

OPTIMIZED SCALABLE NETWORK SWITCH

CROSS-REFERENCE

The present invention claims the benefit of commonly-owned, co-pending United States Provisional Patent Application Serial Number 60/271,124 filed February 24, 2001 entitled MASSIVELY PARALLEL SUPERCOMPUTER, the whole contents
5 and disclosure of which is expressly incorporated by reference herein as if fully set forth herein. This patent application is additionally related to the following commonly-owned, co-pending United States Patent Applications filed on even date herewith, the entire contents and disclosure of each of which is expressly incorporated by reference herein as if fully set forth herein. U.S. patent application
10 Serial No. (YOR920020027US1, YOR920020044US1 (15270)), for “Class Networking Routing”; U.S. patent application Serial No. (YOR920020028US1 (15271)), for “A Global Tree Network for Computing Structures”; U.S. patent application Serial No. (YOR920020029US1 (15272)), for ‘Global Interrupt and Barrier Networks’; U.S. patent application Serial No. (YOR920020030US1
15 (15273)), for ‘Optimized Scalable Network Switch’; U.S. patent application Serial No. (YOR920020031US1, YOR920020032US1 (15258)), for “Arithmetic Functions in Torus and Tree Networks’; U.S. patent application Serial No. (YOR920020033US1, YOR920020034US1 (15259)), for ‘Data Capture Technique for High Speed Signaling’; U.S. patent application Serial No. (YOR920020035US1
20 (15260)), for ‘Managing Coherence Via Put/Get Windows’; U.S. patent application Serial No. (YOR920020036US1, YOR920020037US1 (15261)), for “Low Latency Memory Access And Synchronization”; U.S. patent application Serial No. (YOR920020038US1 (15276)), for ‘Twin-Tailed Fail-Over for Fileservers Maintaining Full Performance in the Presence of Failure’; U.S. patent application
25 Serial No. (YOR920020039US1 (15277)), for “Fault Isolation Through No-Overhead Link Level Checksums’; U.S. patent application Serial No. (YOR920020040US1 (15278)), for “Ethernet Addressing Via Physical Location for Massively Parallel Systems’; U.S. patent application Serial No.

(YOR920020041US1 (15274)), for "Fault Tolerance in a Supercomputer Through Dynamic Repartitioning"; U.S. patent application Serial No. (YOR920020042US1 (15279)), for "Checkpointing Filesystem"; U.S. patent application Serial No. (YOR920020043US1 (15262)), for "Efficient Implementation of Multidimensional Fast Fourier Transform on a Distributed-Memory Parallel Multi-Node Computer"; U.S. patent application Serial No. (YOR9-20010211US2 (15275)), for "A Novel Massively Parallel Supercomputer"; and U.S. patent application Serial No. (YOR920020045US1 (15263)), for "Smart Fan Modules and System".

10 **FIELD OF THE INVENTION**

The present invention broadly relates to a method of switching data traffic in a distributed-memory message-passing parallel computer. More particularly, it relates to a method that efficiently delivers error-free messages over an interconnection network in a massively parallel system.

15

BACKGROUND OF THE INVENTION

A well known topology for massively parallel computer data networking systems, is the 3D torus. a 3D torus is generally a cubic grid of compute nodes that has a ring network at every level. Supercomputing massively parallel systems, such as the system described in the Provisional Application Serial Number 60/271,124, use The 3D torus topology to provide the minimal path route, i.e. the shortest path for communications between hundreds or thousands of nodes. One problem with this topology in a massively parallel system is the inefficient delivery of messages over an interconnection network, particularly when Ethernet or Asynchronous Transfer Mode (ATM) switches are used.

More specifically for example, Ethernet or ATM switches do not generally provide low latency, high throughput, and error free delivery of packets, since these switches typically lose packets if there is not enough buffer space, i.e., holding areas for input and output processing, to hold the packet. Additionally, the problem of contention, i.e., a conflict that arises when two or more requests are made concurrently for a

resource that cannot be shared, such as a communication link, must be overcome if the switching network is to be scalable to the size of tens of thousands of nodes.

Typically, contention issues have been dealt with by employing some sort of arbitration algorithm which mediates which transmitters on a network can transmit packets subsequent to a packet collision detection. Unfortunately, the related art has not addressed the need for a collision detection/arbitration method which is ultra-scalable, thus suitable for massively parallel systems. Additionally, current routing techniques are not suitable for scaling up to massively parallel systems because the routers typically have tables that must be maintained. Overhead for table maintenance becomes unduly burdensome as the number of nodes reaches the tens of thousands.

As stated above, the three-dimensional (3D) torus topology is known. For example, the Cray T3E uses this 3D torus topology. However, the Cray uses routing tables stored in each switch element, an approach that does not scale well to tens of thousands of nodes. Other known technologies are the "Bubble" escape virtual channels (VC's), (Puente et al., "Adaptive Bubble Router: A Design to Balance Latency and Throughput in Networks for Parallel Computers," In Proceedings of the International Conference on Parallel Processing, ICPP '99, Sept. 1999), which provide fully dynamic routing that does not require routing tables.

Another known technique is the use of multiple virtual channels to reduce "head-of-line" blocking, as employed in the SP2 and the Cray computers. The use of a two-stage arbitration approach has been taught by the MIT Reliable Router (William J. Dally, Larry R. Dennison, David Harris, Kinhong Kan, and Thucydidis Xanthopoulos, "Architecture and Implementation of the Reliable Router," In Proceedings of HOT Interconnects II, pp. 122-133, Aug. 1994).

Another related art technology uses virtual cut-through routing in an attempt to optimize throughput and latency. See P. Kermani and L. Kleinrock entitled "Virtual Cut-Through: A New Computer Communication Switching Technique", Computer Networks, Vol. 3, pp267-286, 1979, incorporated herein by reference.

However, the related art references do not adequately solve the problem of packet contention and queueing delays along a selected packet direction of travel and virtual channel, particularly when a switch is scaled up to handle tens of thousands
5 of nodes.

In a related disclosure, U.S. Provisional Application Serial Number 60/271,124 entitled "A Novel Massively Parallel Supercomputer", therein is described a semiconductor device with two electronic processors within each node of the multi-
10 computer. Within the supercomputer, there is a plurality of high-speed internal networks, and an external network employing Ethernet. These networks are expected to service over 64,000 nodes.

While there is no known prior art that attempts to scale a network switch to tens of
15 thousands of nodes for fast, error-free operation, there remains the need for a scalable arbitration method that enables error-free, "low latency, high bandwidth (throughput) data communications, to enhance the message-passing capability for a massively parallel system.

20 **SUMMARY OF THE INVENTION**

Therefore, it is an object of the present invention to provide a method for arbitrating among packet transmitters, wherein latency is decreased, throughput is increased, and processing power has increased in a massively parallel system.

25 A further object of the present invention is to provide a method for preventing packet loss in the massively parallel system.

Yet another object of the current invention is to provide a method of data communications that avoids deadlocking users of the communications channels
30 employed.

A further object of the current invention is to provide a method for increasing the probability of accurate packet reception.

Another object of the current invention is to provide the switch technology on the
5 same chip as the processors, to facilitate scalability.

An even further object of the current invention is to provide a method for routing messages without using a routing table within the routers.

10 These and other objects and advantages may be obtained in the present invention by providing a multilevel arbitration method for decreasing latency, and increasing accurate packet reception in a massively parallel system.

Specifically, there is provided a packet routing method comprising; routing packets
15 to their destinations by inputting packets from at least one of 2^m directions, (where m is an integer representing the dimensionality of the network), via at least one of a corresponding number of input links into a corresponding at least one of a plurality of virtual channel FIFO buffers; generating at least one of a 2^m plurality of output bit vectors containing information derived from downstream nodes; establishing a
20 general arbitration policy wherein arbitration decisions are based on the link and downstream buffer fullness information contained in the output status bit vector; outputting, according to the general arbitration policy, the packets via multiple transfer paths to corresponding output links, wherein the packets are advanced towards their destination nodes.

25

BRIEF DESCRIPTION OF THE DRAWING

The present invention will now be described in more detail by referring to the drawings that accompany the present application. It is noted that in the accompanying drawings like reference numerals are used for describing like and
30 corresponding elements thereof.

Fig. 1 shows the switch architecture according to a preferred implementation of the present invention;

Fig. 2 shows an output status bit vector according to a preferred implementation of
5 the present invention;

Fig. 3 shows an example of an input status bit vector according to a preferred implementation of the present invention;

10 Fig. 4 shows a "possible destination" bit vector according to a preferred implementation of the present invention;

Fig. 5 shows a packet header with hint bits according to a preferred implementation of the present invention; and

15

Fig. 6 shows an embodiment of the current invention, i.e., network switch, located on the same chip as the processor nodes.

20 This invention applies to network communication, i.e., message-passing among nodes in a three dimensional torus network of a massively parallel computing system.

At the outset it should be noted that reference to "FIFO" (First in First Out) is used hereinafter interchangeably with the term "FIFO buffer" as are the plural "FIFO's" used interchangeably with the term "FIFO buffers". References to bit settings and
25 bit counting sequences are exemplary only, and other bit settings, e.g., negative logic, and counting sequences, e.g., negative counting sequence to indicate positive increase in parameter or measurement are within the scope of this invention.

According to a preferred implementation of this invention, the switch architecture,
30 as shown in Fig. 1, which depicts an individual node in a network of processors such as a three dimensional torus network, may be implemented for all such nodes in the network. Additionally, referring to Fig. 6, it is shown that in a preferred

implementation of the current invention, switch 610 may also be integrated onto the same chip 605 as the processors 615. The network delivers variable sized packets of data which, in a preferred implementation may be from one to eight 32 byte sized units, called "chunks". Each packet contains a descriptive packet header 510, as shown in Figure 5. New, i.e., local packets 104, originating from a local node 105, enter the network via processor injection FIFO's 130 and are removed from the network upon reaching their destination by removing them from the processor reception FIFO's 140. In a preferred implementation of the current invention, there is at least one processor reception FIFO 140 dedicated to each input unit, such as 110, 117. Additionally, at least one processor reception FIFO 140 may be dedicated to high priority traffic.

Each node includes input links, one for each neighbor in the torus. Fig. 1, shows for example, the x+ input link 101 and the z- input link 102. Remaining input links, x-, y+, y-, z+ (not shown) represent an example of remaining directions of travel available to a global packet, i.e., a packet arriving from one of the input links, in a system such as the 3D torus system of the present invention. Corresponding to each input link is a local input arbiter such as the x+ input arbiter 116 which is included in x+ input unit 110. The x+ and z- input units 110, 117 are shown as exemplary structure.

Each node also includes output links such as link to the x+ direction 161, and link to the z- direction 162. As implemented in a preferred aspect of the current invention, there exists at least one output link for each neighbor in the torus. The output links are sourced by output units, such as the shown x+ and z- output units 151, 152. Within each output unit is an output arbiter such as the shown z- output arbiter 155.

The output units 151, 152 preferably contain the following sub-units:

- a.) A token unit counter 153 that counts how many chunks, i.e., tokens, are available to be stored at the downstream node. There is preferably one token unit counter for each direction and VC. When a packet wins arbitration, the token units coordinate with

one another so that tokens are returned to the appropriate VC on the upstream node. This is accomplished by sending special, short, token messages.

- 5 b.) An output status bit vector generator 154 that generates output status bits as shown in Figure 2.
- c.) An output arbiter 155 that selects which packet, among all possible requesting packets, shall win control of the corresponding link.
- 10 An input unit, such as shown in input unit 110 comprises a cut-through, i.e., bypass path 120, and waiting buffers, i.e., FIFO's, one for each VC, for storing packets which are blocked, i.e., those packets which cannot move through the bypass path 120. Each FIFO can hold some maximum number of chunks. One or more "dynamic" VC's are implemented, such as dynamic VC's 111 and 112. The
- 15 implementation of "escape" VC 113 for preventing deadlocks, and a high priority VC 114 is a known technique, as described in (Puente et al., "Adaptive Bubble Router: A Design to Balance Latency and Throughput in Networks for Parallel Computers," In. Proceedings of the International Conference on Parallel Processing, ICPP '99, Sept. 1999). Implementation of by-pass path 120 allows packets to be
- 20 routed through the node without being stored in a buffer. However, if a packet in the by-pass path is blocked for any reason, it is transferred into the appropriate FIFO buffer (i.e., the buffer corresponding to the VC associated with the packet). To prevent priority packets from over-taking one another, a priority packet in the bypass path is considered blocked if the corresponding priority VC FIFO is non-empty.
- 25 Similarly, a normal priority packet that is not dynamically routable is considered blocked if the corresponding escape VC FIFO is non-empty. Only the packet header needs to be received in the FIFO before being ready to arbitrate (provided the packet is at the head of the FIFO). An input arbiter, such as input arbiter 116 includes an input status bit vector generator 115, which generates status bits as shown in Figure
- 30 3 and which will be described later.

The input arbiter 116 provides dynamic routing by selecting which packet in the input arbiter gets priority, and which output direction and VC to request (or which processor reception FIFO to request if the packet has reached its destination).

According to a preferred aspect of the present invention, once a packet wins input and output arbitration, i.e., gets priority for transfer to the output unit or processor reception FIFO's, it is transferred to the output link, or reception FIFO via one of a plurality of transfer paths 121. Unlike prior art switches in torus networks which have only a single such path, the plurality of transfer paths 121 enable a packet to be transferred even though one or more of them are being used. According to one aspect of the current invention, some of the transfer paths may be reserved exclusively for the reception FIFO's.

Additionally, a preferred aspect of the present invention includes a by-pass path 120. Associated with each packet at the head of FIFO, or in the by-pass path 120, is a bit vector of possible destination VC's and directions in which the packet may move. An example of such a bit vector 410-415 is shown in Fig. 4.

Output Status Bits

Each output arbiter is responsible for creating an output status bit vector that describes the availability of its link and the downstream buffer space (tokens) associated with this direction. Figure 2 depicts a sample output status bit vector at an instant in time for the x+ 210 to z- 215 links. Each output status bit vector preferably contains

- a.) A link status bit 201 indication whether the link is available (1) or not (0). As the link may also be used to exchange control messages related to the availability of tokens and/or acknowledgement of successfully received packet, such control messages receive absolute priority over data packets. Thus if such a control message is waiting to use the link, or is actually using the link, the link status bit is set to 0 and the control message is transmitted as soon as the physical link becomes free.

- b.) Status bits 202, 203 for each dynamic VC. The left-most bit indicates whether or not there are sufficient tokens downstream to store the largest sized packet. The other 2 bits are determined as follows. The available downstream tokens are broken up into 4 ranges defined by thresholds, e.g., $M \geq T1 > T2 > T3 > 0$ where M is the maximum number of tokens. If the number of tokens available is t, then the bit vector is: 00 if $M \geq t > T1$, 01 if $T1 \geq t > T2$, 10 if $T2 \geq t > T3$, and 11 if $T3 \geq t \geq 0$. (Thus, the bit vector is smaller if more tokens are available. It should be understood to anyone skilled in the art that there exists a generalization to $k > 2$ bits defining 2^k ranges.)
- c.) Status bits for the escape 204 and priority 205 VC. These bits follow the “bubble insertion” rules.
- 15 For the escape 204 and priority 205 status fields, the left bit is 1 if a new packet can be injected into the VC (this requires at least 2 full packet’s worth of tokens) and the right bit is 1 if a packet already on the VC can continue (at least 1 full packet’s worth of tokens). (To prevent deadlocks, the “size” of all packets entering or on these VC’s is considered to be the maximum packet size.)
- 20 For example, as shown in Fig. 2, in the x+ direction 210, the link is free, tokens are available at level 01 on VC 1, no tokens are available on VC 2, the escape VC does not have space to insert a new packet but does have space to forward a packet already “in the bubble”, and the high priority VC can both insert and forward a
- 25 packet.

Input Status Bits

- Each input unit is responsible for creating an input status bit vector that describes the input fullness of its buffers for each VC associated with the direction associated with the input unit. Figure 3 depicts sample status input bit vectors at an instant in time for the x+ 310 to z- 315 links. Each bit vector contains three bits describing each VC 301 - 304. The left-most bit 301 is 1 if a packet is ready to arbitrate and 0

otherwise. The two right-most bits 303, 304 describe token fullness regions, i.e., if the FIFO contains b number of chunks, then the bits are 00 if $0 \leq b < T_1$, 01 if $T_1 \leq b < T_2$, 10 if $T_2 \leq b < T_3$, and 11 if $T_3 \leq b$ for thresholds T_1 , T_2 and T_3 . It should be understood to anyone skilled in the art that ranges may be used to describe the number of packets, rather than chunks.

Each injection FIFO also has a similarly defined input status bit vector 316 as shown in Fig. 3 which illustrates input status bit vectors 317-321, indicating there are 5 injection FIFO's.

10 *Possible Bit Vector*

For each packet at the head of a FIFO or on the by-pass path, a bit vector is created that indicates which directions and VC's a packet may move along, i.e., for a given packet, which directions are still routing goals for each virtual channel. Figure 4 depicts the possible bit vector for a non-priority packet that can move in either the $x+$ or $y-$ directions. A portion of the bit vector corresponds to each direction 410 - 415. In this example, the packet may move on either dynamic VC in both the $x+$ and $y-$ directions. Since the packet is non-priority, the entries for the priority VC's are all 0. In this example, the escape direction is chosen according to the standard dimension order, i.e., first x , then y , then z , so the escape VC is in the $x+$ direction.

20 If the possible bit vector is all zeros, then the packet has reached its destination.

Packet Header

The format for the packet header 510 is shown in Figure 5. In a preferred embodiment, the packet header 510 will be relatively short in order to minimize overhead, e.g., 8 bytes. In the header 510, the first field (e.g., 8 bits) indicates the packet type 512, e.g., normal or priority, and may include a bit indicating whether or not the packet may be routed dynamically, or along a fixed route. In a related patent application 15270, it is described how the packet type can encode "class routing" operations which can also be supported in the current invention with obvious minor modifications. The VC field 514 (e.g., 2 bits) indicates which VC the packet is on and the size field 516 (e.g., 3 bits) indicates the number of chunks in the packet. The use of such headers is standard, however, a preferred implementation of the current

invention includes "hint bits" 520, 522, and 524, which reduce latency by permitting the arbiter to begin selecting a direction earlier than would otherwise be possible. In a preferred implementation of the current invention, these hint bits are 2 bits/dimension and encode which directions are feasible for a packet to move in, e.g., if the x hint bits 518 are "00" then no more packet movement in the x direction is necessary, if the x hint bits 518 are "01" then the packet may move in the x+ direction, and if the x hint bits 518 are "10" then the packet may move in the x- direction. The hint bits must precede the destination fields 524, 526, and 528, and may, alternatively, precede the VC and size fields 514, 516. On a large parallel machine, e.g., a 64x32x32 node torus, the destination fields must be at least 6 bits each. Without the hint bits, 18 bits of destination address would have to be received before arbitration can begin. Thus, since the hint bits precede the destination bits, the hint bits advantageously allow the starting of the arbitration process before the entire packet header is read, i.e., before the destination bits are read, thereby reducing latency. As a packet is routed through the network, the hint bits need to be updated. For minimal hop dynamic routing, for example, the x hint bits, need only be set to 00 as the packet reaches the x coordinate of its destination. A preferred implementation of the current invention accomplishes this by including a register associated with each output link. Within the register is provided the relevant, e.g., the x coordinate of the downstream node. When the packet wins arbitration, the register can be compared with the destination's relevant coordinate; the hint bits are set to 00 if the value in the register and the destination's relevant coordinate are equal. Additionally, as is known in the art, for reliable messaging, the header may include a sequence number, as for example, in the known High Performance Parallel Interface (HIPPI) protocol, and a CRC to ensure that a header is correctly received.

Input/Output Arbitration

In a preferred implementation of the current invention, the arbitration process comprises two phases. First, each input unit selects which packet from among its ready-to-move packets (those at the head of a FIFO or on the by-pass path) will attempt to obtain a link (or reception FIFO). These selected packets are called "input winners". In addition to selecting the input winners, i.e., candidates, a

direction for each candidate and VC for each candidate are selected for the next step in the arbitration. Both input and output arbitration is multilevel, i.e., separate arbitration processing is provided for packets arriving from each node direction to increase throughput of the packets in the system. Packet contention is reduced by the
5 general arbitration policy as described by the rules employed in a preferred implementation of this invention.

Input Arbitration Phase:

A method provided by a preferred implementation of the current invention to select the input winners includes steps to determine which packets not at their destination
10 are completely blocked, i.e., which packets have no links and VC's available in directions the packet needs to move, or no transfer path is available. These steps also determine which candidate packets are unblocked. In one aspect of the current invention the above steps are determined by performing an appropriate logical AND operation of the possible bit vectors 410-415 with the output status bit vectors 210-
15 215. Similarly packets that have reached their destination may be determined to be blocked if there is no space in its reception FIFO or an unblocked candidate vector. Each unblocked candidate is assigned an input priority as follows. Two types of arbitration cycles, a random cycle or a longest queue cycle are provided by a preferred implementation of the current invention. A predefined arbitration cycle
20 criteria of the switch determines what fraction of the arbitration cycles are random. On a random cycle, each candidate packet in a non-priority VC buffer is assigned a random (or based on a round robin algorithm) 2 bit vector. On a longest queue cycle these candidate packets are assigned a priority equal to the right two input status bits. The priorities so generated are used in an overall priority scheme as follows
25 (from highest priority to lowest):

- a.) The unblocked candidate packet at the head of the priority VC.
- b.) The unblocked candidate packet in the by-pass path, if it is a high
30 priority packet and there are no packets in the high priority VC buffer.

- c.) The unblocked candidate packet in a non-priority VC buffer having the highest 2 bit priority vector (ties broken by an arbitrary mechanism, such as random or round robin)
- d.) The unblocked candidate packet in the by-pass path, provided it is a non-priority packet.

For non-priority packets, this algorithm tends to select packets at the head of the most full FIFO's. The winning candidate packet then selects its candidate direction and VC as follows.

- a.) If the packet is high priority it selects the direction and the priority VC for which the priority VC bit of the possible bit vector is set to 1.
- b.) If the packet is non-priority but not dynamically routable, it selects the direction and escape VC for which the escape VC bit of the possible bit vector is set to "1".
- c.) If the packet is non-priority and dynamically routable, it selects, from among those links and dynamic VC's that both are available and possible, the direction and dynamic VC having the smallest right 2 bits in the output status bit vector.
- d.) If no such link and dynamic VC exist, it selects the direction and escape VC for which the escape VC bit of the possible bit vector is set to "1".

The input winners then present their candidate direction and VC to the corresponding output arbiter if the destination is not the current node, or the reception FIFO if the destination is the current node. Each output arbiter /reception FIFO then selects a packet for output from among the presented input winners.

Output Arbitration Phase:

- Each output arbiter is presented with a set of candidate input winners. In addition, there may be candidate packets from the injection FIFO's. Output arbitration also has two kinds of cycles: one in which packets from the injection FIFO's have low

- priority, and one in which packets from the injection FIFO's have high priority. A predefined output phase arbitration cycle criteria, e.g., a parameter of the switch, determines what fraction of cycles are of each type. If injection FIFO's have low priority on this cycle, the output arbiter selects its winner from among all its
- 5 candidates as follows (highest to lowest priority). Ties are preferably broken randomly.
- a.) A high priority winning packet from an input unit that has the highest priority as indicated by the 2-bit status vectors.
 - b.) A high priority packet from an injection queue with the highest 2-
 - 10 bit status vector.
 - c.) A normal priority winning packet from an input unit that has the highest 2-bit status vector.
 - d.) A normal priority from an injection FIFO that has the highest 2-bit status vector.

15

If injection packets have high priority on this cycle, the output arbiter preferably selects its winner from among all candidates as follows (highest to lowest priority):

- a.) A high priority packet from an injection queue with the highest 2-
- 20 bit status vector.
- b.) A high priority winning packet from an input unit that has the highest priority as indicated by the 2-bit status vectors.
- c.) A normal priority from an injection FIFO that has the highest 2-bit status vector.
- 25 d.) A normal priority winning packet from an input unit that has the highest 2-bit status vector.

Injection Arbitration

- Candidate packets at the head of injection FIFO's select their candidate direction
- 30 and VC in the same manner as the winning input candidate packets. These packets also carry a 2 bit status vector, either randomly generated or using the 2 bits to indicate the length of the injection FIFO.

Reception FIFO Arbitration

There may be one or more packets that need to be received at a switch node at the same time. To lessen contention, a separate reception FIFO is assigned to each
5 incoming direction. In addition, there may be one or more FIFO's per direction. If there are multiple processors in each node then a preferable embodiment of the current invention allows for multiple FIFO's for each direction to enable each processor to have its own reception FIFO per direction. Bits in the packet header can be used to indicate which of these FIFO's, for a given direction, is to be used. Also,
10 there may be one or more FIFO's set aside solely for use by high priority packets. If there is contention for these FIFO's, priority may be determined using the 2-bit status vectors.

Now that the invention has been described by way of a preferred embodiment,
15 various modifications and improvements will occur to those of skill in the art. For example, it should be understood to anyone skilled in the art that alternate embodiments of this invention may include a single dimension, or may include m multi-dimensions, where m is any integer. Thus, it should be understood that the preferred embodiment is provided as an example and not as a limitation. The scope
20 of the invention is defined by the appended claims.

WHAT WE CLAIMS IS:

- 1 1. In a massively parallel computing system comprising a plurality of nodes
2 configured in m multi-dimensions, each node including a computing device, a
3 method for routing packets towards their destination nodes comprising the steps of:
4 a.) inputting said packets from at least one of $2m$ directions via at least one
5 of a corresponding number of input links into a corresponding at least
6 one of a plurality of virtual channels;
7 b.) inputting local packets from a local of said nodes to at least one of a
8 plurality of processor injection FIFO's;
9 c.) generating a $2m$ plurality of output status bit vectors, each of said output
10 status bit vectors describing the availability of its link, the availability of
11 downstream virtual channels, and a downstream buffer space indication
12 for all virtual channels associated with its direction;
13 d.) establishing a general arbitration policy wherein arbitration decisions are
14 made based on a contents of said $2m$ plurality of output status bit
15 vectors;
16 e.) outputting according to said general arbitration policy said packets via at
17 least one transfer path to corresponding output links, wherein said
18 packets are advanced towards said destination nodes; and
19 f.) repeating the previous steps for as long as any of said packets and said
20 local packets are available for input.
- 1 2. The method for routing packets as claimed in Claim 1, further comprising
2 using hint bits from packet headers of said packets wherein said hint bits indicate
3 which of said $2m$ directions are feasible for each of said packets to move in,
4 whereby said establishing said general arbitration policy is completed more quickly.
- 1 3. The method for routing packets as claimed in Claim 1, further comprising:
2 creating said $2m$ of a plurality of possible bit vectors for each direction,
3 wherein each of said plurality of possible bit vectors indicates whether further

4 movement of any one of said packets in said each direction for each of said virtual
5 channels is still a routing goal.

1 4. The method for routing packets as claimed in Claim 3, wherein performing
2 an input phase of said establishing said general arbitration policy comprises:

- 3 a.) determining which of said packets not at their destination are completely
4 blocked;
- 5 b.) determining which of said packets not at their destination are unblocked;
- 6 c.) determining which of said packets at their destination are blocked;
- 7 d.) determining which of said packets at their destination are unblocked.

1 5. The method for routing packets as claimed in Claim 4, comprising creating
2 said 2^m of a plurality of input status bit vectors, corresponding to said 2^m directions,
3 wherein each of said input status bit vectors describes a ready to arbitrate indication,
4 and an input fullness of buffers indication for each of said plurality of virtual
5 channels and said plurality of processor injection FIFO's associated with one of said
6 2^m directions.

1 6. The method for routing packets as claimed in Claim 5, further comprising:
2 randomizing said input fullness of buffers indication of said input status bit
3 vectors in a current arbitration cycle if a predefined arbitration cycle criteria has
4 been met;
5 performing a longest queue priority arbitration in said current arbitration
6 cycle if said predefined arbitration cycle criteria has not been met.

1 7 The method for routing packets as claimed in Claim 6, further comprising
2 generating an input arbitration priority for said packets from a highest priority to a
3 lowest priority as follows;

- 4 a.) an unblocked one of said packets at the head of a priority virtual
5 channel;

- 6 b.) an unblocked one of said packets in a by-pass path, if said one of said
7 packets is a high priority packet and there are no said packets in the
8 priority virtual channel;
9 c.) an unblocked one of said packets in a non-priority virtual channel
10 having a highest said fullness of buffers indication;
11 d.) an unblocked one of said packets in the by-pass path, provided it is a
12 non-priority packet.

- 1 8. The method for routing packets as claimed in Claim 7, further comprising
2 selecting for each said unblocked one packet, an output direction and an output
3 virtual channel according to the following steps:
4 a.) for said high priority unblocked packet, selecting a direction and a
5 priority virtual channel for which a priority virtual channel bit of the
6 possible bit vector indicates the direction is still a routing goal of said
7 high priority unblocked packet;
8 b.) for said non-priority unblocked packet which is not dynamically
9 routable, selecting a direction and an escape virtual channel for which
10 an escape virtual channel bit of the possible bit vector indicates the
11 direction is still a routing goal of said non-priority unblocked packet
12 which is not dynamically routable;
13 c.) for said non-priority unblocked packet which is dynamically routable,
14 selecting a direction and a dynamic virtual channel for which the
15 possible bit vector indicates the direction is still a routing goal and for
16 which the output status bit vector indicates the greatest said buffer
17 space indication, wherein if no such channel and direction
18 combination exists then selecting a direction and escape virtual
19 channel for which the escape virtual channel bit of the possible bit
20 vector indicates said direction is still a routing goal.

- 1 9. The method for routing packets as claimed in Claim 8, wherein performing
2 an output phase of said establishing said general arbitration policy comprises:

3 setting the current output phase arbitration cycle to high priority for a
4 processor injection FIFO according to a predefined output phase arbitration cycle
5 criteria;

6 setting the current output phase arbitration cycle to low priority for a
7 processor injection FIFO according to said predefined output phase arbitration cycle
8 criteria.

1 10. The method for routing packets as claimed in Claim 9, further comprising
2 generating an output arbitration priority for said packets from a highest priority to a
3 lowest priority, wherein the current output phase arbitration cycle is low priority for
4 a processor injection FIFO, as follows;

5 a.) said unblocked one of said packets at the head of said priority virtual
6 channel that has the highest input fullness of buffers indication out of
7 a plurality of priority virtual channels;

8 b.) a high priority packet from a processor injection FIFO that has a
9 highest input fullness of buffers indication out of a plurality of
10 processor injection FIFO's;

11 c.) an unblocked one of said packets in a non-priority virtual channel
12 having a highest said input fullness of buffers indication out of a
13 plurality of non-priority virtual channels;

14 d.) a normal priority packet from a processor injection FIFO that has a
15 highest input fullness of buffers indication out of a plurality of
16 processor injection FIFO's.

1 11. The method for routing packets as claimed in Claim 9, further comprising
2 generating an output arbitration priority for said packets from a highest priority to a
3 lowest priority, wherein the current output phase arbitration cycle is high priority for
4 a processor injection FIFO, as follows;

5 a.) a high priority packet from a processor injection FIFO that has a
6 highest input fullness of buffers indication out of said plurality of
7 FIFO's;

- 8 b.) said unblocked one of said packets at the head of said priority virtual
9 channel that has the highest input fullness of buffers indication out of
10 said plurality of priority virtual channels;
11 c.) a normal priority packet from a processor injection FIFO that has a
12 highest input fullness of buffers indication out of said plurality of
13 processor injection FIFO's;
14 d.) an unblocked one of said packets in a non-priority virtual channel
15 having a highest said input fullness of buffers indication out of a
16 plurality of said non-priority virtual channels.

- 1 12. In a massively parallel computing system comprising a plurality of nodes
2 configured in m multi-dimensions, each node including a computing device, a
3 system for routing packets towards their destination nodes comprising:
4 a.) means for inputting said packets from at least one of 2^m directions
5 via at least one of a corresponding number of input links into a
6 corresponding at least one of a plurality of virtual channels;
7 b.) means for inputting local packets from a local of said nodes to at least
8 one of a plurality of processor injection FIFO's;
9 c.) means generating a 2^m plurality of output status bit vectors, each of
10 said output status bit vectors describing the availability of its link, the
11 availability of downstream dynamic virtual channels, and a
12 downstream buffer space indication for all virtual channels associated
13 with its direction;
14 d.) means for establishing a general arbitration policy wherein
15 arbitration decisions are made based on a contents of said 2^m
16 plurality of output status bit vectors;
17 e.) means for outputting according to said general arbitration policy said
18 packets via at least one transfer path to corresponding output links,
19 wherein said packets are advanced towards said destination nodes;
20 and
21 f.) means for repeating the previous steps for as long as any of said
22 packets and said local packets are available for input.

1 13. In a massively parallel computing system comprising a plurality of nodes
2 configured in m multi-dimensions, each node including a computing device, a
3 program storage device readable by machine, tangibly embodying a program of
4 instructions executable by the machine to perform method steps for routing packets
5 towards their destination, said method steps comprising:
6 a.) inputting said packets from at least one of $2m$ directions via at least
7 one of a corresponding number of input links into a corresponding at
8 least one of a plurality of virtual channels;
9 b.) inputting local packets from a local of said nodes to at least one of a
10 plurality of processor injection FIFO's;
11 c.) generating a $2m$ plurality of output status bit vectors, each of said
12 output status bit vectors describing the availability of its link, the
13 availability of downstream dynamic virtual channels, and a
14 downstream buffer space indication for all virtual channels associated
15 with its direction;
16 d.) establishing a general arbitration policy wherein arbitration decisions
17 are made based on a contents of said $2m$ plurality of output status bit
18 vectors;
19 e.) outputting according to said general arbitration policy said packets
20 via at least one transfer path to corresponding output links, wherein
21 said packets are advanced towards said destination nodes; and
22 f.) repeating the previous steps for as long as any of said packets and
23 said local packets are available for input.

1 14. The method for routing packets as claimed in Claim 13, further comprising
2 using hint bits from packet headers of said packets wherein said hint bits indicate
3 which of said $2m$ directions are feasible for each of said packets to move in,
4 whereby said establishing said general arbitration policy is completed more quickly.

1 15. The method for routing packets as claimed in Claim 13, further comprising:
2 creating said $2m$ of a plurality of possible bit vectors for each direction,
3 wherein each of said plurality of possible bit vectors indicates whether further

4 movement of any one of said packets in said each direction for each of said virtual
5 channels is still a routing goal.

1 16. The method for routing packets as claimed in Claim 15, wherein performing
2 an input phase of said establishing said general arbitration policy comprises:

- 3 a.) determining which of said packets not at their destination are
4 completely blocked;
- 5 b.) determining which of said packets not at their destination are
6 unblocked;
- 7 c.) determining which of said packets at their destination are blocked;
- 8 d.) determining which of said packets at their destination are unblocked.

1 17. The method for routing packets as claimed in Claim 16, comprising creating
2 said 2m of a plurality of input status bit vectors, corresponding to said 2m directions,
3 wherein each of said input status bit vectors describes a ready to arbitrate indication,
4 and an input fullness of buffers indication for each of said plurality of virtual
5 channels and said plurality of processor injection FIFO's associated with one of said
6 2m directions.

1 18. The method for routing packets as claimed in Claim 17, further comprising:
2 randomizing said input fullness of buffers indication of said input status bit
3 vectors in a current arbitration cycle if a predefined arbitration cycle criteria has
4 been met;
5 performing a longest queue priority arbitration in said current arbitration
6 cycle if said predefined arbitration cycle criteria has not been met.

1 19. The method for routing packets as claimed in Claim 18, further comprising
2 generating an input arbitration priority for said packets from a highest priority to a
3 lowest priority as follows;
4 a.) an unblocked one of said packets at the head of a priority virtual
5 channel;

- 6 b.) an unblocked one of said packets in a by-pass path, if said one of said
7 packets is a high priority packet and there are no said packets in the
8 priority virtual channel;
9 c.) an unblocked one of said packets in a non-priority virtual channel
10 having a highest said fullness of buffers indication;
11 d.) an unblocked one of said packets in the by-pass path, provided it is a
12 non-priority packet.

- 1 20. The method for routing packets as claimed in Claim 19, further comprising
2 selecting for each said unblocked one packet, an output direction and an output
3 virtual channel according to the following steps:
4 a.) for said high priority unblocked packet, selecting a direction and a
5 priority virtual channel for which a priority virtual channel bit of the
6 possible bit vector indicates the direction is still a routing goal of said
7 high priority unblocked packet;
8 b.) for said non-priority unblocked packet which is not dynamically
9 routable, selecting a direction and an escape virtual channel for which
10 an escape virtual channel bit of the possible bit vector indicates the
11 direction is still a routing goal of said non-priority unblocked packet
12 which is not dynamically routable;
13 c.) for said non-priority unblocked packet which is dynamically routable,
14 selecting a direction and a dynamic virtual channel for which the
15 possible bit vector indicates the direction is still a routing goal and for
16 which the output status bit vector indicates the greatest said buffer
17 space indication, wherein if no such channel and direction
18 combination exists then selecting a direction and escape virtual
19 channel for which the escape virtual channel bit of the possible bit
20 vector indicates said direction is still a routing goal.

- 1 21. The method for routing packets as claimed in Claim 20, wherein performing
2 an output phase of said establishing said general arbitration policy comprises:

3 setting the current output phase arbitration cycle to high priority for a
4 processor injection FIFO according to a predefined output phase arbitration cycle
5 criteria;

6 setting the current output phase arbitration cycle to low priority for a
7 processor injection FIFO according to said predefined output phase arbitration cycle
8 criteria.

1 22. The method for routing packets as claimed in Claim 21, further comprising
2 generating an output arbitration priority for said packets from a highest priority to a
3 lowest priority, wherein the current output phase arbitration cycle is low priority for
4 a processor injection FIFO, as follows;

- 5 a.) said unblocked one of said packets at the head of said priority virtual
6 channel that has the highest input fullness of buffers indication out of
7 a plurality of priority virtual channels;
- 8 b.) a high priority packet from a processor injection FIFO that has a
9 highest input fullness of buffers indication out of a plurality of
10 processor injection FIFO's;
- 11 c.) an unblocked one of said packets in a non-priority virtual channel
12 having a highest said input fullness of buffers indication out of a
13 plurality of non-priority virtual channels;
- 14 d.) a normal priority packet from a processor injection FIFO that has a
15 highest input fullness of buffers indication out of a plurality of
16 processor injection FIFO's.

1 23. The method for routing packets as claimed in Claim 21, further comprising
2 generating an output arbitration priority for said packets from a highest priority to a
3 lowest priority, wherein the current output phase arbitration cycle is high priority for
4 a processor injection FIFO, as follows;

- 5 a.) a high priority packet from a processor injection FIFO that has a
6 highest input fullness of buffers indication out of said plurality of
7 FIFO's;

- 8 b.) said unblocked one of said packets at the head of said priority virtual
9 channel that has the highest input fullness of buffers indication out of
10 said plurality of priority virtual channels;
- 11 c.) a normal priority packet from a processor injection FIFO that has a
12 highest input fullness of buffers indication out of said plurality of
13 processor injection FIFO's;
- 14 d.) an unblocked one of said packets in a non-priority virtual channel
15 having a highest said input fullness of buffers indication out of a
16 plurality of said non-priority virtual channels.

- 1 24. In a massively parallel computing system comprising a plurality of
2 computing processor nodes configured in m multi-dimensions, each node including
3 a computing device, a network switch for routing packets towards their destination
4 nodes comprising:
- 5 a.) at least one input unit for inputting said packets from at least one of $2m$
6 directions via at least one of a corresponding number of input links into a
7 corresponding at least one of a plurality of virtual channels;
 - 8 b.) at least one input link for inputting local packets from a local of said nodes to
9 at least one of a plurality of processor injection FIFO's;
 - 10 c.) an output status bit vector generator for generating a $2m$ plurality of output
11 status bit vectors, each of said output status bit vectors describing the
12 availability of its link, the availability of downstream dynamic virtual
13 channels, and a downstream buffer space indication for all virtual channels
14 associated with its direction;
 - 15 d.) an input arbiter and an output arbiter for establishing a general arbitration
16 policy wherein arbitration decisions are made based on a contents of said $2m$
17 plurality of output status bit vectors;
 - 18 e.) at least one output link for outputting according to said general arbitration
19 policy said packets via at least one transfer path to corresponding output
20 links, wherein said packets are advanced towards said destination nodes.

1 25. The network switch as claimed in Claim 24 comprising an integration of said
2 network switch on the same integrated circuit as said computer processor nodes.

1 26. The network switch as claimed in Claim 24 comprising at least one processor
2 reception FIFO dedicated to said at least one input unit, wherein there is a plurality
3 of processor reception FIFO's for the same plurality of input units, whereby
4 contention among a plurality of input units is decreased.

1 27. In a massively parallel computing system comprising a plurality of computing
2 processor nodes configured in m multi-dimensions, each node including a
3 computing device, a virtual cut through network switch with dynamic routing for
4 routing packets towards their destination nodes comprising:
5 an integration of said network switch on the same integrated circuit as said computer
6 processor nodes.

1 28. The virtual cut through network switch as claimed in Claim 27, comprising:
2 at least one processor reception FIFO dedicated
3 to said at least one input unit, wherein there is a plurality of processor reception
4 FIFO's for the same plurality of input units, whereby contention among a plurality
5 of input units is decreased.