fiber optic media signals representative of said information packets between said at least one remotely disposed second input and
15 second output of said routing interconnect elements.

66. A method according to any one of claims 51-61, comprising selectively responding within at least one
20 of said routing interconnect elements to receipt at said first input of a packet including a response to a request packet previously routed from the second input to the first output of that routing interconnect element for routing that packet from
25 that first input to the second output of that routing interconnect element.

67. A method according to claim 66, comprising
30 selectively responding within at east said second routing interconnect element to receipt at said first input of a packet including a datum or copy thereof for which at least one of

i) a read-only request has previously been transferred from the second input of that routing interconnect element to the first output of that routing interconnect element, and

5

ii) an invalid copy is associated to said first external destination,

10 for generating an information packet including a copy of that datum or copy thereof and for transferring that second duplicate data packet to said second output.

15

68. A method according to claim 66, comprising selectively responding within at least said second routing interconnect element to receipt at said first input of a packet including a datum or copy thereof
20 for which any of an ownership request and an atomic request packet has been transferred from the second input to the first output of that routing interconnect element for transferring that packet received at said first input to the second output of
25 that routing interconnect element.

-254-

69. A method according to any one of claims 51-61, comprising storing within the directory of at least one of said routing interconnect elements association state-defining information for at least selected
5 data, or copies thereof, associated with said first external destination, said association state-defining information including at least one of

10 i) an invalid association state for an invalid datum associated with said first external destination,

15 ii) a read-only association state for a copy of a datum associated with and subject to read-only access within said first external destination, no copy of which datum is modifiable within said first external destination,

20 iii) a nonexclusive owner association state for a datum associated with and modifiable within said first external destination, one or more read-only copies of which datum are associated with at least one
25 of said first and second external destinations, and

30 iv) an exclusive owner association state for a datum exclusively associated with and modifiable within said first external destination, no read-only copies of which datum are associated with either of said first and second external destinations.

-255-

70. A method according to any one of claims 51-61, comprising storing in the directory of at least one of said routing interconnect elements pending
5 state-defining information for at least selected data for which requests are pending to the associated external destination, said pending state-defining information including at least one of
- 10 i) a none-pending pending state for a datum for which no requests are pending,
- ii) a read-only pending state for a datum for which a read-only association
15 request is pending,
- iii) an ownership pending state for a datum for which an ownership association request is pending, or for which an
20 invalidation request is pending, and
- iv) a reissue pending state for a datum for which a request is pending and for which a response other than one generated in
25 response to that request has issued, which response was not routed from the first input to the second output of that routing interconnect element.
- 30
71. A method according to any one of claims 51-61, comprising selectively responding, within at least one of said routing interconnect elements, to receipt, at the second input means of that routing

interconnect element, of multiple information packets requesting association of a datum for transferring an initial one of those request packets to the first output means of that routing interconnect element and
5 for modifying later received ones of such request packets to indicate that a responder is busy and transferring those modified packets to the second output means of that routing interconnect element.

10

72. A method according to claim 71, comprising

A. selectively responding, within said at least one routing interconnect element, to receipt,
15 at the second input means of that routing interconnect element, of multiple information packets requesting read-only allocation of a datum for transferring an initial one of those request packets to the first output means of that associated routing
20 interconnect element and for modifying later received ones of such request packets to indicate that a responder is busy and transferring those modified packets to the second output means of that associated routing interconnect element, and

25

B. selectively responding, within said at least one routing interconnect element, to receipt, at the second input means of that routing interconnect element, of one or more information
30 packets requesting ownership allocation or invalidation of a datum for transferring those packets to the to the first output means of that associated routing interconnect element.

73. A method according to claim 72, comprising

A. selectively responding, within said at
5 least one routing interconnect element, to receipt,
at the second input means of that routing
interconnect element, of multiple information packets
requesting read-only allocation of a datum for
transferring an initial one of those request packets
10 to the first output means of that associated routing
interconnect element and for modifying later received
ones of such request packets to indicate that a
responder is busy and transferring those modified
packets to the second output means of that associated
15 routing interconnect element,

B. selectively responding, within said at
least one routing interconnect element, to receipt,
at the second input means of that routing
20 interconnect element, of multiple information packets
requesting exclusive ownership of a datum for
transferring an initial one of those request packets
to the first output means of that associated routing
interconnect element and for modifying later received
25 ones of such request packets to indicate that a
responder is busy and transferring those modified
packets to the second output means of that associated
routing interconnect element,

30 C. selectively responding, within said at
least one routing interconnect element, to receipt,
at the second input means of that routing
interconnect element, of multiple information packets
requesting, without wait, atomic ownership of a datum

for transferring an initial one of those request packets to the first output means of that associated routing interconnect element and for modifying later received ones of such request packets to indicate
5 that the requested datum is atomically owned and transferring those modified packets to the second output means of that associated routing interconnect element,

10 D. selectively responding, within said at least one routing interconnect element, to receipt, at the second input means of that routing interconnect element, of multiple information packets requesting, with wait, atomic ownership of a datum
15 for transferring an initial one of those request packets to the first output means of that associated routing interconnect element and for modifying subsequently received ones of such request packets, as well as to subsequently received packets
20 requesting atomic ownership without wait, to indicate that the requested datum is atomically owned and transferring those modified packets to the second output means of that associated routing interconnect element, and

25 E. selectively responding, within said at least one routing interconnect element, to receipt, at the second input means of that routing interconnect element, of an unresolved information
30 packet requesting a datum, which information packet was previously issued by said multiple request-handling means, for modifying that packet to

-259-

indicate that a responder is busy and transferring that modified packets to the second output means of that associated routing interconnect element.

5

74. A method according to claim 71, comprising storing within the directory of at least one of said routing interconnect elements pending state-defining information for at least selected data
10 for which requests are pending to the associated external destination, said pending state-defining information including at least one of

15

i) a none-pending pending state for a datum for which no requests are pending,

20

ii) a read-only pending state for a datum for which a read-only association request is pending,

25

iii) an exclusive ownership pending state for a datum for which an exclusive ownership association request is pending,

30

iv) an ownership invalidate pending state for a datum for which an invalidate request is pending;

v) an GET pending state for a datum for which an atomic ownership request, without wait, is pending,

35

vi) a GETW pending state for a datum for which an atomic ownership request, with wait, is pending, and

-260-

vii) an atomic pending state for a datum for which an atomic ownership request, with wait, is pending, which data is
5 previously associated with said atomic state in one of said first and second external destinations, and

iv) a reissue pending state for a datum
10 for which a request is pending and for which a response other than one generated in response to that request has issued, which response was not routed from the first input to the second output of that routing
15 interconnect element.

75. A method according to any one of claims 51-61, comprising
20

A. selectively responding, within said at least one routing interconnect element, to receipt, at the second input means of that routing interconnect element, to multiple information packets
25 requesting allocation of a datum for transferring an initial one of those request packets to the first output means of that associated routing interconnect elements and for deleting later issued ones of said request packets, and

30 B. associating a merged-state signal with a datum for which such later issued ones of said requests have been deleted.

76. A method according to claim 75, comprising selectively responding, within said at least one routing interconnect elements, to receipt, at the
5 second input means of that routing interconnect element, of a packet including a datum associated with said merged-state signals for modifying that packet to include a merged-response signal indicating that a requestor to which that response packet is
10 directed should not delete that packet.

77. A method according to claim 75, comprising

15 A. selectively responding, within said at least one routing interconnect element, to receipt, at the second input means of that routing interconnect element, of multiple information packets requesting like allocation of a datum for
20 transferring an initial one of those request packets to the first output means of that associated routing interconnect elements and for deleting later issued ones of said like request packets, and

25 B. associating a merged-state signal with a datum for which such later issued ones of such like requests have been deleted.

30 78. A method according to claim 75, comprising storing within the directory of at least one of said routing interconnect elements pending state-defining information for at least selected data for which

requests are pending to the associated external destination, said pending state-defining information including at least one of

- 5 i) a none-pending pending state for a datum for which no requests are pending,
- 10 ii) a wait pending state for a datum for which a request packet is pending and for which a response has previously issued, which request packet is to be emptied upon receipt by the routing interconnect cell,
- 15 iii) a wait-merge-read-only pending state for a datum for which multiple information packets requesting read-only association have been received and for which a single read-only request packet is pending, which request packet is to be
20 emptied upon receipt by the routing interconnect element,
- 25 iv) a wait-merge-exclusive pending state for a datum for which multiple information packets requesting exclusive ownership association have been received and for which a single exclusive ownership request packet is pending, which request
30 packet is to be emptied upon receipt by the routing interconnect element,
- v) a wait-merge-atomic pending state for a datum for which multiple information packets requesting atomic ownership

5 association have been received and for which
a single atomic ownership request packet is
pending, which request packet is to be
emptied upon receipt by the routing
interconnect element,

10 vi) a reissue pending state for a datum
for which a request is pending, which
request is to be reissued upon receipt by
the routing interconnect element,

15 vii) a reissue-merge-read-only pending
state for a datum for which multiple
information packets requesting read-only
association have been received by the
routing interconnect element, and for which
a single read-only request is pending, which
request is to be reissued upon receipt by
the routing interconnect element,
20

25 viii) a reissue-merge-exclusive pending
state for a datum for which multiple
information packets requesting exclusive
ownership association have been received by
the routing interconnect element, and for
which a single exclusive ownership request
is pending, which request is to be reissued
upon receipt by the routing interconnect
element,

30 ix) a reissue-merge-atomic pending
state for a datum for which multiple
information packets requesting atomic
ownership association have been received by

5 the routing interconnect element, and for which a single atomic ownership request is pending, which request is to be reissued upon receipt by the routing interconnect element,

10 x) a read-only pending state for a datum for which a read-only association request is pending,

15 xi) a read-only-merge pending state for a datum for which multiple information packets requesting read-only association have been received and for which a single read-only request packet is pending, which request packet is to be marked as merged upon receipt by the routing interconnect element,

20 xii) an exclusive ownership pending state for a datum for which an exclusive ownership association request is pending,

25 xiii) an exclusive-merge pending state for a datum for which an exclusive ownership request packet was pending upon receipt by the routing interconnect unit of a further request for association with that datum, which request packet is to be marked as merged upon receipt by the routing
30 interconnect element,

xiv) an atomic pending state for which an atomic ownership request is pending,

-265-

- 5 xv) an atomic-merge pending state for a datum for which an atomic ownership request packet was pending upon receipt by the routing interconnect unit of a further request for association with that datum, which request packet is to be marked as merged upon receipt by the routing interconnect element,
- 10 xvi) a recombine pending state for a packet including that datum, routed for recombine association with an external destination, is pending, and
- 15 xvii) a merged response pending state for a datum for which a response packet marked as merged is pending and for which a further request has been received by the routing interconnect element.
- 20

79. A method according to any one of claims 51-61, comprising generating within at least one of said
25 routing interconnect elements an association-based routing signal indicative of a packet routing based on correspondence, or lack thereof, between the descriptor in an information packet received at any of said first and second input of that routing
30 interconnect element and an entry signal allocated to that directory.

-266-

80. A method according to any one of claims 51-61, comprising

A. buffering within an extract buffer of
5 at least one of said first and second routing
interconnect elements information packets transferred
between the first input and the second output of that
routing interconnect element, and

10 B. buffering within an insert buffer of at
least one of said first and second routing
interconnect information packets transferred between
said second input and the first output of that
routing interconnect element.

15

81. A method according to any one of claims 51-61,
comprising buffering information packets transferred
from the extract buffer to the second output of at
20 least one of said routing interconnect elements.

82. A method according to any one of claims 51-61,
comprising storing in an error queue in at least one
25 of said routing interconnect elements one or more
packets in which an error is detected.

83. A method according to any one of claims 51-61,
30 comprising buffering in an incoming buffer packet
received at the second input of at least one of said
routing interconnect elements.

-267-

84. A switch for selectively routing digital information packets received from a plurality of external sources to a plurality of external destinations,

5

each said external destination having associated therewith zero, one or more data items, each unique one of said data items having an associated descriptor for use in referencing the
10 associated data item,

said external sources including means for generating information packets, each information packet including a said descriptor of a said data
15 item,

said switch comprising first and second routing interconnect elements, each said routing interconnect element including
20

A. a source input means for receiving an information packet from at least a respective one of said external sources,

25 B. a destination output means for transmitting an information packet to at least a respective one of said external destinations,

C. internal input means for receiving an
30 information packet from the other of said routing interconnect elements, and

-268-

D. internal output means for transmitting an information packet to the other of said routing interconnect elements,

5

E. directory means having allocated therein entries each respectively containing a descriptor corresponding to at least one of

10 i) one of said data items associated with said at least a respective one of said external destinations,

15 ii) a data item requested in an information packet received by the internal input means of the respective routing interconnect element and routed for transmission by the destination output means of the respective routing interconnect
20 element, and

F. control means for selectively transferring an information packet received by at least one of the source input means and the internal
25 input means of the respective routing interconnect element, to at least one of the destination output means and the internal output means of the respective routing interconnect element based on the correspondence, or lack thereof, between the
30 descriptor in that information packet and an entry

-269-

containing a corresponding descriptor in said directory means.

85. A switch according to claim 84, wherein
5 control means of at least one routing interconnect element includes at least one of

A. means responsive to receipt, at the source input means of said at least one routing interconnect element, of a packet including a request
10 for a data item referenced by a descriptor that corresponds to an entry allocated to the directory means of said at least one routing interconnect element for a data item associated with said respective one of said external destinations, for
15 routing that packet from that source input means to the destination output means of said at least one routing interconnect element, and

B. means responsive to receipt, at the
20 source input means of said at least one routing interconnect element, of a packet including a request for a data item referenced by a descriptor that does not correspond to an entry allocated to the directory means of said at least one routing interconnect
25 element for a data item associated with said respective one of said external destinations, for routing that packet from the source input means to the internal output means of said at least one routing interconnect element.

-270-

86. A switch according to claim 85, wherein said control means of the other routing interconnect element includes at least one of

5 A. means responsive to receipt, at the source input means of said other routing interconnect element, of a packet including a request for a data item referenced by a descriptor that corresponds to an entry allocated to the directory means of said other
10 routing interconnect element for a data item associated with said respective one of said external destinations, for routing that packet from the source input means to the destination output means of said other interconnect element, and

15 B. means responsive to receipt, at the source input means of said other routing interconnect element, of a packet including a request for a data item referenced by a descriptor that does not correspond to an entry allocated to the directory
20 means of said other routing interconnect element for a data item associated with said respective one of said external destinations, for routing that packet from the source input means to the internal output means of said other routing interconnect element.

25

87. A switch according to claim 84, wherein said control means of at least one of said routing interconnect elements includes at least one of

30 A. means responsive to receipt, at the internal input means of said at least one of said

-271-

routing interconnect elements, of a packet including a request for a data item referenced by a descriptor that corresponds to an entry allocated to the directory means of said at least one of said routing interconnect elements for a data item associated with said first external destination for routing that packet from that internal input means to the destination output means of said at least one of said routing interconnect elements, and

10

B. means responsive to receipt, at the internal input means of said at least one of said routing interconnect elements, of a packet including a request for a data item referenced by a descriptor that does not correspond to a valid entry allocated to the directory means of said at least one of said routing interconnect elements for a data item associated with said first external destination for routing that packet from that internal input means to the internal output means of that said at least one of said routing interconnect elements.

88. A switch according to claim 87, wherein said control means of the other of said routing interconnect elements includes at least one of

25

A. means responsive to receipt, at the internal input means of said other of said routing interconnect elements, of a packet including a request for a data item referenced by a descriptor that corresponds to an entry allocated to the directory means of said other of said routing interconnect

30

-272-

elements for a data item associated with said first
external destination for routing that packet from that
internal input means to the destination output means
of said other of said routing interconnect elements,
5 and

B. means responsive to receipt, at the
internal input means of said other of said routing
interconnect elements, of a packet including a request
10 for a data item referenced by a descriptor that does
not correspond to a valid entry allocated to the
directory means of said other of said routing
interconnect elements for a data item associated with
said first external destination for routing that
15 packet from that internal input means to the internal
output means of said other of said routing
interconnect elements.

89. A switch according to claim 84, wherein said
20 control means of at least one of said routing
interconnect elements includes means for at least one
of allocating, deallocating or modifying an entry in
the directory means of that element in association
with the packet including a data item from at least
25 one of

i) the source input means to the
internal output means of the associated
routing interconnect element, and

30

-273-

ii) the internal input means to the destination output means of the associated routing interconnect element.

5 90. A switch according to claim 87, wherein said control means of said at least one of said routing interconnect elements includes means selectively responsive to transfer of a packet including a request for a data item from the internal input means to the
10 destination output means of the associated routing interconnect element for at least one of allocating, deallocating and modifying an entry in the directory means of that element.

15 91. A switch according to claim 89, wherein said control means of the other of said routing interconnect elements includes means for at least one of allocating, deallocating or modifying an entry in the directory means of that element in association
20 with the packet including a data item from at least one of

i) the source input means to the internal output means of the associated
25 routing interconnect element, and

ii) the internal input means to the destination output means of the associated routing interconnect element.

30

-274-

92. A switch according to claim 91, wherein said control means of said other of said routing interconnect elements includes means selectively responsive to transfer of a packet including a request
5 for a data item from the internal input means to the destination output means of the associated routing interconnect element for at least one of allocating, deallocating and modifying an entry in the directory means of that element.

10

93. A switch according to claim 84, wherein said first and second routing interconnect elements are disposed physically remotely from one another.

15 94. A switch according to any one of claims 84 to 93, wherein at least one information packet generated by at least one of said first and second external sources includes a requester identifier associated with an external destination to which that packet is
20 directed, at least one of said routing interconnect elements comprising at least one of

A. means responsive to receipt on any of the first input means and internal input means of said at
25 least one of said routing interconnect elements of at least a selected packet including a requester identifier associated with the external destination to which the destination output means transmits information packets for routing that packet from that
30 source input means to that destination output means, and

-275-

B. means responsive to receipt on any of the first input means and internal input means of said at least one of said routing interconnect elements of at least a selected packet including a requester identifier that is not associated with the external destination to which the destination output means transmits information packets for routing that packet from that source input means to the internal output means of said at least one of said routing interconnect elements.

95. A switch according to claim 94, wherein at least one information packet generated by at least one of said first and second external sources includes a requester identifier associated with an external destination to which that packet is directed, the other of said routing interconnect elements comprising at least one of

20

A. means responsive to receipt on any of the first input means and internal input means of said other of said routing interconnect elements of at least a selected packet including a requester identifier associated with the external destination to which the destination output means transmits information packets for routing that packet from that source input means to that destination output means, and

30

-276-

B. means responsive to receipt on any of the first input means and internal input means of that routing interconnect element of at least a selected packet including a requester identifier that is not
5 associated with the external destination to which the destination output means transmits information packets for routing that packet from that source input means to the internal output means of said other of said routing interconnect elements.

10

96. A switch according to any one of claims 84 to 93, wherein at least one of said external sources or external destinations has a copy of at least one of said data items, the copy having an associated
15 descriptor that corresponds to the descriptor of the data item of which it is a copy.

97. A switch according to claim 96 in which said directory means further has allocated therein an entry
20 containing a descriptor corresponding one of said copies associated with said at least a respective one of said external destinations.

98. A switch according to any one of claims 84 to
25 93, wherein at least one of said routing interconnect elements includes means selectively responsive to receipt at said source input means of a packet including a response to a request packet previously routed from the internal input means to the
30 destination output means of said at least one routing interconnect element for routing that packet from that

-277-

source input means to the internal output means of said at least one routing interconnect element.

99. A switch according to any one of claims 84 to 5 93, wherein the directory means of at least one of said routing interconnect elements includes means for storing association state-defining information for at least selected data items associated with said first external destination, said association state-defining 10 information including at least one of

i) an invalid association state for an invalid data item associated with said first external destination,

15 ii) a read-only association state for a copy of a data item associated with and subject to read-only access within said first external destination, no copy of which data item is modifiable within said first external 20 destination,

iii) a nonexclusive owner association state for a data item associated with and modifiable within said first external 25 destination, one or more read-only copies of which data item are associated with at least one of said first and second external destinations, and

30 iv) an exclusive owner association state for a data item exclusively associated with

-278-

and modifiable within said first external destination, no read-only copies of which data item are associated with either of said first and second external destinations.

5

100. A switch according to any one of claims 84 to 93, wherein the directory means of at least one of said routing interconnect elements includes means for storing pending state-defining information for at least selected data for which requests are pending to the associated external destination, said pending state-defining information including at least one of

15 i) a none-pending pending state for a data item for which no requests are pending,

ii) a read-only pending state for a data item for which a read-only association request is pending,

20

iii) an ownership pending state for a data item for which an ownership association request is pending, or for which an invalidation request is pending, and

25

iv) a reissue pending state for a data item for which a request is pending and for which a response other than one generated in response to that request has issued, which response was not routed from the source input means to the internal output means of said at

30

-279-

least one of said routing interconnect elements.

101. A switch according to any one of claims 83 to
5 94, wherein at least one of said routing interconnect elements includes multiple request-handling means selectively responsive to receipt, at the internal input means of said at least one routing interconnect element, of multiple information packets requesting
10 association of a data item for transferring an initial one of those request packets to the destination output means of said at least one of said routing interconnect elements and for modifying later received ones of such request packets to indicate that a
15 responder is busy and transferring those modified packets to the internal output means of said at least one routing interconnect element.

102. A switch according to any one of claims 83 to
20 94, wherein at least one of said routing interconnect elements include

A. multiple request-handling means selectively responsive to receipt, at the internal
25 input means of the associated routing interconnect element, of multiple information packets requesting allocation of a data item for transferring an initial one of those request packets to the destination output means of the associated one of said routing
30 interconnect elements and for deleting later issued ones of said request packets, and

-280-

B. said multiple request-handling means including merge means for associating a merged-state signal with a data item for which such later issued
5 ones of said requests have been deleted.

103. A switch according to any one of claims 83 to 94, wherein the control means of at least one of said routing interconnect elements includes routing
10 directory unit means connected with the directory means of said at least one of said routing interconnect elements for generating an association-based routing indication indicative of a packet routing based on correspondence, or lack thereof,
15 between the descriptor in an information packet received at any of said first input means and internal input means of said at least one of said routing interconnect elements and an entry allocated to that directory means.

20

104. A switch according to any one of claims 83 to 94, wherein the source input means and the destination output means of each said routing interconnect element are arranged together to form a respective terminal
25 pair comprising a bus connecting one said external source to one said external destination.

105. A switch according to any one of claims 83 to 94, wherein at least one of said routing interconnect
30 elements comprise

-281-

A. extract buffer means connected to said source input means and to said internal output means of said at least one of said routing interconnect elements for buffering information packets transferred
5 therebetween, and

B. insert buffer means connected to said internal input means and to said destination output means of said at least one of said routing
10 interconnect elements for buffering information packets transferred therebetween.

106. A switch according to any one of claims 83 to 94, wherein the internal output of at least one of
15 said routing interconnect elements includes outgoing buffer means for buffering information packets transferred from the corresponding extract buffer means to the internal output means of said at least one of said routing connect elements.

20

107. A switch according to any one of claims 83 to 94, wherein the internal output means of at least one of said routing interconnect elements includes error queue means for storing one or more packets in which
25 an error is detected.

108. A switch according to any one of claims 83 to 94, wherein the internal input means of at least one of said routing interconnect element includes incoming
30 buffer means for buffering packets received at the

-282-

internal input means of said at least one of said
routing interconnect elements.

109. A method of routing digital information
5 packets received from a plurality of external sources
to a plurality of external destinations,

each said external destination having
associated therewith zero, one or more data items,
10 each unique one of said data items having an
associated descriptor for use in referencing the
associated data item,

said external sources including means for
15 generating information packets, each information
packet including a said descriptor of a said data
item,

said method comprising:

20

A. providing first and second routing
interconnect elements, each said routing interconnect
element including

25 i) a source input means for receiving an
information packet from at least a respective
one of said external sources,

ii) a destination output means for
30 transmitting an information packet to at least

-283-

a respective one of said external destinations,

5 iii) internal input means for receiving
an information packet from the other of said
routing interconnect elements, and

 iv) internal output means for
transmitting an information packet to the
10 other of said routing interconnect elements,

B. allocating in a directory means
associated with each routing interconnect element
entries each respectively containing a descriptor
15 corresponding to at least one of

 i) one of said data items associated
with said at least a respective one of said
external destinations,
20

 ii) a data item requested in an
information packet received by the internal
input means of the respective routing
interconnect element and routed for
25 transmission by the destination output means
of the respective routing interconnect
element, and

C. transferring an information packet
30 received by at least one of the source input means and
the internal input means of the associated routing

-284-

interconnect element, to at least one of the destination output means and the internal output means of based on the correspondence, or lack thereof, between the descriptor in that information packet and
5 an entry containing a corresponding descriptor in said directory means.

110. A method according to claim 109, wherein said information packet transfer step of at least one
10 routing interconnect element includes at least one of the steps of

A. in response to receipt, at the source input means of said at least one routing interconnect
15 element, of a packet including a request for a data item referenced by a descriptor that corresponds to an entry allocated to the directory means of said at least one routing interconnect element for a data item associated with said respective one of said external
20 destinations, routing that packet from that source input means to the destination output means of said at least one routing interconnect element, and

B. in response to receipt, at the source
25 input means of said at least one routing interconnect element, of a packet including a request for a data item referenced by a descriptor that does not correspond to an entry allocated to the directory means of said at least one routing interconnect
30 element for a data item associated with said respective one of said external destinations, routing

-285-

that packet from the source input means to the internal output means of said at least one routing interconnect element.

5 111. A method according to claim 110, wherein said information packet transfer step of the other routing interconnect element includes at least one of the steps of

10 A. in response to receipt, at the source input means of said other routing interconnect element, of a packet including a request for a data item referenced by a descriptor that corresponds to an entry allocated to the directory means of said other
15 routing interconnect element for a data item associated with said respective one of said external destinations, routing that packet from the source input means to the destination output means of said other interconnect element, and

20

B. in response to receipt, at the source input means of said other routing interconnect element, of a packet including a request for a data item referenced by a descriptor that does not
25 correspond to an entry allocated to the directory means of said other routing interconnect element for a data item associated with said respective one of said external destinations, routing that packet from the source input means to the internal output means of
30 said other routing interconnect element.

-286-

112. A method according to claim 109, wherein said information packet transfer step of at least one of said routing interconnect elements includes at least one of the steps of

5

A. in response to receipt, at the internal input means of said at least one routing interconnect element, of a packet including a request for a data item referenced by a descriptor that corresponds to an entry allocated to the directory means of said at least one of said routing interconnect elements for a data item associated with said first external destination, routing that packet from that internal input means to the destination output means of said at least one of said routing interconnect elements, and

B. in response to receipt, at the internal input means of said at least one routing interconnect element, of a packet including a request for a data item referenced by a descriptor that does not correspond to a valid entry allocated to the directory means of said at least one routing interconnect element for a data item associated with said first external destination for routing that packet from that internal input means to the internal output means of said at least one routing interconnect element.

113. A method according to claim 112, wherein said information packet transfer step of the other of said routing interconnect elements includes at least one of the steps of

30

-287-

A. in response to receipt, at the internal input means of the other routing interconnect element, of a packet including a request for a data item
5 referenced by a descriptor that corresponds to an entry allocated to the directory means of said other routing interconnect element for a data item associated with said first external destination, routing that packet from that internal input means to
10 the destination output means of said other routing interconnect element, and

B. in response to receipt, at the internal input means of said other routing interconnect
15 element, of a packet including a request for a data item referenced by a descriptor that does not correspond to a valid entry allocated to the directory means of said other routing interconnect element for a data item associated with said first external
20 destination, routing that packet from that internal input means to the internal output means of said other routing interconnect element.

114. A method according to claim 109, wherein said
25 information packet transfer step of at least one of said routing interconnect elements includes at least one of the steps of allocating, deallocating or modifying an entry in the directory means of said at least one of said routing interconnect elements in
30 association with the packet including a data item from at least one of

-288-

i) the source input means to the internal output means of the associated routing interconnect element, and

5

ii) the internal input means to the destination output means of the associated routing interconnect element.

10 115. A method according to claim 114, wherein said information packet transfer step of said at least one of said routing interconnect elements includes the step of, in selective response to transfer of a packet including a request for a data item from the internal
15 input means to the destination output means of said at least one of said associated routing interconnect elements, at least one of allocating, deallocating and modifying an entry in the directory means of said at least one of said associated routing interconnect
20 elements.

116. A method according to claim 114, wherein said information packet transfer step of the other of said routing interconnect elements includes at least one of
25 the steps of allocating, deallocating or modifying an entry in the directory means of the other of said routing interconnect elements in association with the packet including a data item from at least one of

-289-

i) the source input means to the internal output means of the associated routing interconnect element, and

5 ii) the internal input means to the destination output means of the associated routing interconnect element.

117. A method according to claim 116, wherein said
10 information packet transfer step of said other of said routing interconnect elements includes the step of, in selective response to transfer of a packet including a request for a data item from the internal input means to the destination output means of said other of said
15 routing interconnect elements, at least one of allocating, deallocating and modifying an entry in the directory means of said other of said routing interconnect elements.

20 118. A method according to claim 109, wherein said first and second routing interconnect elements are disposed physically remotely from one another.

119. A method according to any one of claims 109 to
25 118, wherein at least one information packet generated by at least one of said first and second external sources includes a requester identifier associated with an external destination to which that packet is directed, at least one of said routing interconnect
30 elements performing at least one of the steps of

-290-

A. in response to receipt on any of the first input means and internal input means of said at least one of said routing interconnect elements of at least a selected packet including a requester
5 identifier associated with the external destination to which the destination output means transmits information packets, routing that packet from that source input means to that destination output means, and

10

B. in response to receipt on any of the first input means and internal input means of said at least one of said routing interconnect elements of at least a selected packet including a requester
15 identifier that is not associated with the external destination to which the destination output means transmits information packets, routing that packet from that source input means to the internal output means of said at least one of said routing
20 interconnect elements.

120. A method according to claim 119, wherein at least one information packet generated by at least one of said first and second external sources includes a
25 requester identifier associated with an external destination to which that packet is directed, the other of said routing interconnect elements performing at least one of the steps of

30 A. in response to receipt on any of the first input means and internal input means of said

-291-

other of said routing interconnect elements of at least a selected packet including a requester identifier associated with the external destination to which the destination output means transmits
5 information packets, routing that packet from that source input means to that destination output means, and

B. in response to receipt on any of the
10 first input means and internal input means of said other of said routing interconnect elements of at least a selected packet including a requester identifier that is not associated with the external destination to which the destination output means
15 transmits information packets, routing that packet from that source input means to the internal output means of said other of said routing interconnect elements.

20 121. A method according to any one of claims 109 to 118, wherein at least one of said external sources or external destinations has a copy of at least one of said data items, the copy having an associated descriptor that corresponds to the descriptor of the
25 data item of which it is a copy.

122. A method according to claim 121 in which said directory means further has allocated therein an entry containing a descriptor corresponding one of said
30 copies associated with said at least a respective one of said external destinations.

-292-

123. A method according to any one of claims 109 to 118, wherein at least one of said routing interconnect elements performs the step of, in selective response
5 to receipt at said source input means of a packet including a response to a request packet previously routed from the internal input means to the destination output means of said at least one routing interconnect element, routing that packet from that
10 source input means to the internal output means of said at least one routing interconnect element.

124. A method according to any one of claims 109 to 118, wherein the directory means of at least one of
15 said routing interconnect elements includes means for storing association state-defining information for at least selected data items associated with said first external destination, said association state-defining information including at least one of

20

i) an invalid association state for an invalid data item associated with said first external destination,

25

ii) a read-only association state for a copy of a data item associated with and subject to read-only access within said first external destination, no copy of which data item is modifiable within said first external
30 destination,

-293-

iii) a nonexclusive owner association state for a data item associated with and modifiable within said first external destination, one or more read-only copies of which data item are associated with at least one of said first and second external destinations, and

iv) an exclusive owner association state for a data item exclusively associated with and modifiable within said first external destination, no read-only copies of which data item are associated with either of said first and second external destinations.

15

125. A method according to any one of claims 109 to 118, wherein the directory means of at least one of said routing interconnect elements includes means for storing pending state-defining information for at least selected data for which requests are pending to the associated external destination, said pending state-defining information including at least one of

i) a non-pending pending state for a data item for which no requests are pending,

ii) a read-only pending state for a data item for which a read-only association request is pending,

30

-294-

iii) an ownership pending state for a data item for which an ownership association request is pending, or for which an invalidation request is pending, and

5

iv) a reissue pending state for a data item for which a request is pending and for which a response other than one generated in response to that request has issued, which response was not routed from the source input means to the internal output means of said at least one of said routing interconnect elements.

10

15 126. A method according to any one of claims 109 to 118, wherein at least one of said routing interconnect elements performs the steps of, in selective response to receipt, at the internal input means of said at least one routing interconnect element, of multiple
20 information packets requesting association of a data item,

A. transferring an initial one of those request packets to the destination output means of
25 said at least one of said routing interconnect elements,

B. modifying later received ones of such request packets to indicate that a responder is busy,
30 and

-295-

C. transferring those modified packets to the internal output means of said at least one routing interconnect element.

5 127. A method according to any one of claims 109 to 118, wherein at least one of said routing interconnect elements performs the steps of

A. in selective response to receipt, at the
10 internal input means of the at least one of said routing interconnect elements, of multiple information packets requesting allocation of a data item, transferring an initial one of those request packets to the destination output means of that associated
15 routing interconnect elements and for deleting later issued ones of said request packets, and

B. associating a merged-state signal with a data item for which such later issued ones of said
20 requests have been deleted.

128. A method according to any one of claims 109 to 118, wherein the information packet transfer step of at least one of said routing interconnect elements
25 includes the step of generating an association-based routing indication indicative of a packet routing based on correspondence, or lack thereof, between the descriptor in an information packet received at any of said first input means and internal input means of
30 said at least one of said routing interconnect elements and an entry allocated to the directory means

-296-

of said at least one of said routing interconnect elements.

129. A method according to any one of claims 109 to
5 118, wherein the source input means and the destination output means of each said routing interconnect element are arranged together to form a respective terminal pair comprising a bus connecting one said external source to one said external
10 destination.

130. A method according to any one of claims 109 to
118, further including the step of providing at least one of said routing interconnect elements with
15

A. extract buffer means connected to said source input means and to said internal output means of said at least one of said routing interconnect elements for buffering information packets transferred
20 therebetween, and

B. insert buffer means connected to said internal input means and to said destination output means of said at least one of said routing
25 interconnect elements for buffering information packets transferred therebetween.

131. A method according to any one of claims 109 to
118, further including the step of providing the
30 internal output of at least one of said routing interconnect elements with outgoing buffer means for

-297-

buffering information packets transferred from the corresponding extract buffer means to the internal output means of said at least one of said routing connect elements.

5

132. A method according to any one of claims 109 to 118, further including the step of providing the internal output means of at least one of said routing interconnect elements with error queue means for
10 storing one or more packets in which an error is detected.

133. A method according to any one of claims 109 to 118, further including the step of providing the
15 internal input means of at least one of said routing interconnect element with incoming buffer means for buffering packets received at the internal input means of said at least one of said routing interconnect elements.




FIG. 2A





FIG. 2B

**FIG. 3**






FIG. 4

**FIG. 5**

**FIG. 6**




FIG. 7




FIG. 8





FIG. 9



20

16 | 15

11 | 10

0

CHECK(4:0)

NEXT FREE
L2 ENTRY POINTER

LEVEL2 VENTRY (INVALID)

20

16 | 15

0

CHECK(4:0)

L2 SUBENTRY VALID (15:0)

LEVEL2 VENTRY (VALID)

24

20 | 19

11 | 10

0

CHECK(4:0)

NEXT FREE
L3 ENTRY POINTER

LEVEL3 VENTRY (INVALID)

24

20 | 19

18

17

16

15

0

CHECK(4:0)

DUP
LIMIT

OWNER
LIMIT




CREATED
PEND

PAGE
PEND

L2 SUB
ENTRY VALID(15:0)

LEVEL3 VENTRY (VALID)

FIG. 10

**FIG. 11A****FIG. 11B****FIG. 12**




FIG. 13




FIG. 14A




FIG. 14B





FIG. 15



FULL RING FREEZE

FIG. 16

Ring:1 Pending State	Ring:0 Subpage State			
	Invalid	ReadOnly	Nonexclusive	Exclusive
PINone	InvPINone	RoPINone	NexPINone	ExPINone
PIWait	InvPIWait	RoPIWait	NexPIWait	ExPIWait
PIReissue	InvPIRei	RoPIRei	NexPIRei	ExPIRei
PIReissueMrgRo	—	RoPIReiMrgRo	NexPIReiMrgRo	ExPIReiMrgRo
PIReissueMrgEx	—	—	NexPIReiMrgEx	ExPIReiMrgEx
PIReissueMrgAt	—	—	NexPIReiMrgAt	ExPIReiMrgAt
PIReadOnly	—	RoPIRo	NexPIRo	ExPIRo
PIReadOnlyMrgRo	—	RoPIRoMrgRo	NexPIRoMrgRo	ExPIRoMrgRo
PIExclusive	—	—	NexPIEx	ExPIEx
PIExclusiveMrg	—	—	NexPIExMrg	ExPIExMrg
PIAtomic	—	—	NexPIAt	ExPIAt
PIAtomicMrg	—	—	NexPIAtMrg	ExPIAtMrg
PIRecombine	—	—	—	—
PIMrsp	—	RoPIMrsp	NexPIMrsp	ExPIMrsp
PIMrspMrgRo	—	RoPIMrspMrgRo	NexPIMrspMrgRo	ExPIMrspMrgRo
PIMrspMrgEx	—	—	NexPIMrspMrgEx	ExPIMrspMrgEx
PIMrspMrgAt	—	—	NexPIMrspMrgAt	ExPIMrspMrgAt
PIInvalidate	—	RoPIInv	—	—
PIInvalidateMrg	—	RoPIInvMrg	—	—

FIG. 17

Ring: 0 Pending State	Ring: 0 Subpage State			
	Invalid	ReadOnly	Nonexclusive	Exclusive
PONone	InvPONone	RoPONone	NexPONone	ExPONone
POWait	InvPOWait	RoPOWait	NexPOWait	ExPOWait
POReissue	InvPORei	RoPORei	NexPORei	ExPORei
POReissueMrgRo	InvPOReiMrgRo	—	—	—
POReissueMrgEx	InvPOReiMrgEx	RoPOReiMrgEx	—	—
POReissueMrgAt	InvPOReiMrgAt	RoPOReiMrgAt	—	—
POReadOnly	InvPORo	—	—	—
ROReadOnlyMrgRo	InvPORoMrgRo	—	—	—
POExclusive	InvPOEx	RoPOEx	—	—
POExclusiveMrg	InvPOExMrg	RoPOExMrg	—	—
POAtomic	InvPOAt	RoPOAt	—	—
POAtomicMrg	InvPOAtMrg	RoPOAtMrg	—	—
PORecombine	InvPORec	—	—	—
POMrsp	—	—	—	—
POMrspMrgRo	—	—	—	—
POMrspMrgEx	—	—	—	—
POMrspMrgAt	—	—	—	—
POInvalidate	—	—	NexPOInv	—
POInvalidateMrg	—	—	—	—

FIG. 18




FIG. 19A




FIG. 19B




FIG. 19C




FIG. 19D

