




(86) **Date de dépôt PCT/PCT Filing Date:** 2019/12/12
 (87) **Date publication PCT/PCT Publication Date:** 2020/06/18
 (45) **Date de délivrance/Issue Date:** 2023/10/31
 (85) **Entrée phase nationale/National Entry:** 2021/06/11
 (86) **N° demande PCT/PCT Application No.:** US 2019/065967
 (87) **N° publication PCT/PCT Publication No.:** 2020/123796
 (30) **Priorité/Priority:** 2018/12/12 (US62/778,354)

(51) **Cl.Int./Int.Cl. H04L 12/28** (2006.01),
H04L 12/42 (2006.01), **H04L 12/437** (2006.01),
H04L 12/46 (2006.01), **H04L 12/54** (2013.01),
H04L 45/745 (2022.01), **H04L 49/109** (2022.01),
H04L 69/22 (2022.01)

(72) **Inventeurs/Inventors:**
 REED, COKE S., US;
 MURPHY, DAVID, US;
 DENNY, RONALD R., US;
 ...

(73) **Propriétaires/Owners:**
 INTERACTIC HOLDINGS, LLC, US;
 REED, COKE S., US; ...

(54) **Titre : PROCEDURE ET APPAREIL POUR UN TRANSFERT DE DONNEES AMELIORE ENTRE DES COEURS DE PROCESSEUR**
 (54) **Title: METHOD AND APPARATUS FOR IMPROVED DATA TRANSFER BETWEEN PROCESSOR CORES**



(57) **Abrégé/Abstract:**

Embodiments of an interconnect apparatus enable improved signal integrity, even at high clock rates, increased bandwidth, and lower latency. In an interconnect apparatus for core arrays a sending processing core can send data to a receiving core by forming a packet whose header indicates the location of the receiving core and whose payload is the data to be sent. The packet is sent to a Data Vortex switch on the same chip as an array of processing cores and routes the packet to the receiving core first by routing the packet to the processing core array containing the receiving processing core. The Data Vortex switch then routes the packet to the receiving processor core in a processor core array. The Data Vortex switches are not crossbar switches, eliminating globally setting and resetting the Data Vortex switches as different groups of packets enter the switches. Mounting the Data Vortex switch on the same chip as the array of processing cores reduces power required and latency.

(72) **Inventeurs(suite)/Inventors(continued):** IVES, MICHAEL R., US; DEVANY, REED, US

(73) **Propriétaires(suite)/Owners(continued):**MURPHY, DAVID, US; DENNY, RONALD R., US; IVES, MICHAEL R., US;
DEVANY, REED, US

(74) **Agent:** SMART & BIGGAR LP

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(10) International Publication Number
WO 2020/123796 A1

(43) International Publication Date
18 June 2020 (18.06.2020)

(51) International Patent Classification:

H04L 12/28 (2006.01) *H04L 12/70* (2013.01)
H04L 12/42 (2006.01) *H04L 12/701* (2013.01)
H04L 12/46 (2006.01) *H04L 12/703* (2013.01)
H04L 12/54 (2013.01)

(72) Inventors; and

(71) Applicants: **REED, Coke, S.** [US/US]; 8148 Meandering Way, Austin, TX 78759 (US). **MURPHY, David** [US/US]; 4831 Trail Crest Circle, Austin, TX 78735 (US). **DENNY, Ronald, R.** [US/US]; 10019 Northwood Lane, Brooklyn Park, MN 55443 (US). **IVES, Michael, R.** [US/US]; N1012 Summer Breeze Lane, Greenville, WI 54942 (US). **DEVANY, Reed** [US/US]; 88005 Raintree Place, Austin, TX 78759 (US).

(21) International Application Number:

PCT/US2019/065967

(22) International Filing Date:

12 December 2019 (12.12.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/778,354 12 December 2018 (12.12.2018) US

(71) Applicant: **INTERACTIC HOLDINGS, LLC** [US/US];

2303 Rio Grande Street, Austin, TX 78705 (US).

(74) Agent: **NOWAK, Keith, D.**; Carter Ledyard & Milburn LLP, 2 Wall Street, New York, NY 10005 (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, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,

(54) Title: METHOD AND APPARATUS FOR IMPROVED DATA TRANSFER BETWEEN PROCESSOR CORES

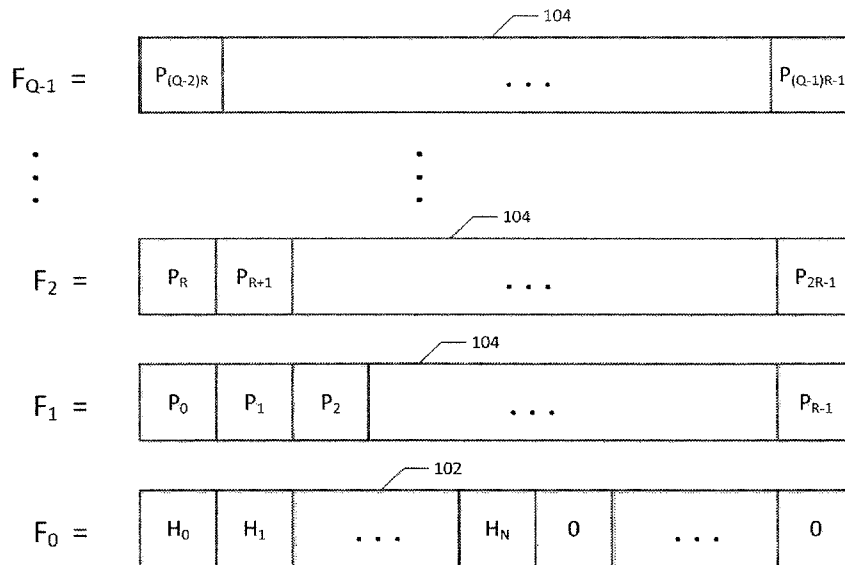


FIG. 1

(57) Abstract: Embodiments of an interconnect apparatus enable improved signal integrity, even at high clock rates, increased bandwidth, and lower latency. In an interconnect apparatus for core arrays a sending processing core can send data to a receiving core by forming a packet whose header indicates the location of the receiving core and whose pay load is the data to be sent. The packet is sent to a Data Vortex switch on the same chip as an array of processing cores and routes the packet to the receiving core first by routing the packet to the processing core array containing the receiving processing core. The Data Vortex switch then routes the packet to the receiving processor core in a processor core array. The Data Vortex switches are not crossbar switches, eliminating globally setting and resetting the Data Vortex switches as different groups of packets enter the switches. Mounting the Data Vortex switch on the same chip as the array of processing cores reduces power required and latency.



WO 2020/123796 A1

WO 2020/123796 A1 

KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, 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, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, 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, KM, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

METHOD AND APPARATUS FOR IMPROVED DATA TRANSFER BETWEEN PROCESSOR CORES

CROSS-REFERENCE TO RELATED PATENTS

[0001] This application claims priority to U.S. Provisional Patent Application No. 62/778,354 filed December 12, 2018, entitled "Data Transfer Between The Cores In A Microprocessor". The disclosed system and operating method is also related to subject matter disclosed in the following patents: (1) U.S. Patent No. 5,996,020 entitled, "A Multiple Level Minimum Logic Network", naming Coke S. Reed as inventor; (2) U.S. Patent No. 6,289,021 entitled, "A Scalable Low Latency Switch for Usage in an Interconnect Structure", naming John Hesse as inventor; (3) U.S. Patent No. 6,754,207 entitled, "Multiple Path Wormhole Interconnect", naming John Hesse as inventor, U.S. Patent No. 9,954,797 entitled " Parallel Data Switch", naming Coke S. Reed and Davis Murphy as inventors.

BACKGROUND

[0002] Components of large computing and communication systems can be configured with interconnect structures of switch chips connected by interconnect lines. Increasing the switch-chip port count decreases the number of chip-to-chip hops, resulting in lower latency and lower cost. What is needed in these systems is switch chips that have high port count and are also able to handle short packets.

[0003] In present day multi-core processors, data is transferred between cores using a mesh. The cores are tiles arranged in a mesh structure. These techniques have been used in connecting the cores on a chip, but are not effective in transferring data from a core on a first processor to a core

on a second processor. In addition to the difficulties due to the mesh structure, the use of long packets passing through crossbar switches carrying data between chips presents additional difficulties in multi-chip applications. The long packets cause low bandwidth, high latency, limited scalability, and high congestion. The purpose of the invention claimed herein is to provide a high bandwidth and low latency method and apparatus to exchange information between processor computing cores. This is accomplished by mounting a Data Vortex switch and an array of processing cores on the same chip.

SUMMARY

[0004] Embodiments of an interconnect apparatus enable improved signal integrity, even at high clock rates, increased bandwidth, and lower latency. In an interconnect apparatus for core arrays a sending processing core can send data to a receiving core by forming a packet whose header indicates the location of the receiving core and whose pay load is the data to be sent. The packet is sent to a Data Vortex switch described herein and in the patents referenced herein. The Data Vortex switch is on the same chip as an array of processing cores and routes the packet to the receiving core first by routing the packet to the processing core array containing the receiving processing core. The Data Vortex switch then routes the packet to the receiving processor core in a processor core array. Since the Data Vortex switches are not crossbar switches, there is no need to globally set and reset the Data Vortex switches as different groups of packets enter the switches. Mounting the Data Vortex switch on the same chip as the array of processing cores reduces the power required and reduces latency.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] Embodiments of the invention relating to both apparatus and method of operation may best be understood by referring to the following description and accompanying drawings:

FIGURE 1 is a data structure diagram illustrating an embodiment of a format of a packet comprising a number of sub-packet flits;

FIGURE 2A is a schematic block diagram depicting a high-level view of a switch described in an embodiment of referenced United States patent numbers 6,289,021 and 6,754,207 including node arrays, connection lines, and FIFOs;

FIGURE 2B is a schematic block diagram showing a high-level view of a switch in an embodiment of the system disclosed herein comprising node arrays interconnects;

FIGURE 3A is a schematic block diagram illustrating a pair of simple switch nodes as described in referenced United States patent numbers 6,289,021 and 6,754,207;

FIGURE 3B is a schematic block diagram showing a connected pair of nodes as described in the "flat latency" or "double-down" version of the switches described in referenced United States patent numbers 6,289,021 and 6,754,207;

FIGURE 4 is a schematic block diagram depicting a building block for a node in an embodiment of the disclosed system (referred to herein as an LDM module) including a one-tick delay logic element, a one-tick delay FIFO element, and a multiplexing device for combining busses;

FIGURE 5 is a block diagram showing four interconnected LDM modules in accordance with an embodiment of the disclosed system;

FIGURE 6 is a block diagram illustrating a switching node used in an embodiment of the disclosed system;

FIGURE 7 is a block diagram showing the source and timing of the control signals sent to the logic elements in an embodiment of the switching node of the disclosed system;

FIGURE 8 is a block diagram illustrating control registers used in the logic elements of the switching nodes of an embodiment of the system disclosed herein;

FIGURES 9A, 9B, and 9C are schematic block diagrams that illustrate interconnections of nodes on various levels of the interconnect structure;

FIGURE 10 is a timing diagram which illustrates timing of message communication in the described interconnect structure; and

FIGURE 11 is a pictorial representation illustrating the format of a message packet including a header and payload.

FIGURE 12 is a block diagram of a chip containing two components. The first component is a Data Vortex switch and the second component is an array of processing cores. The Data Vortex switch receives data packets and routes them to the appropriate cores in the core array.

FIGURE 13 consists of the chip in **FIGURE 12** with connections that transfer packets from the cores in the array of processing cores to the Data Vortex switch. This provides a mechanism for a first processing core in the array of processing cores to send data to a second processing core in the array of processing cores.

FIGURE 14 shows four processor core arrays, each having a Data Vortex switch and an array of processing cores on the same chip. The transfer of packets from each processor core array is directed through a master Data Vortex switch.

DETAILED DESCRIPTION

[0006] The devices, systems, and methods disclosed herein describe a network interconnect system that is extremely effective in connecting a large number of objects, for example line cards in a router, network interface cards in a parallel computer, or other communication systems and devices. The described network interconnect system has extremely high bandwidth as well as extremely low latency.

[0007] Computing and communication systems attain highest performance when configured with switch chips that have high port count and are also able to handle short packets. The Data Vortex switch chips described in United States patent numbers 5,996,020 and 6,289,021 have extremely high port counts and have the ability to transmit short message packets.

[0008] The systems and methods disclosed herein include several improvements over United States patent numbers 6,289,021 and 6,754,207, attained by one or more of a number of enhancements, including the following two basic improvements: 1) the bandwidth is increased and the first-bit-in to last-bit-out latency is decreased by us in parallel data lines between nodes; and 2) the bandwidth is further increased and the latency is further reduced by a logic that sets up a data path through the switch that contains a one-bit-long parallel FIFO at each level, enabling the use of a much faster clock than was possible in United States patent numbers 6,289,021 and 6,754,207.

[0009] United States patent number 6,289,021 describes a switch that is suitable to be placed on a chip. In that system, data (in the form of packets) passes through the switch in wormhole fashion on one-bit wide data paths. The packets include a header and a payload. The first bit of the header is a status bit (set to a value of 1 in most embodiments) indicating the presence of a message. In a simple arrangement, the remaining header bits represent the binary address of a target output port. The topology of the switch includes a richly connected set of rings. A $(2^N \times 2^N)$ switch includes rings arranged in $(N+1)$ levels with connections between rings on different levels. Packets enter the switch at level N and exit the switch at level 0 . The header of a message packet entering the switch on level N has one status bit and N target address bits. The logic at a node on level N makes routing decisions based on: 1) the status bit; 2) the first bit of the address in the header; 3) a control signal sent from a node on level $N-1$; and 4) (in the basic embodiment) a control signal from a node on level N . The first bit of the address in the header is used by the logic on level N . When the logic on a level N node directs a packet to a node on level $N-1$, the first bit of the address is discarded. This is done for several reasons: 1) The first address bit is not needed for routing decisions on lower levels; 2) the discarding of this bit allows the message packets on level $N-1$ to travel ahead of the packets on level N so that, based on incoming packets, level $N-1$ nodes can send control signals to level N nodes, thus enabling the level $N-1$ nodes to direct level N traffic; 3) the discarding of the first header bit ensures that the most significant bit of the remaining header bits is the bit that is needed to route the packet on level $N-1$. This process continues throughout the switch so that a packet on level K has one status bit followed by K address bits.

[0010] A consequence of this design is that data paths can be established that cut directly between levels. The timing of the system is such that two clock ticks are required for a status bit to move

between two logic nodes on the same ring, but only one tick is required for the status bit to move between two nodes on different levels (a node on a level K ring is referred to as a level K node). Therefore, if the path of a packet through the switch contains N downward steps (steps between rings on different levels) and J steps between two nodes on a ring at a given level, then $(N+2J+1)$ ticks are required before the first payload bit arrives at the output level 0. When the status bit is on level 0, there are 2J one-tick delays on different levels, with one data bit in each of the one-bit FIFO delay elements. The passing of information through multiple transistors on nodes at different levels necessarily limits the clock rate of the system. In fact, if a packet passes down at each step, the status bit arrives at level 0 while the first payload bit is on level N, the top entry level.

[0011] In contrast, for the system described herein, each bit of the packet passes through at least one single-tick FIFO on each level, advantageously enabling the signal to be reconstituted at each node and enabling the system described herein to operate at higher clock rates than systems described in United States patent number 6,289,021.

[0012] The switching systems described in United States patent numbers 5,996,020, 6,289,021, and 6,754,207 provide low latency with high bandwidth and also support short packets. The topology of the switches in United States patent numbers 5,996,020, 6,289,021, and 6,754,207 includes a richly interconnected set of rings. **FIGURE 2A** is a high-level block diagram of an embodiment of a switch described in United States patent numbers 6,289,021 and 6,754,207. The entire structure illustrated in **FIGURE 2A** fits on a single chip. The interconnection lines **210, 212, 214, 216** and **218** are bit serial. The specification that a packet fits on a single ring calls for inclusion of FIFO elements **220**. To decrease the probability of a packet passing through FIFO **220**, all of the packets are inserted into a single input node array through

lines **202**. The switch illustrated in **FIGURE 2A** can be built using simple nodes illustrated in **FIGURE 3A** or by using "double-down" or "flat latency" nodes illustrated in **FIGURE 3B**.

[**0013**] First consider the simple switch U illustrated in **FIGURE 3A**. One aspect of the operation of the switches of United States patent numbers 6,289,021 and 6,754,207 is that the first bit of a data packet can enter the switch node U only at specific packet entry times. At a given packet entry time T, no more than one packet can enter the switch node U. This is the case because of the novel use of control lines described in United States patent numbers 5,996,020, 6,289,021, and 6,754,207.

[**0014**] For each switch node U of **FIGURE 3A** on level N-K, there is a K long poly-bit $PB_u = (b_0, b_1, b_2, \dots, b_{K-1})$ such that each packet PK entering U has a target destination whose binary representation has leading bits $(b_0, b_i, b_2, \dots, b_{K-1})$. Each packet that exits switch node U through line **306** has a target destination whose binary representation has leading bits $(b_0, b_i, b_2, \dots, b_{K-1}, 1)$. Next consider the switch node L on level K such that each packet PK entering L has a poly-bit $PB_L = PB_u$ so that each packet PK entering L has a target destination whose binary representation has leading bits $(b_0, b_i, b_2, \dots, b_{K-1})$. Each packet that exits switch node L through line **316** has a target destination whose binary representation has leading bits $(b_0, b_1, b_2, \dots, b_{K-1}, 0)$.

[**0015**] In case a packet PK enters U and the target address of PK has leading bits $(b_0, b_1, b_2, \dots, b_{K-1}, 1)$ and the control line **310** indicates a non-busy condition, then PK will exit U through line **306**. Otherwise, PK must exit U through line **304**. In case a packet PK enters L and the target address of PK has leading bits $(b_0, b_i, b_2, \dots, b_{K-1}, 0)$ and the control line **320** indicates a non-busy condition, then PK will exit L through line **316**. Otherwise, PK must exit L through line **314**.

[0016] The "double-down" switch **380** illustrated in FIGURE 38 combines nodes U and L using additional logic and interconnection lines **342** and **344**. The double-down switch DD **380** is positioned so that each packet that enters switch DD on level N-K has a destination whose target address has K leading bits $PB_{DD} = (b_0, b_1, b_2, \dots, b_{K-1})$. Switch DD is positioned in the network so that each packet that exits node U of DD through line **326** has a target address with leading bits $(b_0, b_1, b_2, \dots, b_{K-1}, 1)$ and each packet that exits node L of DD through line **336** has a target address with leading bits $(b_0, b_1, b_2, \dots, b_{K-1}, 1)$.

[0017] The switch DD operates as follows when a packet PKT enters node U:

- 1) If a packet PKT enters node U of DD through line **328** and the leading address bits of the target of PKT are $(b_0, b_1, b_2, \dots, b_{K-1}, 1)$, then a) if there is no busy signal on line **330**, then the logic of U directs the packet PKT down line **326**; or b) if there is a busy signal on line **330**, then the logic of U directs the packet PKT through line **324** to another double-down switch in the network. In either case, the logic of U sends a busy signal on line **344** to node L to indicate that line **326** is busy.

- 2) If a packet PKT enters node U of DD through line **328** and the leading address bits of the target of PKT are $(b_0, b_1, b_2, \dots, b_{K-1}, 0)$, then a) if there is no busy signal on line **344**, then the logic of U directs the packet PKT through line **342** to node L so that the logic of L can send PKT down line **336**; or b) if there is a busy signal on line **344**, then the logic of U directs the packet PKT through line **324** to another double-down switch in the network. The logic of U sends a busy signal on line **344** to node L to indicate that line **326** is busy only if there is a busy signal on line **330**.

[0018] Packets entering node L behave similarly. Thus, when a packet PKT enters node L of the switch **DD**, the following set of events occur:

1) If a packet PKT enters node L of **DD** through line **338** and the leading address bits of the target of PKT are $(b_0, b_1, b_2, \dots, b_{K-1}, 0)$, then a) if there is no busy signal on line **348**, then the logic of L directs the packet PKT down line **336**; or b) if there is a busy signal on line **348**, then the logic of L directs the packet PKT through line **334** to another double-down switch in the network. In either case, the logic of L sends a busy signal on line **344** to node U to indicate that line **336** is busy.

2) If a packet PKT enters node L of **DD** through line **338** and the leading address bits of the target of PKT are $(b_0, b_1, b_2, \dots, b_{K-1}, 1)$, then a) if there is no busy signal on line **344**, then the logic of L directs the packet PKT through line **342** to node U so that the logic of U can send PKT down line **326**; or b) if there is a busy signal on line **344**, then the logic of L directs the packet PKT through line **334** to another double-down switch in the network. The logic of L sends a busy signal on line **344** to node U to indicate that line **336** is busy only if there is a busy signal on line **348**.

[0019] The switching system described in the system disclosed here. represents important improvements over the switching system described in United States patent numbers 6,289,021 and 6,754,207. The main improvements include

1) the addition of parallel data paths through the system that enable higher bandwidths than were possible in the systems built in the United States patent numbers 6,289,021 and 6,754,207;

2) a modified timing system that simultaneously enables the creation of data paths having FIFOs at each logic node with the advantages of said FIFOs including the ability to use high clock rates; 3) a timing system that employs only one tick for a packet flit to move between nodes on different levels and two ticks for a packet flit to move between nodes on the same level. Advantageously, the FIFO lines 214 of **FIGURE 2A** are eliminated.

[0020] **FIGURE 1** illustrates the layout of a packet. The packet is decomposed into Q sub-packets referred to as flits. Each of the flits has R bits. The flits are designed to travel through an R wide bus. The first flit, F_0 102, is the header flit. It consists of a status bit H_0 , N routing bits H_1, H_2, \dots, H_N , and additional bits (set to 0 in **FIGURE 1**) that can be used to carry other information, e.g. error correction bits or QOS. The status bit H_0 is set to 1 to indicate the presence of a packet. The N routing bits represent the binary address of the target. Therefore, a switch with N routing bits, is used in a switch of radix 2^N . A logic L on level K uses bit H_K in conjunction with control bits to route F_0 .

[0021] **FIGURE 2A** represents a high level block diagram of the switch presented in United States patent numbers 6,289,021 and 6,754,207. **FIGURE 2B** represents a high level block diagram of the switch of the present patent wherein all header bits can be used to route the packet through the switch. In referenced United States patent numbers 6,289,021 and 6,754,207 and also in the present disclosure, the node arrays are arranged in rows and columns. The node array $NA(U,V)$ in a radix 2^N switch includes 2^N switching nodes on level U in column V. In United States patent numbers 6,289,021 and 6,754,207 and also in the present system, in an example embodiment, a switch of radix 2^N has (N+1) rows (or levels) of node arrays from an entry level N to an exit level 0. The switch has some number M of columns (where the number of columns is a design parameter). There are important novel improvements in the system introduced in the

disclosed system. In the disclosed system, the lines between the node arrays include busses that replace the single-bit-wide lines used in Patents No. 6,289,021 and Patent No. 6,754,207. The nodes in the node array have novel new designs that enable lower latencies and higher bandwidths. Additionally, in United States patent numbers 6,289,021 and 6,754,207, the system includes a FIFO shift register **220** of sufficient length so that an entire packet can fit on a given ring on level 0; two packets can fit on a ring on level 1, and in general 2^R packets can fit on a ring on level R. This FIFO can be eliminated in switches using the technology of the present disclosure because the width of the parallel bus will make it possible to have a sufficient number of active nodes, thus enabling entire messages to fit in the active elements of a row of nodes. Referring to **FIGURE 2B**, each data carrying line can include a bus with a width that is equal to the length of an R-long sub-packet as illustrated in **FIGURE 1**. The R long sub-packet can be referred to as a flit. The busses connect nodes in which each node is a module containing a one-tick logic element followed by a one tick delay element followed by a multiplexer (an LDM module). Two flits fit in an LDM module. Therefore, a packet containing Q flits can fit in $Q/2$ LDM modules on a single row of node arrays in **FIGURE 2B**. In one illustrative embodiment, the top level (level N) contains LDM modules on a single ring. The LDM modules on level N-1 are arranged on two rings. The bottom node array includes 2^N rings with each ring including M nodes connected by an R-wide bus. The arrangement of the nodes and busses in the node arrays make the ring structure possible. Each ring on level 0 must be long enough to hold an entire packet consisting of Q flits. Each ring on level 0 delivers packets to a single target address. A plurality of output ports **238** can connect from a single level 0 ring. Each of these outputs from a single ring delivers data to the same target output. Data is input into the structure through input busses **222**. In the simple embodiment, as shown, each logic node on a given level W is connected to forward packets to the next logic node

on the associated ring and also is positioned to forward packets to one of two logic nodes on level W-1 . In the present configuration (as depicted in United States Patent numbers 6,289,021 and 6,754,297) a logical element on level W-1 has priority over a logic node on level W to forward data to a logic element on level W-1. In the illustrative example presented herein, each level contains the same number of logic nodes.

[0022] In embodiments of United States patent numbers 6,289,021 and 6,754,207, all of the packets enter node array NA(N,0) (where N denotes the level of the array and 0 denotes the entry angle of the node array) in order to minimize the probability of a given packet entering FIFO 220. The elimination of the FIFO 220 enables the embodiments of the present disclosure to provide for the insertion of packets at a plurality of angles. The plurality of insertion angles reduces the number of packets inserted at any given angle, thereby reducing congestion. In the present invention, the first flit F_0 of a packet can enter a node only at specified entry times that are different for different nodes in the system.

[0023] An embodiment of the presently-disclosed switch has connections that correspond to the connections illustrated in **FIGURE 3B**. Therefore, the wiring pattern topology for the disclosed system, as illustrated in **FIGURE 2B**, can be the same as the wiring pattern topology illustrated in **FIGURE 2A** and described in United States patent numbers 5,996,020, 6,289,021, and 6,754,207.

[0024] Referring to **FIGURE 4**, a block diagram of an individual switching module 400 is shown including a logic element 402, a delay element 404 and a multiplexing element 406. The module with logic, delay and multiplexer (mux) is referred to as an LDM module. The LDM modules are components used in the construction of the switching nodes in the node arrays illustrated in **FIGURE 2B**. Timing of the packet entry into a logic element 402 is important for correct operation

of the high-speed parallel switch of the disclosed system. Each logic device **402** is capable of holding exactly one flit of a packet. Moreover, each delay device **404** is also capable of holding exactly one packet flit. At the conclusion of a time step, the logic device **402** may or may not contain one flit of a packet. Similarly, at the conclusion of a time step, the delay device **404** also may or may not contain one flit of a packet. In case the delay device contains one flit F_L of a packet P and F_L is not the last flit of the packet (i.e., $L < Q$), then the logic device contains flit F_{L+1} of the packet. If $L \neq 0$ and F_L is in a logic unit or delay unit, then the location of flit F_L at time T_s is equal to the location of flit $F_{(L-1)}$ at time T_{s-1} . In this manner, packets travel in wormhole fashion through the switch.

[0025] In the illustrative embodiment, an LDM module comprises logic, delay, and multiplexer devices configured to synchronize timing of a packet.

[0026] Referring to **FIGURE 5** in conjunction with **FIGURE 4**, basic operations and timing of the LDM modules are described. The LDM module with logic element L_1 is in the node array $NA(W,Z)$; logic element L_2 is in $NA(W, Z+1)$; L_3 is in $NA(W-1, Z)$; and logic element L_4 is in $NA(W-1,Z+1)$.

[0027] At each packet insertion time for L_1 , the logic unit L_1 checks for a flit arrival by checking for a status bit H_0 set to 1. In case, at a flit arrival time T_s for logic unit L_1 , the logic senses that H_0 is set to 0, the logic unit identifies that no packet is arriving in this time slot and takes no action until the next first flit packet arrival time. If at a first flit arrival time T_s for L_1 (as identified by a clock or counter) the logic unit senses a 1 in the status bit slot H_0 in F_0 , the logic unit ascertains that it contains the first flit of a valid packet PKT and precedes as follows:

- A. If 1) based on the bit H_w of F_0 , L_1 determines if a path exists from L_4 to a target output of PKT, and 2) the control signal from L_3 indicates that L_1 is free to use line **412**, then L_1 sends the first flit F_0 of PKT through M_3 to arrive at L_4 at time T_{s+1} .
- B. If one or both of the above conditions 1) and 2) is not met, then L_1 sends the first flit F_0 of PKT to D_1 in order to arrive at D_1 at time T_{s+1} and subsequently to arrive at L_2 at time T_{s+2} (one unit of time after first flits arrive at L_4).

[0028] A detailed discussion of the use of the routing bit H_w and the control signals is included in the discussions of **FIGURES. 6, 7, and 8**. Timing of first flit arrival at logic units on different levels in the same column is important for the proper operation of the switching system. Since first flits arrive at logic elements in $NA(W-1, Z+1)$ one time unit before first flits arrive at logic elements in $NA(W, Z+1)$. sufficient time is available for control signals generated on level $W-1$ to control the routing of packets on level W .

[0029] Given that an initial flit F_0 of a packet PKT arrives at a logic unit L in time step T_s , then the next flit F_1 of PKT will arrive at L in time step T_{s+1} . This continues until the last flit F_{Q-1} of PKT arrives at the logic unit L at time T_{s+Q-1} . Similarly, given the arrival of an initial flit F_0 of a packet PKT arrives at a delay unit D in time step T_{s+1} , then the next flit F_1 of PKT will arrive at D in time step T_{s+2} . This continues until the last flit F_{Q-1} of PKT arrives at D at time T_{s+Q} . Each time that a flit of a packet arrives at a logic or delay unit, the signal of the flit is regenerated. This signal regeneration at each tick allows for higher chip clock rates. In a simple "single-down" embodiment, an LDM module can be used as a node in the switch.

[0030] Referring to **FIGURE 6** a switching node used in the switch of an embodiment of disclosed system is shown. The switching node contains two LDM modules **610** and **620** and a 2x2 crossbar **602**. Suppose the switching node represented in **FIGURE 6** is in node array $NA(W, Z)$. The LDM

modules are of two types: 1) a type α LDM module **620** that is able to control the crossbar **602** by sending a signal down line **608** and 2) a type β module **610** that is not able to control the crossbar **602**. Suppose the interesting case where the level W is not equal to either 0 or N . The set of packet entry times for the type α module **620** is equal to the set of packet entry times for the type β module **610**.

[0031] Referring to **FIGURE 7**, a schematic block diagram illustrates an embodiment of an interconnect apparatus which enables improved signal integrity, even at high clock rates, increased bandwidth, and lower latency.

The illustrative interconnect apparatus comprises a plurality of logic units and a plurality of buses coupling the plurality of logic units in a selected configuration of logic units which can be considered to be arranged in triplets comprising logic units LA **624**, LC **724**, and LD **710**. The logic units LA **624** and LC **724** are positioned to send data to the logic unit LD **710**. The logic unit LC **724** has priority over the logic unit LA **624** to send data to the logic unit LD **710**. For a packet PKT divided into subpackets, a subpacket of the packet PKT at the logic unit LA **624**, and the packet specifying a target either: (A) the logic unit LC **724** sends a subpacket of the packet PKT to the logic unit LD **710** and the logic unit LA **624** does not send a subpacket of the packet PKT to the logic unit LD **710**; (B) the logic unit LC **724** does not send a subpacket of data to the logic unit LD **710** and the logic unit LA **624** sends a subpacket of the packet PKT to the logic unit LD **710**; or (C) the logic unit LC **724** does not send a subpacket of data to the logic unit LD **710** and the logic unit LA **624** does not send a subpacket of the packet PKT to the logic unit LD **710**.

[0032] In the illustrative interconnect structure, the logic, delay, and multiplexer units can be configured with insufficient memory to hold the entire packet, and thus have only a bus-wide first-in-first-out (FIFO) buffer. Thus, packets are communicated on a bus wide data path.

[0033] A logic node does not reassemble a packet. A first subpacket, called a flit, of a packet PKT arrives at a logic node LA at a given time T_1 . At time T_2 , the first flit of PKT arrives at the next downstream logic unit or delay unit. Also at time T_2 , the second flit of PKT arrives at logic unit LA 624. In fact, the packet is never reassembled in the switch it leaves the switch one flit at a time. As a detail, a flit formed of R bits (See FIGURE 1) travels through the switch in wormhole fashion and exits through a SER-DES (serializer-deserializer) module connected to a sequential interconnect that runs R times as fast.

[0034] Logic unit LA 624 will send PKT to logic unit LD 710 provided that 1) a path exists from logic unit LD 710 to the target output port for PKT; and 2) logic unit LA 624 is not blocked from traveling to logic unit LD 710 by a logic element LC with a higher priority than logic unit LA 624 to send to logic unit LD 710. Referring to FIGURE 7, only logic unit LC 724 has a higher priority than logic unit LA 624 to send to logic unit LD 710 whereas both logic units LB and LE have a higher priority than logic unit LA 624 to send to logic unit LF. In the example, a flit of PKT is at logic unit LA 624. Logic unit LA 624 routes the first flit of PKT based on header information and incoming control bits. Logic unit LA 624 sends the second flit of PKT to the same element that it sent the first flit. In any case, logic unit LA 624 cannot hold onto a packet, if a flit arrives at time T_1 , it is forwarded at time T_2 .

[0035] The interconnect structure transfers the packets and subpackets in a sequence of time step. With a sequence of flits of the packet PKT entering the logic unit LA 624 at the operating time at an instant. Accordingly, the data communication operation can be considered to operate at instances in time. In the illustrative embodiment, a first flit, or subpacket, of the packet PKT contains routing information through a switch to the target.

[0036] The logic unit LC 724 uses a control signal sent to the logic unit LA 624 to enforce priority over the logic unit LA 624 to send packets to logic unit LD 710.

[0037] A logic unit routes a packet based on packet header information and also based on control signals from other logic units.

[0038] The interconnect structure can further comprise a one-tick first-in-first-out (FIFO) buffer. A flit (subpacket) entering a logic unit passes through the one tick FIFO at the logic unit, regenerating the signal at each logic unit.

[0039] In some embodiments, the interconnect structure can operate so that, for the logic unit LA 624 positioned to send packets to a plurality of logic units including the logic unit LD 710, either Case 1 or Case 2 hold. In Case 1, the logic unit LA 624 determines that LD is the logic unit that is most appropriate to receive packet PKT, and either the logic unit LC 724 sends a packet to logic unit LD 710 and the logic unit LA 624 sends the PKT to a logic unit LG distinct from LD; or no logic unit with higher priority than the logic unit LA 624 to send packets to the logic unit LD 710 sends a packet to the logic unit LD 710 and the logic unit LA 624 sends the packet PKT to the logic unit LD 710. In Case 2, the logic unit LA 624 determines that sending the packet PKT to the logic unit LD 710 is unacceptable, and the logic unit LA 624 sends the packet PKT in the logic unit LG distinct from the logic unit LD 710 or to the logic unit LF 720 distinct from the logic unit LD 710.

[0040] For the logic unit LA 624 receiving a first subpacket of the packet PKT at a time T_s , if the logic unit LA 624 sends the first subpacket of the packet PKT to the logic unit LD 710, then logic unit LD 710 receives the first subpacket of packet PKT at a time T_{s+1} . If the logic unit LA 624 sends the first subpacket of the packet PKT to the logic unit LG, then the first subpacket passes

through a delay unit DA and arrives at the logic unit LG at a time T_{s+2} . If the logic unit LC 724 sends a first subpacket of a packet QKT to the logic unit LD 710 and the first subpacket of a packet QKT blocks packet PKT from traveling to the logic unit LD 710, then the subpacket QKT arrives at the logic unit LD 710 at time T_{s+1} .

[0041] In some embodiments, if the logic unit LA 624 determines that the logic unit LD 710 is a most appropriate logic unit to receive the packet PKT, then the logic unit LD 710 reaches that determination based on the routing information in the packet PKT. If the logic unit LA 624 determines that sending the packet PKT to the logic unit LD 710 is not acceptable, then the logic unit LD 710 reaches the determination based on the routing information in the packet PKT.

[0042] Referring to **FIGURE 7** in conjunction with **FIGURES 6** and **3B**, data and control lines that govern the high-speed parallel data path switching node 600 are described. Packets arrive at logical element LA on line 704 from another logic element on level W or on level W+1. At a given packet entry time T_s for logical element LA, exactly one of the following conditions is satisfied:

- 1) no first flit F_0 arrives at logic unit LA;
- 2) exactly one first flit arrives at logic unit LA from a logic element on level W, but no first flit F_0 arrives at logic unit LA from a logic element on level W+1; and
- 3) exactly one first flit arrives at logic unit LA from a node on level W+1, but no first flit F_0 arrives at logic unit LA from a node on level W.

[0043] Similarly, packets arrive at logic element LB from a logic element on level W+1 or from a logic element on level W. At logic node LB packet entry time T_s , either no first flit arrives at logic

element LB or exactly one first flit arrives at logic element LB. Importantly, given that a first flit F_0 of a packet PKT arrives at a logic element LA at time T_s , the next flit F_1 of PKT arrives at L at time T_{s+1} , followed by the other flits of PKT so that the last flit F_{Q-1} of PKT arrives at LA at time T_{s+Q-1} . Similarly, given that flit F_C (with $C < Q$) of PKT is in delay element DEL at time T_D then flit F_{c+1} of PKT is in delay element DEL at time T_{D+1} . Thus at each logical element and each delay element, the signal is reconstructed. This feature of the presently disclosed system, which is not shown in United States patent numbers 6,289,021 and 6,754,207, enables the switch chip clock to run faster than the clock in the systems depicted in United States patent numbers 6,289,021 and 6,754,207.

[0044] With continued reference to **FIGURE 7** in conjunction with **FIGURE 6**, the node 600 is on level W of the $N+1$ level switch. Define t_i to be $N - W$. Corresponding to the node 600, a Δ long binary sequence $BS = (b_0, b_1, \dots, b_{\Delta-1})$ exists such that each packet entering logic unit LA or logic unit LB of node 600 has a target address whose leading bits are BS. Each time that a packet PKT moves down a level in the switch, an additional bit of the target address of PKT is set. Each packet entering logic element LD has a target whose binary address has leading bits $(b_0, b_1, \dots, b_{\Delta-1}, 1)$. Each packet entering logic element LF has a target whose binary address has leading bits $(b_0, b_1, \dots, b_{\Delta-1}, 0)$. The bus 622 connects LA and LB to a logic element LD 710 in a node on level $W-1$ so that a packet P with first flit F_0 in logic unit LA or logic unit LB is able to progress to its target output through LD provided that the bit H_w of F_0 is equal to 1 and other control line conditions (discussed below) are satisfied. The bus 612 connects logic unit LA and logic unit LB to a logic element LF 720 in a node on level $W-1$ so that a packet PKT with first flit F_0 in logic unit LA or logic unit LB is able to progress to its target output through logic element LF provided that the bit H_w of F_0 is equal to 0 and other control line conditions (discussed below) are satisfied.

[0045] A logic element LC 724 exists in a LDM module 722 on level W-1 such that the logic element LC is positioned to send data to logic element LD 710 through delay unit DC. Also, a logic element LE 714 exists on level W-1 such that the logic element LE 714 is able to send data to LF 720 through delay unit DE. Suppose that T_S is a packet arrival time for logic elements LA 624 and LB 614. Then T_{S+1} is a packet arrival time at logic unit LF. A packet PKT traveling from logic element LE to LF must have its first flit F_0 in DE at time T_S and therefore must have its first flit in LE at time T_{S-1} . Similarly, a packet PKT traveling from LC to LD must have its first flit arrive in LC at time T_{S-1} . Therefore, T_{S-1} is a packet arrival time for both logic elements LC and LE.

[0046] The lack of static buffers in the switch can be compensated for a priority scheme for competing messages to travel to logic element LD or LF. The priority scheme gives highest priority to level W-1 packets and gives the bar setting (where the packet travels horizontally on the same path) of crossbar 602 priority over the cross setting (where the packet travels diagonally to an alternate path) of that switch. Therefore, the priority scheme for the first flits F_0 of packets entering LD 710 at time T_{S+1} is as follows:

- 1) A packet whose first flit F_0 is in DC at time T_S has priority one to travel to logic unit LD and such a packet will always arrive at logic unit LD at time T_{S+1} ;
- 2) A packet whose first flit F_0 is in logic unit LA 624 at time T_S and whose F_0 bit H_W is 1 has priority two and will travel through switch 602 set to the bar state to arrive at logic unit LD at time T_{S+1} , provided there is no priority one packet arriving at logic unit LD at time T_{S+1} ; and

- 3) A packet whose first flit F_0 is in logic unit **LB 614** at time T_s and whose F_0 bit H_w is 1 has priority three and will travel through switch **602** set to the cross state to arrive at logic unit **LD** at time T_{s+1} , provided that no priority one or priority two packet will arrive at logic unit **LD** at time T_{s+1} .

[0047] The priority scheme guarantees that lines **732** and **622** cannot possibly carry information at the same time. Therefore, the signals from those two lines can be joined in the multiplexer **MC** with no loss of fidelity. Notice that it is not necessary to designate a tick for multiplexer **MC**. A similar situation exists for multiplexer **ME**.

[0048] Similarly, the priority scheme for the first flits F_0 of packets entering **LF 720** at time T_{s+1} is as follows:

- 1) A packet whose first flit F_0 is in delay **DE** at time T_s has priority one to travel to logic **LF** and such a packet will always arrive at logic **LF** at time T_{s+1} ;
- 2) A packet whose first flit F_0 is in logic **LB 614** at time T_s and whose F_0 bit H_w is 0 has priority two and will travel through switch **602** set to the bar state to arrive at logic **LF** at time T_{s+1} , provided there is no priority one packet arriving at logic **LF** at time T_{s+1} ; and
- 3) A packet whose first flit F_a is in logic **LA 624** at time T_s and whose F_0 bit H_w of F_a is 0 has priority three and will travel through switch **602** set to the cross state to arrive at logic **LF** at time T_{s+1} , provided that no priority one or two packet will arrive at **LF** at time T_{s+1} .

[0049] Refer to **FIGURE 6** in conjunction with **FIGURE 7** and **FIGURE 8**. The integer W denotes the integer such that, in **FIGURE 7**, the logical elements **LA**, **LB**, **LG**, and **LH** are on level W of the switch and the logical elements **LC**, **LD**, **LE**, and **LF** are on level $W-1$ of the switch. The priority scheme is enforced by logic unit control registers **CR** and **CL** which are set by control

packets from logic or delay units in the LDM modules. Each of the logic elements in the parallel double-down switch contains two control registers CR (remote control) and CL (local control) and each control register contains two bits, thus allowing each register to store the binary representation of any one of the integers 0, 1, 2, or 3. T_s is packet time arrival at the logical elements LA and LB; T_{s+1} is an arrival time at LD, and LF; T_{s+2} is an arrival time at LG and LH. T_s is an arrival time at DC and DE and T_{s-1} is an arrival time at LC and LE. Prior to time T_{s-1} , the registers CR and CL in LA and LB are set to 0. The register CR in LA is set to a value distinct from 0 by a control signal on line 728 from LC. The register CR in LB is set to a value distinct from 0 by a control signal on line 718 from LE. The register CL in LB is set to a value distinct from 0 by a control signal on line 604 from LA. The register CL in LA is set to a value distinct from 0 by a control signal on line 606 from LB. The LDM module 620 is able to set the crossbar switch 602 to the bar or cross state by sending a signal from LA down line 608. The crossbar controlling LDM module 620 is referred to as the type α LDM module. The LDM module 610 has no means to control the crossbar 602 and is referred to as the type α LDM module. The crossbar has three states: state 0 indicates that the crossbar is not ready to receive data; state 1 indicates that the crossbar is in the bar state and is ready to receive data; state 2 indicates that the crossbar is in the cross state and is ready to receive data. If a packet flit arrives at the crossbar while the crossbar is in state 0, the flit will be stored in a flit wide buffer until the crossbar state is set to 1 or 2. There is a logic that keeps track when the last flit of a packet exits the crossbar. The state of the crossbar remains constant while flits of a given packet are passing through it. When the last flit exits the crossbar, the crossbar is placed in state 0.

[0050] In an illustrative embodiment, functionality of the control registers in logical element LA can be defined as follows:

- 1) $CR = 1$ implies logical element LO is blocked by a packet from logical element LC;
- 2) $CR = 2$ implies logical element LO is not blocked and can receive a packet from logical element LA;
- 3) $CL = 1$ implies logical element LB is sending a packet to logical element LF through the crossbar in the bar state;
- 4) $CL = 2$ implies logical element LF is not blocked and can receive a packet from logical element LA; and
- 5) $CL = 3$ implies logical element LF is blocked by a packet from logical element LE.

[0051] In an illustrative embodiment, functionality of the control registers in logical element LB can be defined as follows:

- 1) $CR = 1$ implies logical element LF is blocked by a packet from logical element LE;
- 2) $CR = 2$ implies logical element LF is not blocked and can receive a packet from logical element LB;
- 3) $CL = 1$ implies logical element LA is sending a packet to logical element LO through the crossbar in the bar state;
- 4) $CL = 2$ implies logical element LO is not blocked and can receive a packet from logical element LB; and
- 5) $CL = 3$ implies logical element LO is blocked by a packet from logical element LC.

[0052] The switch illustrated in **FIGURE 6**, **FIGURE 7**, and **FIGURE 8** can operate as follows:

[0053] In a first action, at or before packet arrival time T_s , the CR registers are set by a signal on line **728** from logical element LC and by a signal on line **718** from logical element LE.

[0054] In a second action, at packet arrival time T_s , the logical unit LA proceeds as follows:

- 1) Case 1: The first flit of a packet PKT_A arrives at logical element LA on line **704**; the header bit H_w of PKT_A is set to one indicating that logical element LD is on a path to the target output of PKT_A ; and the logical element LA register $CR = 2$ indicating that logical element LD will not receive a packet from logical element LC at time T_{s+1} . In this case, logical element LA sends a signal down line **608** to set the crossbar to the bar state. Then logical element LA sends the first flit of packet PKT_A through the crossbar to arrive at logical element LD at time T_{s+1} . Then logical element LA sends a signal through line **604** to set the CL register of logical element LB to 1.
- 2) Case 2: The conditions of Case 1 do not occur and the CR register of logical element LA is set to 2 indicating that logical element LC is not sending a packet to arrive at logical element LO at time T_{s+1} . In this case, logical element LA sends a signal through line **604** to set the CL register of logical element LB to 2.
- 3) Case 3: The conditions of Case 1 do not occur and the CR register of logical element LA is set to one indicating that logical element LC is sending a packet to arrive at logical element LO at time T_{s+1} . In this case, logical element LA sends a signal through line **604** to set the CL register of logical element LB to 3.

[0055] In a third action, at packet arrival time T_s , the logical unit LB proceeds as follows:

- 1) Case 1: The first flit of a packet PKT_s arrives at logical element LB on line **702**; the header bit H_w of PKT_s is set to zero and the logical element LB register $CR = 2$ indicating that logical element LF will not receive a packet from logical element LE at time T_{s+1} . In this case, logical element LB sends a first flit of PKT_s to the crossbar to travel through the crossbar after the crossbar has been set to the bar state so as to arrive at logical element LF

at time T_{s+1} . Then logical element LB sends a control signal through line **604** to set the CL register of logical element LA to 1.

- 2) Case 2: The conditions of Case 1 do not occur and the CR register in logical element LB is set to 2 indicating that logical element LE is not sending a packet to arrive at logical element LF at time T_{s+1} . In this case, logical element LB sends a signal through line **606** to set the CL register of logical element LA to 2.
- 3) Case 3: The conditions of Case 1 do not occur and the CR register in logical element LB is set to 3 indicating that logical element LE is sending a packet to arrive at logical element LF at time T_{s+1} . In this case, logical element LB sends a signal through line **606** to set the CL register of logical element LA to 3.

[0056] In a fourth action, if logical element LA has already set the crossbar to the bar state, then logical element LA takes no further action. If logical element LA has not set the crossbar to the bar state, then logical element LA examines its CL register after the CL register has been set to a non-zero value. If the CL register contains a 1, then logical element LA sets the crossbar to the bar state. If the CL register contains a number distinct from 1, then logical element LA sets the crossbar to the cross state.

[0057] In a fifth action, at this point the logic at logical element LA has information of the state of the crossbar and logical element LA proceeds as follows:

- 1) Case 1: There is no packet flit in logical element LA at time T_s , implying that no further action is required of logical element LA.
- 2) Case 2: The first flit of a packet PKT_A arrived at logical element LA at time T_s . The crossbar is in the bar state, the first flit of a packet PKT_A was sent through the crossbar as described hereinabove implying that no further action is required of logical element LA.

- 3) Case 3: The first flit of a packet PKT_A arrived at logical element LA at time T_s . The crossbar is in the bar state, the first flit of a packet PKT_A was not sent through the crossbar in the second action as described hereinabove implies that the first flit of packet PKT_A is to be sent to delay unit DA. Therefore, logical element LA sends the first flit of PKT_A to delay unit DA.
- 4) Case 4: The first flit of a packet PKT_A arrived at logical element LA at time T_s . The crossbar is in the cross state, the header bit H_w of PKT_A is set to 0, and the register CL of logical element LA is set to 2 then the first flit of PKT_A will be sent through the crossbar to arrive at logical element LF at time T_{s+1}
- 5) Case 5: The first flit of a packet PKT_A arrived at logical element LA at time T_s . The crossbar is in the cross state, but the conditions of case 4 do not occur. Then the first flit of PKT_A will be sent to delay unit DA.

[0058] In a sixth action, which can be performed simultaneously with fifth action, if either the CL register of logical element LB is set to 1, or LB sets the CL register of logical element LA to 1, then the logic at logical element LA has information that the crossbar is set to the bar state. If neither of these conditions is met, then logical element LA is aware that the crossbar is set to the cross state. Logical element LB proceeds as follows:

- 1) Case 1: No packet flit is in logical element LB at time T_s , implying that no further action is required of logical element LB.
- 2) Case 2: The first flit of a packet PKT_B arrived at logical element LB at time T_s . The crossbar is in the bar state, the first flit of a packet PKT_B was sent through the crossbar as described hereinabove, implying that no further action is required of logical element LB.

- 3) Case 3: The first flit of a packet PKT_B arrived at logical element LB at time T_s . The crossbar is in the bar state, the first flit of a packet PKT_B was not sent through the crossbar in the second action as described hereinabove, implying that the first flit of packet PKT_B is to be sent to delay unit DB.
- 4) Case 4: The first flit of a packet PKT_B arrived at logical element LB at time T_s . The crossbar is in the cross state, the header bit H_W of PKT_B is set to 1, and the register CL of logical element LB is set to 2 then the first flit of PKT_B will be sent through the crossbar to arrive at logical element LD at time T_{s+1} .
- 5) Case 5: The first flit of a packet PKT_B arrived at logical element LB at time T_s . The crossbar is in the cross state, but the conditions of case 4 do not occur. The first flit of PKT_A will be sent to delay unit DB.

[0059] In the illustrative example, the priority is given to the bar state over the cross state. In another example priority can be given to the cross state. In still another example priority can be given to logical element LA over logical element LB or to logical element LB over logical element LA.

[0060] The multiplexer elements improve structure compactness and performance by reducing the amount of interconnection paths between nodes. In a different embodiment, the multiplexers may be omitted. Referring to **FIGURE 7**, notice that MC **738** and interconnect line **734** can be eliminated by connecting interconnect line **732** to a first input of logic unit LD and connecting interconnect line **622** to a second input of logic unit LD. In another simplified embodiment, a single LDM module can serve as a node of a switch. In this case, the crossbar in the node switching node can be omitted. In another more complex embodiment, a switching node can include a

number on LDM modules and a switch where the number N of LDM modules is not equal to one or two and the switch is of radix N .

[0061] The structures and systems disclosed herein include significant improvements over the systems described in the referenced United States patent numbers 5,996,020, 6,289,021, and 6,754,207, including one or more of the following advantageous properties: 1) improved signal integrity even at high clock rates, 2) increased bandwidth, and 3) lower latency.

[0062] Improvements include one or more of: 1) a bus-wide data path; 2) all header bits sufficient to route data through the switch are contained in flit F_0 ; and 3) the signal is cleaned up at each logic unit and each delay unit of an LDM module.

[0063] **FIGUREs 9A, 9B and 9C** are schematic block diagrams that show interconnections of nodes on various levels of the interconnect structure. **FIGURE 9A** shows a node AR_J **920** on a ring R of outermost level J and the interconnections of node AR_J **920** to node BR_J **922**, device C **924**, node DR_J **926**, node $ER_{(J-1)}$ **928**, node $FR_{(J-1)}$ **930** and device G **932**. **FIGURE 9B** shows a node AR_T **940** on a ring R of a level J and the interconnections of node AR_T **940** to node BR_T **942**, node $CR_{(T+1)}$ **944**, node DR_T **946**, node $ER_{(T-1)}$ **948**, node $FR_{(T-1)}$ **950** and node $GR_{(T+1)}$ **952**. **FIGURE 9C** shows a node AR_0 **960** on a ring R of innermost level 0 and the interconnections of node AR_0 **960** to node BR_0 **962**, node CR_1 **964**, node DR_0 **966**, device E **968** and node GR_1 **972**.

[0064] **FIGUREs 9A, 9B and 9C** show topology of an interconnect structure. To facilitate understanding, the structure can be considered a collection of concentric cylinders in three dimensions r , Θ , and z . Each node or device has a location designated (r, Θ, z) which relates to a position $(r, 2\pi, \Theta/K, z)$ in three-dimensional cylindrical coordinates where radius r is an integer which specifies the cylinder number from 0 to J , angle Θ is an integer which specifies the spacing

of nodes around the circular cross-section of a cylinder from 0 to $K-1$, and height z is a binary integer which specifies distance along the z -axis from 0 to 2^{J-1} . Height z is expressed as a binary number because the interconnection between nodes in the z -dimension is most easily described as a manipulation of binary digits. Accordingly, an interconnect structure can be defined with respect to two design parameters J and K .

[0065] In **FIGURES 9A, 9B** and **9C** interconnections are shown with solid lines with arrows indicating the direction of message data flow and dashed lines with arrows indicating the direction of control message flow. In summary, for nodes A, B and D and nodes or devices C, E, F, G:

- 1) A is on level $t=r$;
- 2) B and C send data to A;
- 3) D and E receive data from A;
- 4) F sends a control signal to A;
- 5) G receives a control signal from A;
- 6) B and D are on level T ;
- 7) B is the immediate predecessor of A;
- 8) D is the immediate successor to A; and
- 9) C, E, F and G are not on level T .

[0066] Positions in three-dimensional cylindrical notation of the various nodes and devices is as follows:

- 1) A is positioned at node $N(r, \Theta, z)$;
- 2) B is positioned at node $N(r, \Theta - 1, H_T(z))$;
- 3) C is either positioned at node $N(r+1, \Theta - 1, z)$ or is outside the interconnect structure;

- 4) D is positioned at node $N(r, \ominus + 1, h_T(z))$;
- 5) E is either positioned at node $N(r-1, \ominus + 1, z)$ or is outside the interconnect structure and the same as device F;
- 6) F is either positioned at node $N(r-1, \ominus, H_{T-1}(z))$ or is outside the interconnect structure and the same as device E;
- 7) G is either positioned at node $N(r+1, \ominus, h_T(z))$ or is outside the interconnect structure.

[0067] Note that the terms $\ominus + 1$ and $\ominus - 1$ refer to addition and subtraction, respectively, modulus K.

[0068] In this notation, $(\ominus - 1) \bmod K$ is equal K when \ominus is equal to 0 and equal to $\ominus - 1$ otherwise. The conversion of z to $H_r(z)$ on a level r is described for $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_{r-1}, \dots, z_2, z_1, z_0]$ by reversing the order of low-order z bits from z_{r-1} to z_0 into the form $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_0, z_1, z_2, \dots, z_{r-1}]$, subtracting one (modulus 2^r) and reversing back the low-order z bits. Similarly, $(\ominus + 1) \bmod K$ is equal 0 when \ominus is equal to $K-1$ and equal to $\ominus + 1$ otherwise. The conversion of z to $h_r(z)$ on a level r is described for $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_{r-1}, \dots, z_2, z_1, z_0]$ by reversing the order of low-order z bits from z_{r-1} to z_0 into the form $z = [z_{J-1}, z_{J-2}, \dots, z_r, z_0, z_1, z_2, \dots, z_{r-1}]$, adding one (modulus 2^r) and reversing back the low-order z bits.

[0069] In accordance with one embodiment of the system depicted in **FIGURES 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes arranged in a structure comprising a hierarchy of levels from a source level to a destination level; plurality of nodes spanning a cross-section of a level; a plurality of nodes in a cross-section span. The level of a node can be determined entirely by the position of the node in the structure; and a plurality of interconnect lines coupling the nodes in the structure. For a node N on a level L: (1) a plurality of message input interconnect lines are coupled to a node on a previous level L + 1; (2) a plurality of message input interconnect

lines are coupled to a node on the level L; (3) a plurality of message output interconnect lines are coupled to a node on the level L; (4) a plurality of message output interconnect lines are coupled to a node on a subsequent level L-1; (5) a control input interconnect line is coupled to the message output interconnect line of a node on the level L-1; and (6) a switch is coupled to receive a message on the control input interconnect line and, in accordance with the message, to selectively transmit a message without buffering on the plurality of message output interconnect lines coupled to the subsequent level L-1 node or on the plurality of message output interconnect lines coupled to the level L.

[0070] In accordance with another embodiment of the system depicted in **FIGURES 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes and a plurality of interconnect lines coupling the nodes. A node X of the plurality of nodes can include a plurality of message input interconnect lines coupled to a node A distinct from the node X; and a plurality of message input interconnect lines coupled to a node B distinct from the node A and the node X. The node X accepts a message input from the node A and a message input from the node B with a control signal communicating between the node A and the node B for determining a priority relationship between conflicting messages. The control signal can enforce the priority relationship between the sending of a message from the node A to the node X and the sending of a message from the node B to the node X.

[0071] In accordance with a further embodiment of the system depicted in **FIGURES 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes and a plurality of interconnect lines in an interconnect structure selectively coupling the nodes in a hierarchical multiple level structure. The hierarchical multiple level structure can be arranged to include a plurality of J+i levels in an hierarchy of levels arranged from a lowest destination level L_0 to a highest level L_J which is farthest

from the lowest destination level L_0 . The level of a node can be determined entirely by the position of the node in the structure with the interconnect structure transmitting a message M in a plurality of discrete time steps. A message M moving in a time step and the interconnect structure having interconnections to move the message M in one of three ways in the time step including: (1) the message M enters a node in the interconnect structure from a device external to the interconnect structure; (2) the message M exits the interconnect structure to a designated output buffer; and (3) the message M either moves from a node U on a level L_k to a different node V on the same level L_k or moves from the node U to a node W on a level L_i ; where k is greater than i so that the level L_i is closer to the destination level L_0 than the level L_k .

[0072] In accordance with other embodiments of the system depicted in **FIGUREs 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes; and a plurality of interconnect lines selectively coupling the nodes in a hierarchical multiple level structure with the level of a node being determined entirely by the position of the node in the structure in which data moves only unilaterally from a source level to a destination level or laterally along a level of the multiple level structure. Data messages can be transmitted through the multiple level structure from a source node to a designated destination node. A level of the multiple levels can include one or more groups of nodes. The data message can be transmitted to a group of the one or more groups of nodes that is en route to the destination node. A group of the one or more groups can include a plurality of nodes. The data message can be transmitted to a node N of the plurality of nodes of a group unilaterally toward the destination level if the node is not blocked and otherwise the data message being transmitted laterally if the node is blocked.

[0073] In accordance with further embodiments of the system depicted in **FIGUREs 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes and a plurality of interconnect lines

L interconnecting communication devices at the plurality of nodes. The nodes can include communication devices that communicate messages in a sequence of discrete time steps including receiving messages and sending messages. A node N of the plurality of nodes can include: (1) a connection to a plurality of interconnect lines LUN for transmitting a message from a device U to the node N; (2) a connection to a plurality of interconnect lines L VN for transmitting a message from a device V to the node N; and (3) the network which has a precedence relationship PN(U,V) relating to the node N and the devices U and V such that the device U has precedence over the device V in sending a message to the node N so that for a message MU at the device U that is directed to the node N via the plurality of interconnect lines LUN at a time step t and a message MV at the device V that is directed to the node N via the plurality of interconnect lines LVN also at a time step t. The message MU is successfully sent to the node N and the node V uses a control signal to decide where to send the message MV.

[0074] In accordance with still further embodiments of the system depicted in **FIGUREs 9A, 9B** and **9C**, the interconnect structure can include a plurality of nodes N and a plurality of interconnect lines L connecting the plurality of nodes N in a predetermined pattern. The interconnect lines carry messages M and control signals C. The messages M and control signals C can be received by a node of the plurality of nodes at a discrete time step t and the messages M can be moved to subsequent nodes of the plurality of nodes in an immediately subsequent discrete time step t+1. The plurality of interconnect lines L connecting the plurality of nodes N can include: (1) a node NA having a message input interconnection for receiving a message MA, (2) a control input interconnection for receiving a control signal CA, (3) a direct message output interconnection to a node ND, (4) a direct message output interconnection to a node NE, (5) a direct control output interconnection to a device G. A control logic for determining whether the message MA is sent to

the node ND or the node NE can be based on: (1) the control signal CA; (2) a location of the node NA within the plurality of interconnect lines L; and (3) a routing information contained in the message MA.

[0075] In still another embodiment, the interconnect structure can comprise a plurality of nodes N and a plurality of interconnect lines L connecting the plurality of nodes N in a predetermined pattern. The plurality of interconnect lines L connecting the plurality of nodes N can include a node NA having a direct message input interconnection for receiving a message MA and having a plurality of direct message output interconnections for transmitting the message MA to a plurality of nodes including a selected node NP being most desired for receiving the message MA. The selected node NP can be determined only by routing information in a header of the message MA and the position of the node NA within the plurality of interconnect lines L. The selected node NP has a plurality of direct message input interconnections for receiving a message MP from a plurality of nodes including a priority node NB which has priority for sending a message to the selected node NP. The priority node NB can be determined by position of the node NB within the plurality of interconnect lines L so that: (1) if the node NA is the same as the node NB, then the message MA is the message MP and is sent from the node NA to the node NP; and (2) if the node NA is not the same as the node NB and the node NB directs a message MB to the node NP, then the message MB is sent from the node NB to the node NP.

[0076] In additional embodiments, the interconnect structure can comprise a network capable of carrying a plurality of messages M concurrently comprising a plurality of output ports P; a plurality of nodes N, the individual nodes N including a plurality of direct message input interconnections and a plurality of direct message output interconnections; and a plurality of interconnect lines. The individual nodes N pass messages M to predetermined output ports of the plurality of output ports

P. The predetermined output ports P are designated by the messages M . The plurality of interconnect lines can be configured in an interconnect structure selectively coupling the nodes in a hierarchical multiple level structure arranged to include a plurality of $J+1$ levels in an hierarchy of levels arranged from a lowest destination level LO to a highest level LJ which is farthest from the lowest destination level LO , the output ports P being connected to nodes at the lowest destination level LO . The level of a node can be determined entirely by the position of the node in the structure. The network can include a node NA of the plurality of nodes N , a control signal operating to limit the number of messages that are allowed to be sent to the node NA to eliminate contention for the predetermined output ports of the node NA so that the messages M are sent through the direct message output connections of the node NA to nodes NH that are a level L no higher than the level of the node NA , the nodes NH being on a path to the designated predetermined output ports P of the messages M .

[0077] In accordance with an embodiment of the system depicted in **FIGURES 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes and a plurality of interconnect lines in an interconnect structure selectively coupling the nodes in a hierarchical multiple level structure. The multiple level structure can be arranged to include a plurality of $J+1$ levels with J an integer greater than 0 in an hierarchy of levels arranged from a lowest destination level $L0$ to a highest level LJ with the level of a node being determined entirely by the position of the node in the structure, the interconnect structure transmitting a plurality of multiple-bit messages entering the interconnect structure unsorted through a plurality of input ports. An individual message M of the plurality of messages can be self-routing. The individual message M moves in a plurality of ways including three ways which are sufficient for the message M to exit the interconnect structure through an output port designated by the message M . The three ways are: (1) the message M enters

a node in the interconnect structure from a device external to the interconnect structure, the message M designating one or more designated output ports; (2) the message M moves through a node in the interconnect structure without buffering to a designated output port; and (3) the message M moves either through a node U on a level L_k of the interconnect structure without buffering to a different node V on the same level L_k or moves through the node U on a level L_k of the interconnect structure without buffering to a node W on a level L_i nearer in the hierarchy to the destination level L_0 than the level L_k .

[0078] In accordance with still other embodiments of the system depicted in **FIGUREs 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes and a plurality of interconnect lines in an interconnect structure selectively coupling the nodes in a structure. The interconnect structure transmits a plurality of multiple-bit messages entering the interconnect structure unsorted through a plurality of input ports. An individual message M of the plurality of messages can be self-routing. The interconnect structure can include: (1) a node NE having a first data input interconnection from a node NA and a second data input interconnection from a node NF distinct from the node NA; and (2) a control interconnection between the node NA and node NF for carrying a control signal to resolve contention for sending messages to the node NE. The control signal can be supplied from the node NA or the node NF each distinct from the node NE with which messages are communicated.

[0079] In accordance with further other embodiments of the system depicted in **FIGUREs 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes and a plurality of interconnect lines in an interconnect structure selectively coupling the nodes in a hierarchical multiple level structure. The multiple level structure can be arranged to include a plurality of $J+1$ levels with J an integer greater than 0 in an hierarchy of levels arranged from a lowest destination level L_0 to a

highest level L_j . The interconnect structure receives a plurality of multiple bit messages unsorted through a plurality of input ports and transmitting the multiple bit messages. An individual message M of the plurality of messages can be self-routing and moves through nodes using wormhole routing in which only a portion of the multiple-bits of a message are in transit between two nodes. The multiple-bit message extends among multiple nodes. An individual message M moves in a plurality of ways including four ways which are sufficient for the message M to exit the interconnect structure through an output port designated by the message M . The four ways are: (1) the message M enters a node in the interconnect structure from a device external to the interconnect structure, the message M designating one or more designated output ports; (2) the message M moves through a node in the interconnect structure without buffering to a designated output port; (3) the message M moves through a node on a level L_k of the interconnect structure without buffering to a different node on the same level L_k ; and (4) the message M moves through a node on a level L_k of the interconnect structure without buffering to a node on a level L_i nearer in the hierarchy to the destination level L_o than the level L_k .

[0080] In accordance with further other embodiments of the system depicted in **FIGUREs 9A, 9B and 9C**, the interconnect structure can include a plurality of nodes and a plurality of interconnect lines in an interconnect structure selectively coupling the nodes in a structure. The interconnect structure receives a plurality of multiple bit messages unsorted through a plurality of input ports and transmits the multiple bit messages. An individual message M of the plurality of messages can be self-routing and move through nodes using wormhole routing in which only a portion of the multiple bits of a message are in transit between two nodes with the multiple-bit message extending among multiple nodes. The interconnect structure can include a node N_E that has a first data input interconnection from a node N_A and a second data input interconnection from a node N_F ; and a

control interconnection between the node N_A and node N_F that resolves contention for sending messages to the node N_E .

[0081] In accordance with further other embodiments of the system depicted in **FIGUREs 9A, 9B** and **9C**, the interconnect structure can include a plurality of nodes with each node having a plurality of input ports and a plurality of output ports, a logic associated with the plurality of nodes, and a node X included in the plurality of nodes and having an output port opx . The node X can have a set of input terminals IPX wherein a logic associated with the node X can send messages entering an input terminal of the set IPX to the output port opx . The logic associated with the node X can be operational wherein if a message M_P arrives at an input port p of input port set IPX and a path exists from the output port opx to a target of message M_P , then one of the messages arriving at input port set IPX will be sent to output port opx so long as the output port opx is not blocked by a message that is not traveling through the node X .

[0082] In still further embodiments, the interconnect structure can comprise a plurality of interconnected nodes including distinct nodes F_w , F_B , and F_x ; means for sending a plurality of messages through the plurality of nodes including sending a set of messages S_w through the node F_w ; and means for sending information I concerning routing of the messages in the message set S_w through the node F_w including routing a portion of the messages in the message set S_w through the node F_w to the node F_x . The interconnect structure can further comprise means associated with the node F_s for using the information I to route messages through the node F_B .

[0083] In other embodiments, the interconnect structure can comprise a plurality of nodes including a node X , a node set T , and a node set S including nodes Y and Z ; a plurality of interconnect paths connecting the nodes; a plurality of output ports coupled to the plurality of nodes; and logic that controls flow of data through the nodes to the output ports. The logic controls

data flow such that: (1) the node X is capable of sending data to any node in the set S; (2) the node set T includes nodes that can alternatively pass data that are otherwise controlled by the logic to flow through the node X; (3) any output port that can access data passing through the node X can also access data passing through the node Y; (4) the plurality of output ports include an output port O that can access data passing through the node X but cannot access data passing through the node Z; and (5) the logic controls flow of data through the node X to maximize the number of data messages that are sent through a node in the set T such that the number of output ports accessible from the node in the set T is less than the number of output ports that are accessible from the node X.

[0084] Referring to **FIGURE 10**, a timing diagram illustrates timing of message communication in the described interconnect structure. In various embodiments of the interconnect structure, control of message communication is determined by timing of message arrival at a node. A message packet, such as a packet **1100** shown in **FIGURE 11**, includes a header **1110** and a payload **1120**. The header **1110** includes a series of bits **1112** designating the target ring in a binary form. When a source device $CU(\Theta_1, z_1)$ at an angle Θ_1 , and height z_1 , sends a message packet M to a destination device $CU(\Theta_2, z_2)$ at an angle Θ_2 and height z_2 , the bits **1112** of header **1110** are set to the binary representation of height z_2 .

[0085] A global clock servicing an entire interconnect structure keeps integral time modulus K where, again, K designates the number of nodes n at a cylinder height z. There are two constants α and β such that the duration of α exceeds the duration of β and the following five conditions are met. First, the amount of time for a message M to exit a node $N(T, \Theta + 1, hT(z))$ on level T after exiting a node $N(T, \Theta, z)$ also on level T is α . Second, the amount of time for a message M to exit a node $N(T-1, \Theta + 1, z)$ on level T-1 after exiting a node $N(T, \Theta, z)$ on level T is $\alpha - \beta$. Third, the

amount of time for a message to travel from a device CU to a node $N(r, \Theta, z)$ is $\alpha\beta$. Fourth, when a message M moves from a node $N(r, \Theta, z)$ to a node $N(r, \Theta + 1, h_r(z))$ in time duration a , the message M also causes a control code to be sent from node $N(r, \Theta, z)$ to a node $N(r+1, \Theta, h_r(z))$ to deflect messages on the outer level $r+1$. The time that elapses from the time that message M enters node $N(r, \Theta, z)$ until the control bit arrives at node $N(r+1, \Theta, h_{r+1}(z))$ is time duration β . The aforementioned fourth condition also is applicable when a message M moves from a node $N(J, \Theta, z)$ to a node $N(J, \Theta + 1, h_J(z))$ at the outermost level J so that the message M also causes a control code to be sent from node $N(J, \Theta, z)$ to the device D outside of the network such that D sends data to $N(J, \Theta + 1, h_J(z))$. In one embodiment, $D = CU(\Theta + 1, h_J(z))$. The time that elapses from the time that message M enters node $N(r, \Theta, z)$ until the control bit arrives at device $CU(\Theta, z)$ is time duration 13. Fifth, the global clock generates timing pulses at a rate of α .

[0086] When the source device $CU(\Theta_1, z_1)$ sends a message packet M to the destination device $CU(\Theta_2, z_2)$, the message packet M is sent from a data output terminal of device $CU(\Theta_1, z_1)$ to a data input terminal of node $N(J, \Theta_1, z_1)$ at the outermost level J. Message packets and control bits enter nodes $N(T, \Theta, z)$ on a level T at times having the form $n\alpha + L\beta$ where n is a positive integer. The message M from device $CU(\Theta_1, z_1)$ is sent to the data input terminal of node $N(J, \Theta_1, z_1)$ at a time $t_0 - \beta$ and is inserted into the data input terminal of node $N(J, \Theta_1, z_1)$ at time t_0 so long as the node $N(J, \Theta_1, z_1)$ is not blocked by a control bit resulting from a message traversing on the level J. Time t_0 has the form $(\Theta_2 - \Theta_1) \alpha + \beta$. Similarly, there is a time of the form $(\Theta_2 - \Theta_1) \alpha + J\beta$ at which a data input terminal of node $N(J, \Theta_1, z_1)$ is receptive to message packet from device $CU(\Theta_1, z_1)$.

[0087] Nodes $N(T, \Theta, z)$ include logic that controls routing of messages based on the target address of a message packet M and timing signals from other nodes. A first logic switch (not shown) of node $N(T, \Theta, z)$ determines whether the message packet M is to proceed to a node $N(T-1, \Theta+1, z)$

on the next level $T-1$ or whether the node $N(T-1, \Theta+1, z)$ is blocked. The first logic switch of node $N(T, \Theta, z)$ is set according to whether a single-bit blocking control code sent from node $N(T-1, \Theta, H_{T-1}(z))$ arrives at node $N(T, \Theta, z)$ at a time t_0 . For example, in some embodiments the first logic switch takes a logic 1 value when a node $N(T-1, \Theta+1, z)$ is blocked and a logic 0 value otherwise. A second logic switch (not shown) of node $N(T, \Theta, z)$ determines whether the message packet M is to proceed to a node $N(T-1, \Theta+1, z)$ on the next level $T-1$ or whether the node $N(T-1, \Theta+1, z)$ is not in a suitable path for accessing the destination device $CU(\Theta_2, z_2)$ of the header of the message packet M . The header of the message packet M includes the binary representation of destination height $z_2(z_2(J), z_2(J-1), \dots, z_2(T), \dots, z_2(1), z_2(0))$. The node $N(T, \Theta, z)$ on level T includes a single-bit designation z_T of the height designation $z(z_J, z_{J-1}, \dots, z_T, \dots, z_1, z_0)$. In this embodiment, when the first logic switch has a logic 0 value and the bit designation $z_2(T)$ of the designation height is equal to the height designation z_T , then the message packet M proceeds to the next level at node $N(T-1, \Theta+1, z)$ and the destination height bit $z_2(T)$ is stripped from the header of message packet M . Otherwise, the message packet M traverses on the same level T to node $N(T, \Theta+1, h_T(z))$. If message packet M proceeds to node $N(T-1, \Theta+1, z)$, then message packet M arrives at a time $t_0 + (\alpha - \beta)$ which is equal to a time $(z_2 - z_1 + 1) \alpha + (J-1) \beta$. If message packet M traverses to node $N(T, \Theta+1, h_T(z))$, then message packet M arrives at a time $t_0 + \alpha$, which is equal to a time $(z_2 - z_1 + 1) \alpha + J \beta$. As message packet M is sent from node $N(r, \Theta, z)$ to node $N(T, \Theta+1, h_T(z))$, a single-bit control code is sent to node $N(T+1, \Theta, h_T(z))$ (or device $CU(\Theta, z)$ which arrives at time $t_0 + \beta$. This timing scheme is continued throughout the interconnect structure, maintaining synchrony as message packets are advanced and deflected.

[0088] The message packet M reaches level zero at the designated destination height z_2 . Furthermore, the message packet M reaches the targeted destination device $CU(\Theta_2, z_2)$ at a time

zero modulus K (the number of nodes at a height z). If the targeted destination device $CU(\Theta_2, z_2)$ is ready to accept the message packet M , an input port is activated at time zero modulus K to accept the packet. Advantageously, all routing control operations are achieved by comparing two bits, without ever comparing two multiple-bit values. Further advantageously, at the exit point of the interconnect structure as message packets proceed from the nodes to the devices, there is no comparison logic. If a device is prepared to accept a message, the message enters the device via a clock-controlled gate.

[0089] An embodiment of the interconnect structure can comprise a plurality of nodes arranged in a topology of three dimensions and means for transmitting a message from a node N to a target destination. The means for transmitting a message from a node N to a target destination can comprise means for determining whether a node en route to the target destination in the second and third dimensions and advancing one level toward the destination level of the first dimension is blocked by another message; and means for advancing the message one level toward the destination level of the first dimension when the en route node is not blocked, and means for moving the message in the second and third dimensions along a constant level in the first dimension otherwise. The means for transmitting a message from a node N to a target destination can further comprise means for specifying the first dimension to describe a plurality of levels, the second dimension to describe a plurality of nodes spanning a cross-section of a level, and the third dimension to describe a plurality of nodes in the cross-section of a level; means for sending a control signal from a node on the level of the en route node to the node N in the first dimension, the control signal specifying whether the node en route is blocked; means for timing transmission of a message using a global clock specifying timing intervals to keep integral time modulus the number of nodes in a cross-section of a level; and means for setting a first time

interval α for moving the message in the second and third dimensions. The means for transmitting a message from a node N to a target destination can still further comprise means for setting a second time interval α - β for advancing the message one level toward the destination level, the global clock specifying a global time interval equal to the second time interval, the first time interval being smaller than the global time interval; and means for setting a third time interval for sending the control signal from the node on the level of the en route node to the node N, the third time interval being equal to β .

[0090] In **FIGURE 12**, a single chip **1200** contains two components. The first component is a Data Vortex switch **1220** as described in all previous figures, and the second component is an array **1240** of processing cores. Refer to **FIGURE 13**. The Data Vortex switch receives data from an external source on line **1210**. The processing cores **1250** of an array **1240** of processing cores can receive data from the Data Vortex switch on line **1230**. A sending processing core **1250** in the array **1240** of processing cores can send data to a receiving core in the array of processing cores by forming a packet whose header indicates the location of the receiving core and whose payload indicates the data to be sent. In **FIGURE 13**, this packet is sent down line **1310** and enters the Data Vortex switch **1220**. The Data Vortex Switch routes the packet to the receiving core on line **1230**. A core in the array **1240** of processing cores can also send a packet to an address, not pictured, in **FIGURE 13** on an output line **1260**.

[0091] Refer to Figure 14 that pictures four processor core arrays **1240**. In general, there can be an arbitrary number of processor core arrays. Conversely, with present-day crossbar or mesh topologies, there is an upper limit of processor core arrays given the limited number of end points of those topologies.

[0092] A sending processing core in an array **1240** of processing cores can send data to a receiving core in one of the arrays of processing cores by forming a packet whose header indicates the location of the receiving core and whose payload indicates the data to be sent. This packet is sent down line **1450** and enters the Data Vortex switch **1410**. The Data Vortex switch **1410** routes the packet to the receiving core first by routing the packet to the processing core array containing the receiving processing core. The Data Vortex Switch **1220** routes the packet to the receiving processing core in a processor core array **1240**. Since the Data Vortex Switches **1410** and **1220** are not crossbars, the switches do not need to be globally set and reset as different group of packets enter the switches. In present-day technology as the number of inputs in a crossbar switch increase the time to set the switch increases as a function of the number of inputs. This setting problem in other technologies causes long packets. There is no setting of the Data Vortex switch as packets simply enter and leave advantageously.

[0093] The number of Data Vortex switches **1410** that need to be deployed depends on the total number of processor cores and the bandwidths of the transmission lines.

[0094] In case a processing core in a given processing core array **1240** sends a packet PKT to another processing core in its same array, the sent packet passes through the Data Vortex Switch **1410** where it travels with other packets passing through the system. This shuffling of packets provides a randomness that has proven effective in other hardware systems enabled by Data Vortex switches. This is advantageous because, unlike chips and systems that connect processor cores using a crossbar or mesh, there is fine-grained parallelism. Fine-grained parallelism allows for

short packet movement (no longer than a cache line) that avoids congestion. This is ideal for applications that require small data packets.

[0095] It is an important fact that there can be a large number of chips **1200** on a silicon substrate **1400** it is not necessary for packets traveling between these chips to pass through SerDes modules. In present-day hardware, data packets travelling through SerDes modules add significant latency. Packets traveling between chips **1200** on a silicon substrate **1400** won't suffer from this latency as there are no SerDes modules at the edges of the chips **1200**.

[0096] In case the modules **1200** are placed on a printed circuit board **1400**, then packets traveling from one module **1200**, using line **1450** through Data Vortex switch **1410** then through line **1440** must travel through SerDes modules on each chip boundary. Even though in this implementation, packets suffer from the latency caused by the SerDes modules, the system still benefits from the increased number of cores, the shorter packet lengths, and the fine-grain parallelism enabled by the Data Vortex switch.

[0097] The plurality of processor core arrays **1240** allows for a greater total number of processing cores and allows for each of the cores to be a larger size. In present day technology using crossbars, as the number of cores goes up, the packets sizes go up. Using the Data Vortex as described in **Figures 12, 13, & 14** as the number of cores go up, the packet sizes remain the same.

[0098] In other embodiments of **Figure 14**, memory controllers adjacent to the processor core arrays are accessed through the Data Vortex switch **1410**.

[0099] There are numerous advantages to placing a Data Vortex network and processor array on the same module (e.g. a silicon substrate). Doing so removes the serializer/deserializer block (“SerDes”) from the Data Vortex path, thus reducing the power required and the latency. Present day Data Vortex-enabled systems are also slowed down by commodity network-on-chips. Having a Data Vortex network on the same module replaces those traditional network-on-chips (NoCs) and allows an entire system to benefit from all of the advantages of the Data Vortex topology (i.e. congestion-free, small packet movement throughout the entire ecosystem). Non-Data Vortex NoCs can therefore be removed from the core to core data path and therefore packets can remain small compared to prior art where packets travelling off commodity microprocessors get broken apart as they go through off-chip Data Vortex networks. This also provides more consistent core to core or core to memory latencies compared to present Data Vortex-enabled systems. On the next level up (board to board), an on-Module Data Vortex network provides a common socket to socket and core to core architecture across an entire system, removing the necessity for different topologies within the same system. All of this enables a common programming model across the cores, sockets, and servers, making it easier for the end user.

[0100] Terms "substantially", "essentially", or "approximately", that may be used herein, relate to an industry-accepted variability to the corresponding term. Such an industry-accepted variability ranges from less than one percent to twenty percent and corresponds to, but is not limited to, materials, shapes, sizes, functionality, values, process variations, and the like. The term "coupled", as may be used herein, includes direct coupling and indirect coupling via another component or element where, for indirect coupling, the intervening component or element does not modify the operation. Inferred coupling, for example where one element is coupled to another element by

inference, includes direct and indirect coupling between two elements in the same manner as "coupled".

[0101] The illustrative pictorial diagrams depict structures and process actions in a manufacturing process. Although the particular examples illustrate specific structures and process acts, many alternative implementations are possible and commonly made by simple design choice. Manufacturing actions may be executed in different order from the specific description herein, based on considerations of function, purpose, conformance to standard, legacy structure, and the like.

[0102] While the present disclosure describes various embodiments, these embodiments are to be understood as illustrative and do not limit the claim scope. Many variations, modifications, additions and improvements of the described embodiments are possible. For example, those having ordinary skill in the art will readily implement the steps necessary to provide the structures and methods disclosed herein, and will understand that the process parameters, materials, shapes, and dimensions are given by way of example only. The parameters, materials, and dimensions can be varied to achieve the desired structure as well as modifications, which are within the scope of the claims.

WHAT IS CLAIMED IS:

- 1) An interconnect apparatus configured to communicate data packets through a network, the data packets arranged in a plurality of sub-packets, including an address sub-packet that identifies a target processing core for receiving a data packet, the interconnect apparatus comprising:
 - a Data Vortex switch on a chip, and
 - an array of processing cores also on the e chip , wherein the Data Vortex switch receives data from an external source and the array of processing cores receives data from the Data Vortex switch;and
 - wherein a data packet sent to the Data Vortex switch is not reassembled in the Data Vortex switch but travels through the Data Vortex switch in wormhole fashion.
- 2) An interconnect apparatus in accordance with claim 1 wherein the array of processing cores includes the target core and a sending core wherein the sending core can send a data packet to the target core by forming the data packet with an address header that includes an address for a target core's location and a payload with data to be sent to the target core.
- 3) An interconnect apparatus in accordance with claim 2 wherein the data packet from the sending core is sent from the sending core to the Data Vortex switch and the Data Vortex switch routes the data packet to the target core based on the address header.
- 4) An interconnect apparatus in accordance with claim 3 wherein bits included in said address header are contained in a flit and travel through said Data Vortex switch in a bus-wide data path.
- 5) An interconnect apparatus in accordance with claim 4 wherein the Data Vortex switch is used for core-to-core communication either between two or more processing cores on a processor core array on the same chip as the Data Vortex switch or between two or

more processing cores on a processor core array on a different chip than the Data Vortex switch.

6) An interconnect apparatus in accordance with claim 5 wherein data packets sent from a sending processor core to the Data Vortex switch go through one or two paths depending on the target core destination;

i. if a sending processor core in an array on the same chip as the Data Vortex switch sends a data packet to a target processor core in the same array on the same chip, the data packet is sent on a path to the Data Vortex switch and back down to the target processor core in same array as the sending processor core;

ii. if a sending processor core on a chip sends a data packet to a target processor core in a different array on a different chip the packet is sent on a path from the sending processor core to the Data Vortex switch on the chip that includes the target processing core, and the Data Vortex switch sends the data packet to the target processing core.

7) An interconnect apparatus in accordance with claim 6 wherein the number of processor cores on a single processing chip can range from 1 to N processing cores in an array on a processing chip communicating with 1 to K processing cores where $K > N$.

8) An interconnect apparatus configured to communicate data packets through a network, the data packets arranged in a plurality of sub-packets, including an address sub-packet that identifies a target processing core for receiving a data packet, the interconnect apparatus comprising;

a plurality of chips, each chip including a Data Vortex switch and an array of processing cores, and

a master Data Vortex switch connected to each Data Vortex switch and each array of processing cores on each of said plurality of chips, wherein the master Data Vortex switch

communicates data packets between each Data Vortex switch or each array of processing cores;

and

wherein a data packet sent to the Data Vortex switch is not reassembled in the Data Vortex switch but travels through the Data Vortex switch in wormhole fashion.

- 9) An interconnect apparatus in accordance with claim 8, wherein a sending processor core in any one of the array of processing cores included in any one of the plurality of chips can send a data packet to a target processing core included in an array of processing cores in another one of the plurality of chips by forming a core packet for the sending processing core having a header that identifies a location for the target processing core and a payload including the data to be sent.
- 10) An interconnect apparatus in accordance with claim 9, wherein said data packet is sent from the sending processing core to the master Data Vortex switch, from the master Data Vortex switch to the Data Vortex switch on the chip including the target processing core and from the Data Vortex switch on the chip including the target processing core to the target processing core.
- 11) An interconnect apparatus in accordance with claim 9, wherein said data packet is sent from the sending processor core to the master Data Vortex switch, and from the master Data Vortex switch to the target processing core.
- 12) An interconnect apparatus in accordance with claim 9, wherein said data packet is sent from the sending processor core to the Data Vortex switch on the same chip as the sending processing core, and from the Data Vortex switch on the same chip as the sending processing core to the master Data Vortex switch which can send the data packet either to the Data Vortex switch on the same chip as the target processing core or directly to the target processing core.
- 13) An interconnect apparatus in accordance with claim 9, wherein said data packet is sent from the sending processor core to the Data Vortex switch on the same processing core as

the sending processor core and from that Data Vortex switch directly to the target processing core.

14) An interconnect apparatus in accordance with claim 9, wherein neither the master Data Vortex switch or the Data Vortex switches on each chip need to be globally set and reset as different groups of packets enter the switches.

15) A method of communicating data packets in a network comprising:

arranging the data packets in a plurality of sub-packets, including an address sub-packet that identifies a target processing core for receiving a data packet,

placing a Data Vortex switch and an array of processing cores on a same chip, receiving data at the Data Vortex switch from an external source and receiving data at the array of processing cores from the Data Vortex switch;

wherein a data packet sent to the Data Vortex switch is not reassembled in the Data Vortex switch but travels through the Data Vortex switch in wormhole fashion.

16) A method of communicating data packets in accordance with claim 15 further comprising a formation of a data packet in a sending core having a header that includes an address for a target core and a payload with data to be sent to the target core.

17) A method of communicating data packets in a network comprising:

arranging the data packets in a plurality of sub-packets, including an address sub-packet that identifies a target processing core for receiving a data packet,

placing both a Data Vortex switch and an array of processing cores on a plurality of chips,

connecting both the Data Vortex switch and the array of processing cores on each of the plurality of chips to a master Data Vortex switch, wherein the master Data Vortex switch communicates data packets between each Data Vortex switch or each array of processing cores;

and

wherein a data packet sent to the Data Vortex switch is not reassembled in the Data Vortex switch but travels through the Data Vortex switch in wormhole fashion.

- 18) A method of communicating data packets in accordance with claim 17, further comprising sending a data packet from a sending processor core in any one of the array of processing cores included in any one of the plurality of chips to a target processing core included in an array of processing cores in another one of the plurality of chips and forming a data packet for the sending processing core having a header that identifies a location for the target processing core and a payload including the data to be sent.
- 19) A method of communicating data packets in accordance with claim 18, further comprising sending the data packet from the sending processor core to the master Data Vortex switch and from the master Data Vortex switch on the chip to the target processing core.




FIG. 1

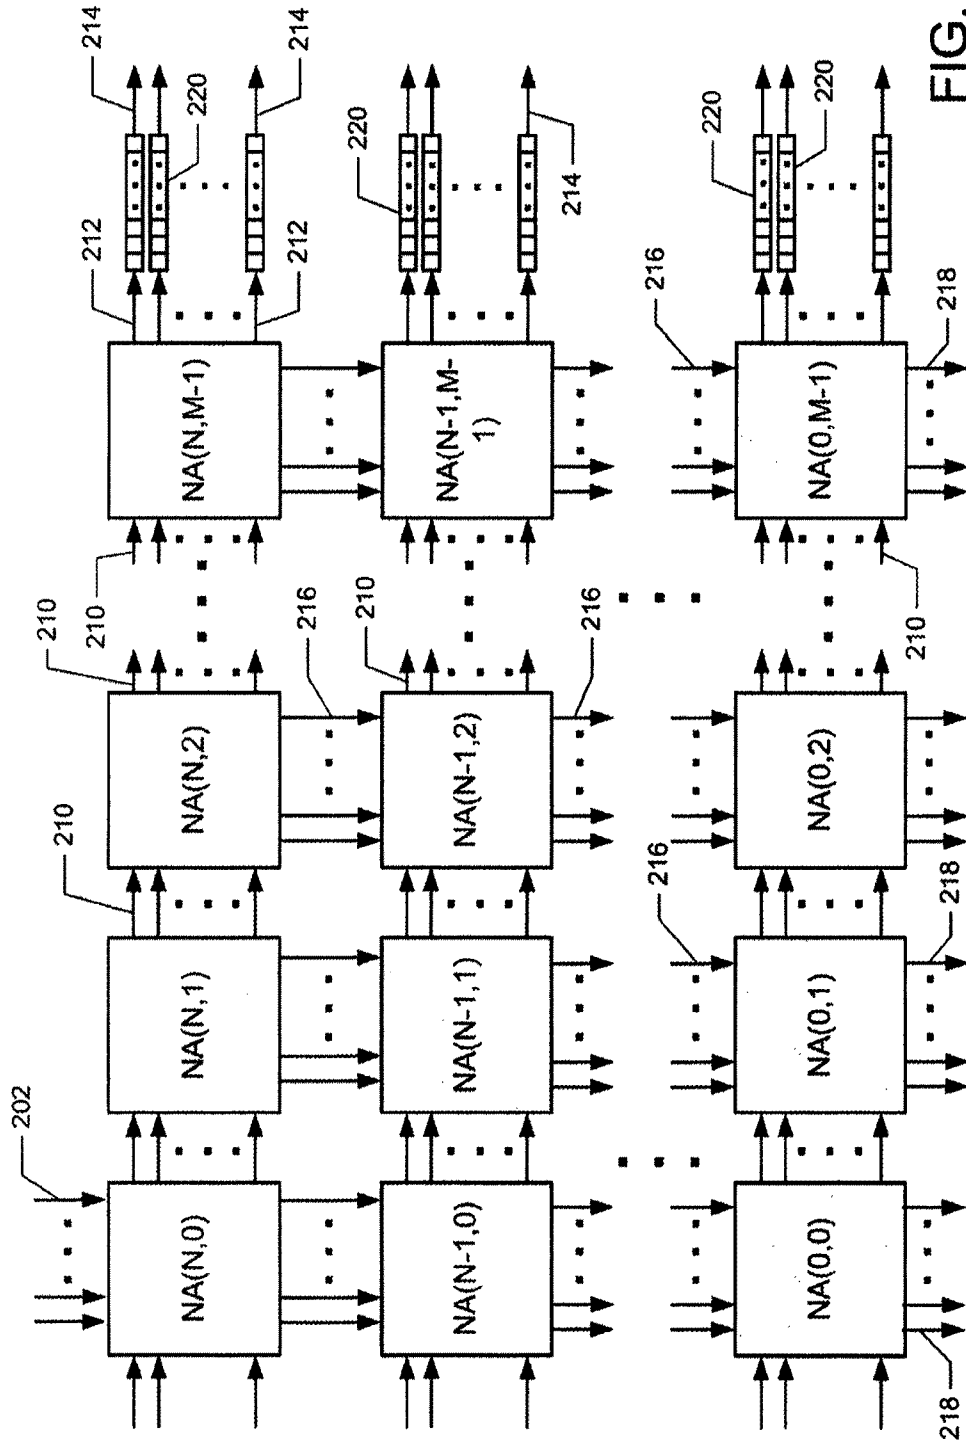


FIG. 2A
(PRIOR ART)

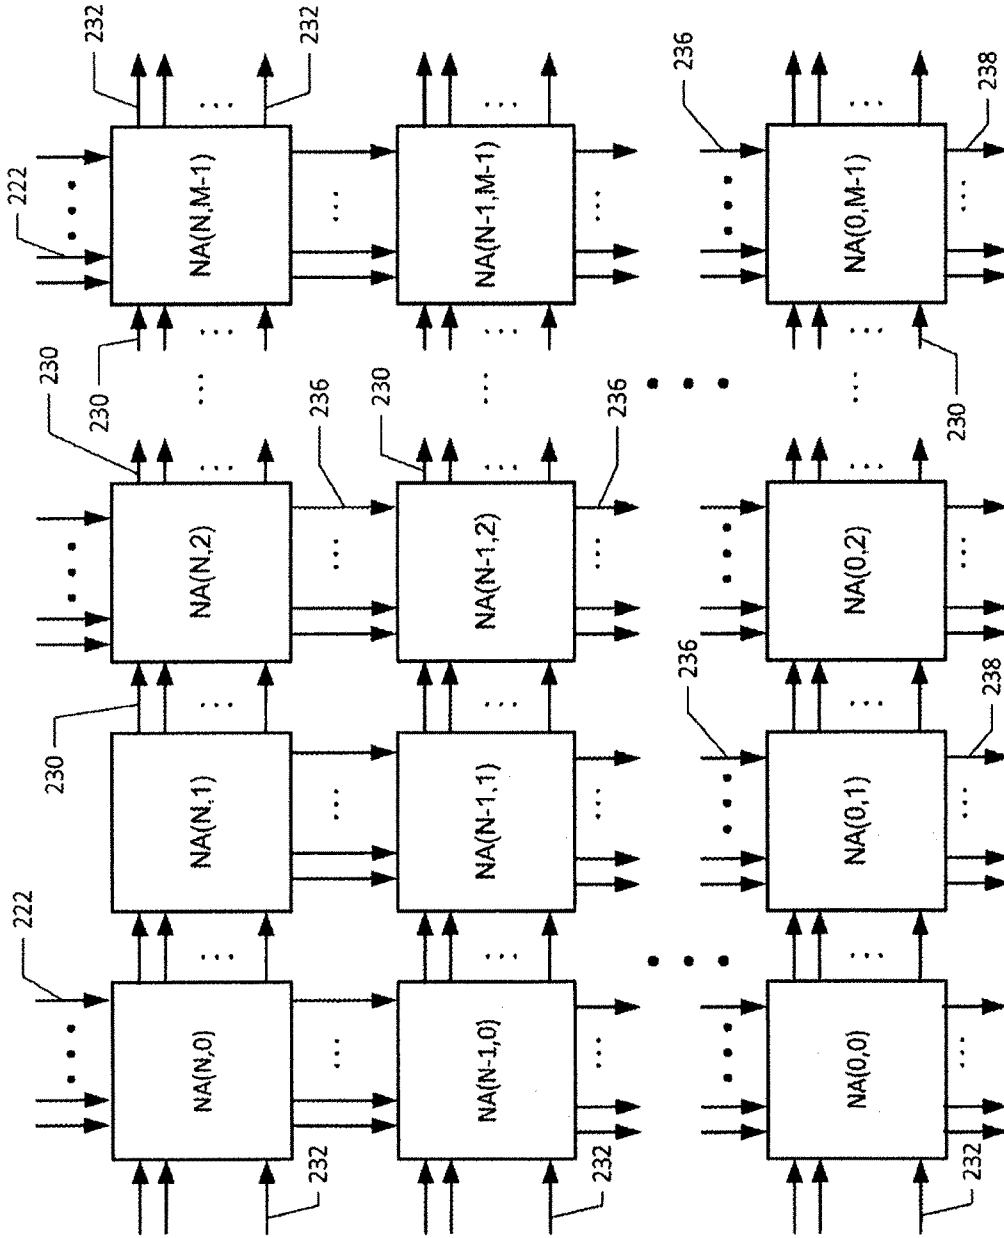




FIG. 2B






FIG. 4




FIG. 5

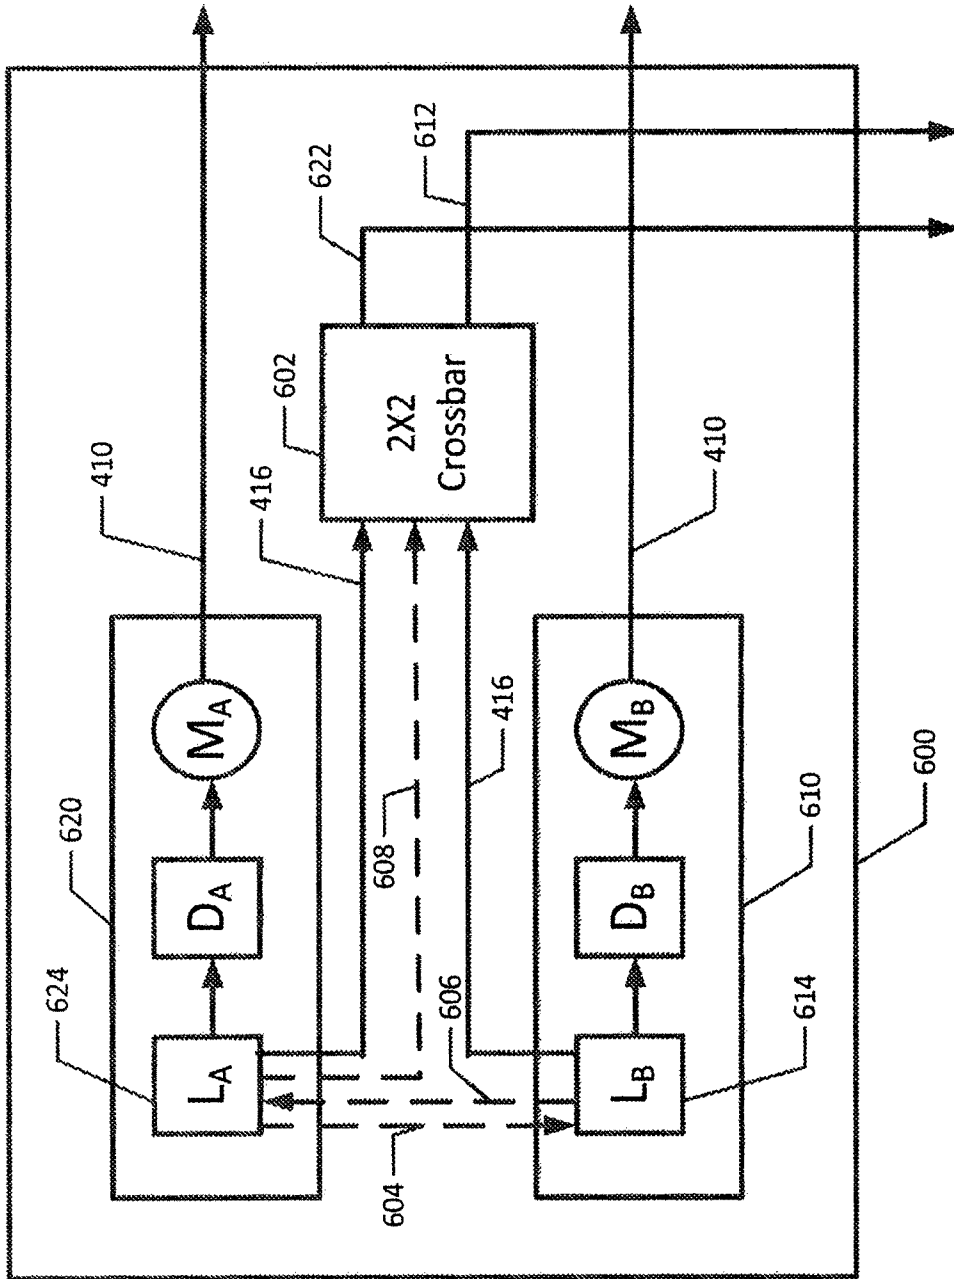


FIG. 6




FIG. 7




FIG. 8

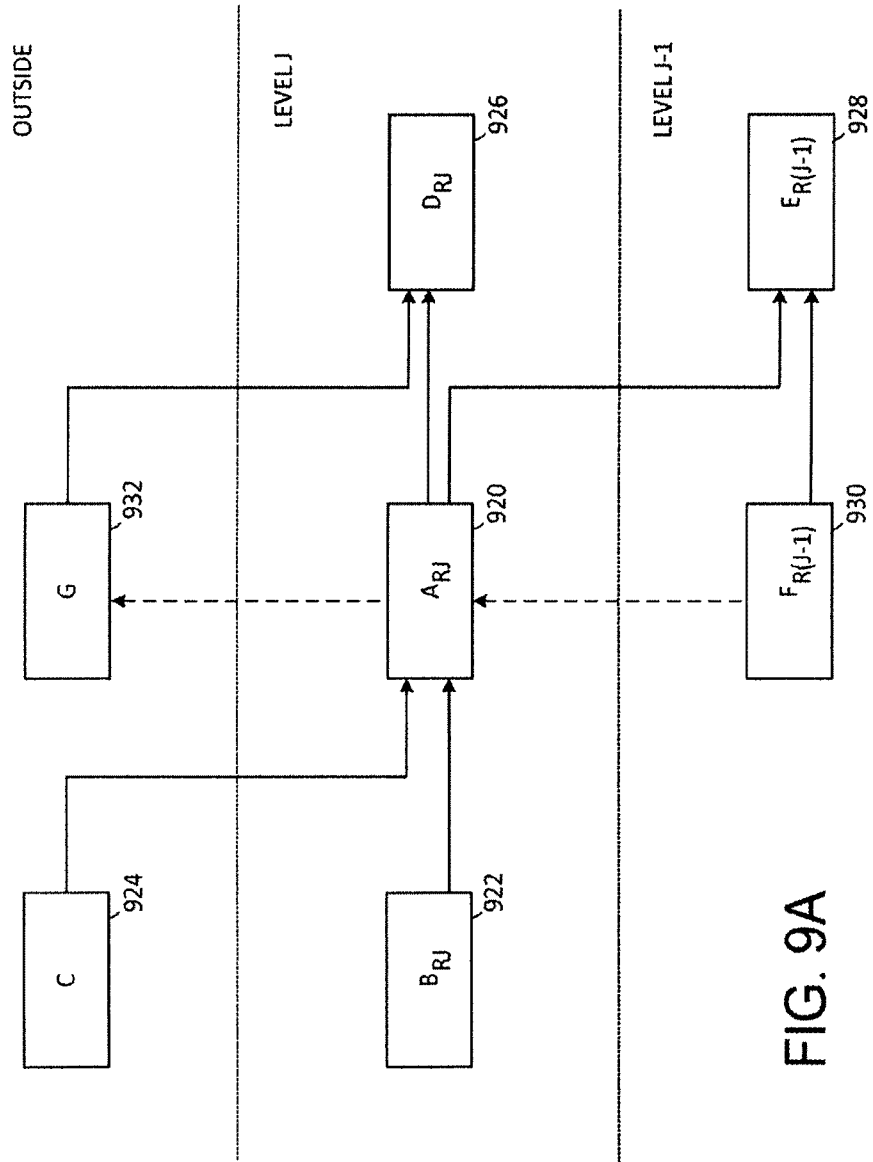


FIG. 9A




FIG. 9B




FIG. 9C




FIG. 10




FIG. 11




FIG. 12

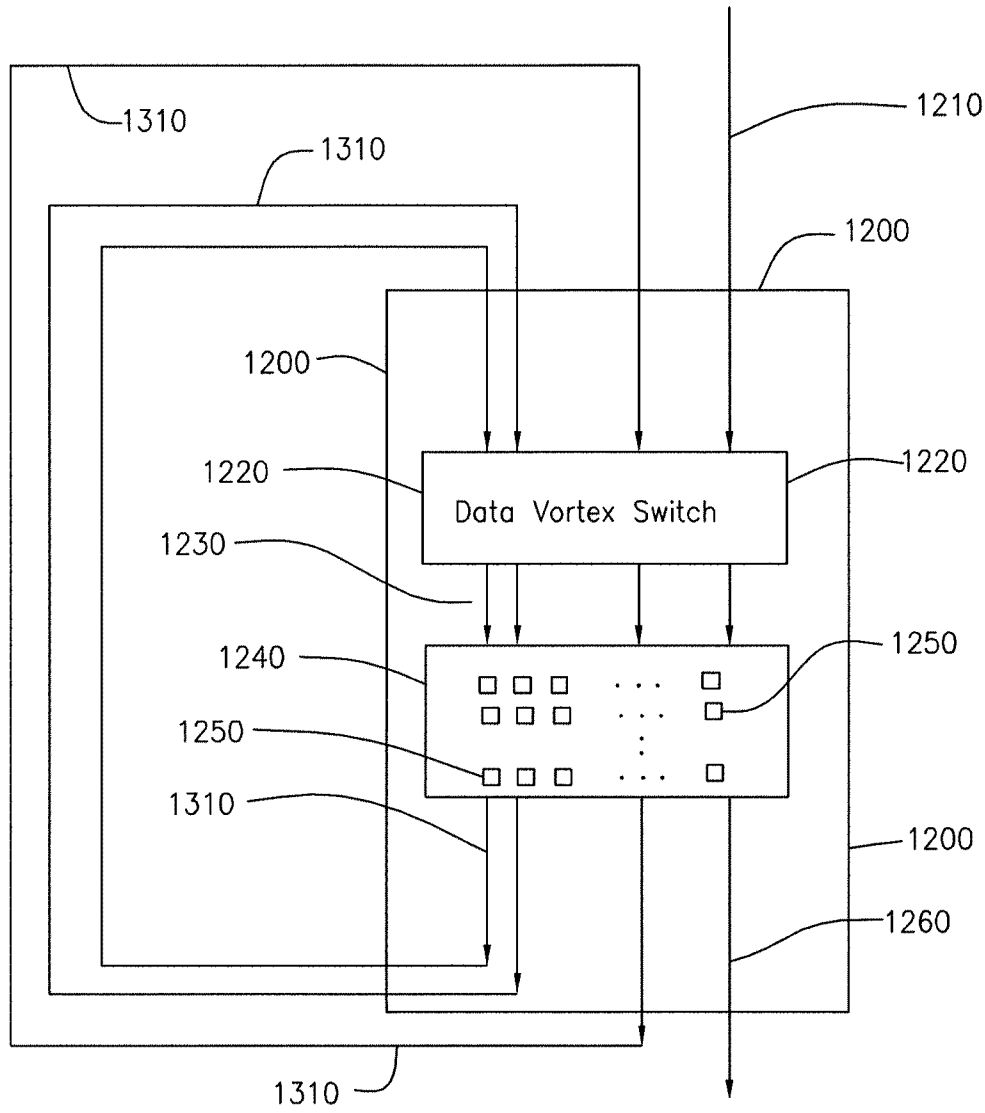


FIG. 13




FIG. 14

