



(51) International Patent Classification:

H04L 12/28 (2006.01) **G06F 13/14** (2006.01)
H04L 12/56 (2006.01)

(21) International Application Number:

PCT/US2011/042431

(22) International Filing Date:

29 June 2011 (29.06.2011)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

12/827,495 30 June 2010 (30.06.2010) US

(71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **KAUSCHKE, Michael** [DE/DE]; Am Meierhof 19, 38122 Braunschweig (DE). **DOSHI, Gautam B.** [US/IN]; O-73 Diamond District, Airport Road, Bangalore 560008 (IN).

(74) Agents: **TROP, Timothy N.** et al.; TROP, PRUNER & HU, P.C., 1616 S. Voss Rd., Ste. 750, Houston, Texas 77057-2631 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: PROVIDING A BUFFERLESS TRANSPORT METHOD FOR MULTI-DIMENSIONAL MESH TOPOLOGY

100




FIG. 1

(57) Abstract: In one embodiment, the present invention includes a method for determining whether a packet received in an input/output (I/O) circuit of a node is destined for the node and if so, providing the packet to an egress queue of the I/O circuit and determining whether one or more packets are present in an ingress queue of the I/O circuit and if so, providing a selected packet to a first or second output register according to a global schedule that is independent of traffic flow. Other embodiments are described and claimed.

PROVIDING A BUFFERLESS TRANSPORT METHOD FOR MULTI-DIMENSIONAL MESH TOPOLOGY

Background

As semiconductor technology advances in the form of smaller and faster circuitry, more and more components can be implemented in a single semiconductor die. One type of die is a so-called system-on-chip (SoC) in which multiple agents are present that can be used to process data and perform other operations. Interconnectivity between these agents is achieved by a network-on-chip (NoC), which can be based on various methods including a packet-routing network, a bus, ring or so forth. However, these approaches can have issues with regard to design complexity, limited scalability, or area/power efficiency. Furthermore, some of the known interconnections may exhibit complex deadlock/contention problems to be resolved.

Most implementations include some type of buffer medium as part of the transport mechanism between agents, given that due to traffic considerations, latencies and so forth, multiple packets may be present in the transport mechanism between agents. While a ring can provide for a bufferless medium, such rings have historically been of a single dimension such that the ring concept does not leverage to multi-dimensions.

Brief Description of the Drawings

FIG. 1 is a block diagram of a mesh interconnection system in accordance with one embodiment of the present invention.

FIG. 2 is a block diagram of a mesh interconnection system that is configured in a pass-through state in accordance with an embodiment of the present invention.

FIG. 3 is a block diagram of a mesh interconnection system that is configured in a turn state in accordance with an embodiment of the present invention.

FIG. 4 is a block diagram of an interconnection system in accordance with another embodiment of the present invention.

FIG. 5 is a block diagram of a node in accordance with an embodiment of the present invention.

FIG. 6 is a flow diagram of a method for communicating packets according to a static global cyclic schedule in accordance with an embodiment of the present invention.

Detailed Description

In various embodiments, bufferless and routerless intersecting rings, so-called “routerless I-rings,” may be provided to enable interconnections of on-chip nodes such as may be present in a system on chip (SoC) or other semiconductor device. Such an SOC
5 may include various types of designs including multi-core and many-core central processing unit (CPU) designs with high agent counts including agents such as cores, caches, accelerators, memory controllers and so forth.

An interconnect network in accordance with an embodiment of the present invention can be designed to be of two or higher dimensions. In general, the network
10 includes: (i) a bufferless transport medium; (ii) an ingress access point (for scheduling packets into the transport medium); and (iii) an egress receiver point (to pull packets from the transport medium). In many embodiments, the bufferless transport medium may include a set of synchronously clocked nodes arranged in an nxm mesh. Each such node may include networking related hardware, and in turn may be coupled to a corresponding
15 agent such as a core, cache, accelerator, etc.

Referring now to FIG. 1, shown is a block diagram of a mesh system in accordance with one embodiment of the present invention. As shown in FIG. 1, system 100 may be implemented as a two-dimensional (2D) link interconnect mesh including a plurality of nodes $110_{1,1} - 110_{n,m}$. In various embodiments, system 100 may be a system-on-chip
20 (SoC), or other such single chip semiconductor device that is formed on one semiconductor die, while in other embodiments a system may be formed of multiple-stacked die, thus enabling a mesh having greater than two dimensions. In different implementations, the agent may be heterogeneous or homogeneous such as processor cores, cache blocks, memory controllers, dedicated processing units such as graphics
25 processors, packet processors and so forth. In general, a node may be associated with an agent (e.g., core or other processing or control unit), and may include interconnect related input/output (I/O) circuitry as described below.

Still referring to FIG. 1, each node 110 may be coupled to multiple other nodes of the 2D configuration via a plurality of interconnects 115, which may be bidirectional
30 interconnects. In addition, the edge nodes (that is, the nodes at the periphery of the system) may include feedback interconnects 120 which couple outputs ports of the corresponding nodes to input ports of the same or other edge node. In the high level view of FIG. 1, only I/O circuitry (e.g., routing structures) are shown. That is, there is no

illustration in FIG. 1 of the actual functional units of the agents. In general, each node may include a plurality of input and output ports. In general, these input and output ports may include logic and other circuitry to effect the selection and transfer of received information (such as packets) through the mesh interconnection. Note that the
5 implementation shown in FIG. 1 is representative, and does not fully reflect the full switching or configurations available within a node. That is, as will be discussed further below, each node can be configured according to multiple routing patterns to route data to different nodes to which a given node may be connected.

As seen in FIG. 1, each node 110 has four bidirectional links to its neighboring
10 nodes (with special feedback rules at the edge nodes). The number of links may increase with the number of dimensions. In each clock cycle, at least one packet can be received by each of a plurality of input ports of each node. The incoming packets may be locally consumed or passed to one of a plurality of output ports of each node where they are transmitted to adjacent nodes in the next cycle. In this way, the transport medium is
15 therefore bufferless. Note that details of the I/O circuitry are not shown in FIG. 1. Instead, only a plurality of output registers 105_{a-d} are shown, each of which is associated with a given output port and thus is the component from which a selected packet is sent from a node. Since an output port may also receive packets from the local agent coupled to the node aside from the respective plurality of input ports, a multiplexer (not shown in
20 FIG. 1) may be present in front of each output register 105.

The route taken by a packet is determined by the a priori defined and traffic independent sequence of connection states of the single nodes. The sequence may be chosen to be identical for all nodes or individually composed for each node. At any time, as determined by the node specific sequences, all nodes of the mesh are in one of a
25 plurality of predefined, traffic independent connection states mapping input ports to output ports. While the scope of the present invention is not so limited, in one embodiment, all nodes may follow the same sequence built from two connection states, namely a pass-through state and a turn state, each of which enables communication in a different dimension.

Referring now to FIG. 2, shown is a block diagram of a mesh interconnection
30 system 100' with all nodes in pass-through state. In this pass-through state, each node may be configured (e.g., by control of a multiplexer or other switch logic) to directly pass data received from a first neighboring node to a second neighboring node. In the

implementation shown in FIG. 2, these nodes may be the horizontal or vertical neighboring nodes. As used herein, a neighboring node is a node immediately adjacent to a node in a given dimension. Note in the implementation shown in FIG. 2, edge nodes are coupled to interconnects 120 that provide for direct feedback such that in a pass-through mode of operation, data output by the edge node is also directly received in that same edge node. For ease of illustration in FIG. 2, other interconnections such as internal routing within the nodes are not shown to not obscure the state shown in FIG. 2.

In this pass-through state, the nodes simply forward packets from opposite sides and thus the mesh is decomposed into a set of horizontally and vertically overlayed, unconnected, rings. In some embodiments, pass-through communications occur in both directions based on control of the system to thus enable potentially more efficient communications, although complexity may increase in this instance.

While the pass-through state as represented in FIG. 2 enables communication of data between neighboring agents of a single dimension, to enable communication between nodes of different dimensions embodiments may configure the nodes to provide for one or more turn states. Such turn states may be used to communicate data between neighboring cores of different dimensions. Note that regardless of the selected state, each node may input and output equal numbers of packets per cycle, as the transport mechanism is bufferless.

Referring now to FIG. 3, shown is a block diagram of a mesh interconnection system that is configured in a turn state. As seen in FIG. 3, system 100' is now configured such that each incoming packet is routed in a different direction than the pass-through direction. In this state each node routes an incoming packet left. This results in the incoming packets at each node being interchanged between horizontal and vertical rings (that were established in the pass-through state). In the illustration of FIG. 3, this turn state may be a left turn such that packets are communicated to the next neighboring node. Note that again with respect to the illustration of FIG. 3, edge nodes may be configured via a feedback interconnect 120 to receive directly a packet output from an output port of the node via an input port of the node. While shown with a left turn state in the illustration of FIG. 3, the scope of the present invention is not limited in this regard, and in other implementations, this turn state may be a right turn state.

In one embodiment, states are selected according to a cyclic schedule that can be completely independent from actual traffic on the medium. The schedule has to guarantee

that at each sending node a time slot for insertion exists such that eventually the inserted packet arrives at the intended destination. One such scheme meeting this constraint is a cyclic schedule of S clocks (where, $S = \max(n, m)$) comprising $S-1$ clock cycles with all nodes in the pass through state (FIG. 2) and 1 cycle with all nodes in the turn state (e.g., left, FIG. 3). More generally, a cyclic schedule can be of S clocks (where, $S = \max(n, m)$) and $S-x$ clock cycles are in the pass-through state (FIG. 2) and x cycles are in the turn state (e.g., left, FIG. 3). Such a cyclic schedule can, provably, establish at least one Manhattan route from each sending node to every other node. As used herein, the term "Manhattan route" is used to define a route with minimal travel distance between agents, without any roadblocks, detours, backwards tracking or so forth. Typically, additional viable paths not meeting the Manhattan property may exist. Other schedules are also feasible, like $S = 2 * \max(n, m)$ with again all nodes in pass-through configuration for $S-1$ cycles and a single cycle in turn configuration or so forth. While just one sequence applied to all nodes can suffice, multiple sets of sequences applied to respective subsets of nodes may be considered to improve certain characteristics of the medium. As an example, nodes may have individual sequences of connection states leading to multiple different connection states visible on a network layer for a given cycle. Further additional types of connection states may be considered like right turns or a combination of turn and pass-through configurations within one node. In general, additional states may increase the number of inputs to a multiplexer of the output ports.

It is to be understood that different topological variations exist, such as other feedback connection schemes at the edges of a node or reducing the mesh to a unidirectional transport mechanism. Referring now to FIG. 4, shown is a block diagram of an interconnection system in accordance with another embodiment of the present invention. As shown in FIG. 4, system 200 includes a plurality of nodes $210_{1,1} - 210_{n,m}$. Note in this implementation, only uni-directional interconnects 115 couple neighboring nodes together. In addition, note that instead of feedback paths for edge nodes, interconnects 120 are present to couple neighboring nodes of a common dimension. In other aspects, system 200 may be configured similarly to that of FIG. 1 described above.

With reference back to FIG. 1, at the source node, for each destination node a table of insertion time slots (slot table) exists for each output stage to ensure delivery of the inserted packet to the intended destination node. For the mesh operation of FIG. 1 and assuming a fixed schedule period of S clocks, the sequence of valid insertion options

becomes also cyclic with the period of S and thus may vary between 1 and S depending on the source/destination combination. The table is only dependent on the mesh topology and the cyclic schedule, and thus can be precomputed by different methods. In some embodiments, the table may be derived using an analytical expression or via exhaustive self-discovery. This table may in some embodiments be implemented in a non-volatile storage of each node. Also the table may be constructed on the fly by a respective function implemented by logic gates taking clock cycle count as well as source and destination information as input parameters. Or, the predetermined table can be dynamically loaded on power up to a volatile storage of the node. In particular embodiments, each output port of a node may include its own table, which is predetermined to reflect its particular connection within the mesh system. In one embodiment, each entry of the table may include a destination identifier and a corresponding cycle slot at which a packet for that destination is to be inserted via an output port. In some other embodiment a single table entry is provided per destination containing an S bit wide vector, where each bit n marks if insertion in the respective clock cycle n of the cyclic pattern is viable. In some embodiments, the table may further include heuristic data to aid output port logic in determining which packet is to be provided in a given slot. For example, for destinations where only one or two dedicated slots can reach the destination, a heuristic to select packets can be established for those precious slots over packets that have more options available for reaching their intended destinations. In this way, priority for a given packet may dynamically change based on packets pending for transmission. The global synchronization of the connection state sequences of the single node may be based on a master or global control signal provided to each of the multiplexers of the system. In some other embodiment synchronization may be achieved by special packets sent at network initialization.

When a packet is passed by agents associated with the respective node for packet insertion via the I/O circuitry, it can be intermediately stored in an ingress buffer until an appropriate time slot for transmission is reached, with no data allocated at the respective output port. That is, the system is configured such that due to its bufferless operation, incoming packets to an input port take precedence over locally generated packets, as there is no ability to maintain the received packets in a buffer or other storage structure.

To improve the maximum injection rate per port, multiple packets provided by the local agents may be processed in parallel, in some embodiments. As there typically exist

multiple combinations of time slot and output ports for a packet to enter the network (or multiple packets may enter at the same timeslot), various heuristic algorithms can be applied to improve network performance. In these algorithms, options leading to optimal Manhattan routes may be preferred. As one example, other viable but inefficient routes
5 (due to long paths in the network) may be removed (or deprioritized) from the tables. In cases of multiple packets contending for the same output port at the same timeslot, packets with fewer insertion options may be prioritized. Still further, such heuristic rules can be dynamically adapted. Meeting real-time requirements with this network is easy to achieve due to the deterministic behavior of the medium. To this end, deterministic slot pre-
10 allocation mechanisms may be present.

For each input port, incoming packets are checked for their destination address. This checking can be based on different destination encodings, e.g., an absolute value (such as a destination identifier) or a down-counter decremented at each passing hop. In case of a match, the incoming packet is locally stored in a per-input port instantiated
15 egress receiver queue and the respective slot on the mesh is marked as empty. To realize such operation, a receiver may have a local over-speed clock to allow reception of packets from all input ports in parallel at a given cycle. To eliminate such over-speed, in some other embodiment receivers may be restricted to only observe a deterministic subset of input ports at a given time, which is reflected accordingly in the respective entries of the
20 slot tables. In addition, the input port logic may mark the slot as empty so that a subsequent node can use the slot for data insertion (assuming the current node does not). In one embodiment, a single bit indicator may be used to indicate whether the slot is actually allocated to data or is an empty slot. Note that a receiver port may choose to not sink an incoming packet destined for it (e.g., for power, resource, or other reason). The
25 fabric allows for such “bounce” conditions and ensures that the packet is redelivered to its destination again after a finite period. That is, the fabric can support the ability to bring a packet back to its destination node in a deterministic manner. In different situations, this originally non-delivered packet may go back to its source, where information regarding its non-delivery can be used for flow control or other purposes.

30 Various micro-architectural improvement techniques may be applied to an I-ring interconnect in accordance with an embodiment of the present invention. For example, physical timing may be improved by allowing control signals that flag empty slots to run ahead on an identical operated mesh by a fixed cycle count to enable receiver pipelining.

Also destination decoding may be pipelined by moving it to the preceding node, as will be described further below.

Embodiments may also enable scaling of rings beyond two-dimensions. For example, a topology of $p \times n \times m$ meshes can be “stacked” together. In such
5 implementations, each $n \times m$ mesh may be located on one or more different layers of a semiconductor die. Alternately, multiple stacked die may be present, with each die including a $m \times n$ mesh in these implementations. At each node, a “switch” state may be provided to allow “packet swapping” between the 3 rings that intersect in a vertical dimension. Just as a 2-state minimal switching mechanism suffices for a 2-D mesh, a 3-
10 state minimal switching mechanism will suffice for a 3-D stack of meshes. Of course, additional states and cyclic schedules can be created to optimize various network parameters.

Referring now to FIG. 5, shown is a block diagram of a single input data path of a network node in accordance with an embodiment of the present invention. Multiple
15 instantiations of this path may be present in each node of a NoC. Specifically, in one embodiment, four such circuits may be present to form a node, each configured to receive incoming data from a first direction and to output data to a second direction opposite the first direction or a third direction, namely a turn direction.

In the implementation shown in FIG. 5, incoming information may be received in
20 an input port register 310. In some embodiments, based on a presence or validity indicator (active when no packet is to be output), control of a clock for an input register can occur to reduce power consumption by disabling the register when no valid data is incoming. The information received may include payload data and status/control information. Status/control information may be used to indicate whether incoming data is valid as well
25 as other control information such as a destination identifier for the corresponding packet, quality of service (QoS) related information such as a reservation bit keeping the slot free for a given connection and so forth. As seen, if incoming payload data is for one of the agents associated with the node, e.g., as indicated by destination indicator present in the control information, the payload data is provided to an egress queue 320. From there, the
30 data may be provided to one of the agents associated with the node such as a core, a memory controller or other such processing unit.

Otherwise the data and control information both pass to an output data selector 330, which generally is used to select either the data received via input port register 310,

or data generated in the node, received from an agent of the node. Thus as seen in FIG. 5, an ingress queue 350 is coupled to receive data of the agents desired to be sent to another node. In turn, the ingress queue is coupled to an ingress controller 340. In general, ingress controller 340 may select one of multiple data packets for insertion into an open slot, based on information present in a slot table 345. As discussed above, slot table 345 may be accessed by a destination identifier for the corresponding packet and may be used to indicate an appropriate slot for insertion of the packet. As seen, additional information used by ingress controller 340 to select a data packet for insertion is a slot identifier, received from a static global cyclic scheduler 365, further details of which are discussed below. Accordingly, based on control of ingress controller 340, output data selector 330 will output either the payload data received from input port register 310, or data received from ingress controller 340 (namely data originating in the node itself).

Still referring to FIG. 5, output data selector 330 is coupled to an output port selector 360 which is controlled by a static global cyclic scheduler 365 representing the sequence of connection states of the node. Specifically, in an implementation in which two states are present, namely a pass-through state and a turn state, for each time slot the scheduler will generate a control signal to cause output port selector to forward the control and payload data to one of a first or second output port 370 and 375. In one embodiment, output port 370 may output the information for a pass through slot, while an output port 375 may output the information for a turn slot. While shown with this particular implementation in the embodiment of FIG. 5, understand the scope of the present invention is not limited in this regard.

For example, as discussed above, additional states may be present and accordingly, more output ports may be present, each configured to output information to a different neighboring node to which the present node is coupled. Further, while in many implementations incoming traffic takes priority such that when it is present, output logic only passes the incoming packet and does not insert a new packet in the next slot, embodiments are not so limited. That is, in other embodiments an override configuration may be present such that an incoming packet can be replaced in favor of insertion of a new packet (presumably of a higher traffic class). For example, the node can sink the packet (e.g., to a cache) and based on a flow control scheme later inject the packet, or the original sender of the replaced packet can be informed, e.g., by a no acknowledgment message, so that the packet can be resent.

Referring now to FIG. 6, shown is a flow diagram of a method for communicating packets according to a static global cyclic schedule in accordance with an embodiment of the present invention. As shown in FIG. 6, method 400 may begin by determining whether a valid incoming packet has been received (diamond 410). This determination
5 may be made, for example, in an input port register that receives payload data and various control information and from this, can determine whether the incoming payload data is valid. If so, it may further be determined whether the packet is destined for an agent associated with the current node (diamond 415). This determination, which may also be made in the input port register, may be based on a destination identifier associated with the
10 packet. If so, control passes to block 420, where the packet may be provided to an egress queue. In addition, status information associated with this slot may be updated. For example, a valid indicator may be reset to indicate that no valid data is present in this slot and accordingly, this slot is available for insertion of data.

Still referring to FIG. 6, if instead it is determined that the incoming packet is not
15 destined for the current node, control passes from diamond 415 to block 430 where the packet may be passed to a selected output port according to a global schedule (block 430). This global schedule may dictate whether the packet is to be passed out of the node in a first direction or a second direction. Based on this global schedule then, control passes to block 435 where the packet may be output on an interconnect.

Still with reference to FIG. 6, if at diamond 410 it is determined that a valid packet
20 is not received (or where a received packet is destined for the associated agent), control passes to diamond 450, where it may be determined whether one or more packets are present in an ingress queue. If not, control passes to block 470, where a void packet may be passed to the output port. This void packet, namely a null data payload and control
25 information to indicate an invalid packet, may thus be provided to another node, which can then insert a packet of its own in that slot.

If at diamond 450 it is determined that one or more packets are present, control passes to diamond 460 where it may be determined whether a global schedule slot matches one of the packets. If not, control passes to block 470, discussed above. If there is a
30 match, control instead passes to block 480 where a packet may be provided from an ingress queue to a selected output port according to the global schedule. Note that the determination as to which of multiple pending packets is to be transmitted can be based on the information in the table, as well as an understanding of the current slot of the global

schedule. For example, the table information may associate slot identifiers with destinations. Various heuristics also may be taken into account in selecting one of multiple packets for output. In this way, priority information can be part of the basis of the determination. Still further, heuristics regarding the availability of multiple slots (or not) for a given packet to be inserted into the network can be used as part of the basis for the determination. Accordingly, as between two packets, a packet that has fewer slots into which it can be inserted may be selected ahead of a packet having greater available slots into which it can be inserted.

Thus in basic implementations the determination of which of multiple packets is to be selected can be based on slot identifier and destination identifier, while in more advanced implementations additional details regarding various heuristics can be taken into account. Then as seen, control passes to block 435 for output of this packet on the interconnect. While shown with this particular implementation for FIG. 6, understand the scope of the present invention is not so limited.

For example, in some embodiments a pipelining implementation may be present in which the determination of whether a packet is destined for a given node is performed in a neighboring (i.e., a previous) node. That is, this determination may be made in the previous node and can be forwarded, e.g., via a one bit wire either in a look ahead manner or as part of the normal control information. In either event, using a pipelining implementation, it can be determined ahead of time that an ingress controller can start selecting a packet for insertion into an available slot. One mechanism to make the determination in the previous node as to whether a packet is destined for the next node may be via a count down mechanism. For example, rather than a destination identifier that corresponds to an intended node, a destination identifier may take the form of a count of nodes through which a packet is to pass between source and destination. Thus in each node through which the packet passes, this count down value is decremented. Accordingly, when this packet is received in the previous node, the count value may be decremented, e.g., to zero or one, indicating that the packet is intended for the next node. Accordingly, in a look ahead manner, a one bit wire may indicate that the next received packet by the intended node is for that node. To this end, some implementations may include separate pathways to pass this destination information as separate control information that can be provided through the IO circuitry such that these decisions can be made in advance.

Embodiments may be implemented in code and may be stored on a storage medium having stored thereon instructions which can be used to program a system to perform the instructions. The storage medium may include, but is not limited to, any type of disk including floppy disks, optical disks, optical disks, solid state drives (SSDs),
5 compact disk read-only memories (CD-ROMs), compact disk rewritables (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories
10 (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications
15 and variations as fall within the true spirit and scope of this present invention.

Claims

What is claimed is:

1. A system comprising:

5 a nxm mesh system including a plurality of nodes, the plurality of nodes arranged in a first dimension and a second dimension; and

a plurality of interconnects each to couple a pair of the plurality of nodes, wherein the mesh system is configured in a first traffic independent connection state for a first cycle of a traffic schedule period and in a second traffic independent connection state for a second cycle of the traffic schedule period.

10 2. The system of claim 1, wherein the first traffic independent connection state comprises a pass-through state in which packets are communicated between neighboring nodes in the first dimension.

3. The system of claim 2, wherein the second traffic independent connection state comprises a turn state in which packets are communicated between neighboring
15 nodes in the second dimension.

4. The system of claim 1, wherein the mesh system comprises a bufferless transport medium.

5. The system of claim 1, wherein the traffic schedule period is a fixed cyclic schedule of S clocks, where S is max (n,m) and the mesh system is configured in the first
20 traffic independent connection state for S-x cycles of the fixed cyclic schedule and in the second traffic independent connection state for x cycles of the fixed cyclic schedule.

6. The system of claim 1, wherein each of the plurality of nodes includes a plurality of input ports and a plurality of output ports, each of the plurality of input ports coupled to at least one of the plurality of interconnects.

25 7. The system of claim 6, wherein each of the plurality of input ports is to direct a packet destined for the corresponding node to an egress queue coupled to an agent associated with the corresponding node.

8. The system of claim 6, wherein each of the plurality of output ports is to select a packet received from the corresponding node for output from the output port based
30 on an entry in a slot table that stores a destination identifier corresponding to a destination of the packet and a slot identifier.

9. The system of claim 8, wherein each of the plurality of output ports is to select a packet further based on control of a static global cyclic scheduler.

10. The system of claim 8, wherein the output port includes a first output port to output the packet to a first interconnect coupled between a first node and a first neighboring node, and a second output port to output the packet to a second interconnect coupled between the first node and a second neighboring node.

5 11. A method comprising:

determining whether a packet received in a node is destined for an agent associated with the node and if so, providing the packet to an egress queue of the node;

if the received packet is destined for the agent, determining whether one or more packets are present in an ingress queue of the node and if so, providing a selected one of
10 the one or more packets from the ingress queue to a first output port or a second output port of the node according to a global schedule, the global schedule independent of traffic flow; and

otherwise providing the received packet to the first output port or the second output port according to the global schedule.

15 12. The method of claim 11, further comprising providing a void packet to the selected first or second output port, if none of the one or more packets present in the ingress queue match the global schedule for a current slot.

20 13. The method of claim 11, further comprising receiving an indication that the received packet is destined for the agent from a neighboring node prior to receipt of the received packet in the node.

14. The method of claim 11, further comprising receiving an indication that the received packet is destined for the agent in a slot of the global schedule prior to receipt of the received packet in the node.

25 15. The method of claim 11, further comprising forwarding received packets from the first output port for S-x cycles of the global schedule and from the second output port for x cycles of the global schedule.

16. An apparatus comprising:

an input port register of a node to receive payload data and corresponding control information from a first neighboring node;

30 an egress queue coupled to the input port register to receive the payload data and to forward the payload data to an agent associated with the node;

an output data selector to select payload data from the input port register or from the agent for output from the node; and

an output port selector coupled to the output data selector to output the selected payload data and corresponding control information to a first output port coupled to a second neighboring node during a first cycle of a traffic schedule period and to a second output port coupled to a third neighboring node during a second cycle of the traffic schedule period, wherein the traffic schedule period is independent of traffic between the nodes.


17. The apparatus of claim 16, wherein the first cycle corresponds to a pass-through state in which packets are communicated between the node and the second neighboring node in a first dimension, and the second cycle corresponds to a turn state in which packets are communicated between the node and the third neighboring node in a second dimension.

18. The apparatus of claim 17, wherein a third cycle corresponds to another turn state in which packets are communicated between the node and a fourth neighboring node in a third dimension.

19. The apparatus of claim 16, wherein the apparatus comprises a first input/output (I/O) circuit of the node, and wherein the node further includes a plurality of I/O circuits each including an input port register, an egress queue, an output data selector and an output port selector.

20. The apparatus of claim 16, wherein the apparatus comprises a processor including a plurality of stacked semiconductor dies, wherein the node is on a first semiconductor die, and further comprising a fourth neighboring node configured on a second semiconductor die stacked on the first semiconductor die.

100






FIG. 2




FIG. 3




FIG. 4


300

FIG. 5


400

FIG. 6