



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

(51) International Patent Classification <sup>5</sup> : <b>G06F 15/80</b>	<b>A1</b>	(11) International Publication Number: <b>WO 93/19432</b> (43) International Publication Date: 30 September 1993 (30.09.93)
--	-----------	--

(21) International Application Number: PCT/US93/02416

(22) International Filing Date: 17 March 1993 (17.03.93)

(30) Priority data:  
852,603 17 March 1992 (17.03.92) US

(71) Applicant: MASSACHUSETTS INSTITUTE OF TECHNOLOGY [US/US]; 77 Massachusetts Avenue, Cambridge, MA 02139 (US).


(72) Inventors: WARD, Stephen, A. ; 95 South Street, Chestnut Hill, MA 02167 (US). PRATT, Gill, A. ; 60 Denton Road, Wellesley, MA 02181 (US). NGUYEN, John, N. ; 550 Memorial Drive, No. 6E, Cambridge, MA 02139 (US). PEZARIS, John, S. ; 366 Broadway, Cambridge, MA 02139 (US).

(74) Agents: SMITH, James, M. et al.; Hamilton, Brook, Smith &amp; Reynolds, Two Militia Drive, Lexington, MA 02173 (US).

(81) Designated States: AU, JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

**Published***With international search report.**Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.*

(54) Title: LOW-NEIGHBORHOOD THREE-DIMENSIONAL INTERCONNECT



## (57) Abstract

Circuit supporting modules form a three-dimensional communication interconnect mesh. Each module has fewer than six sets of connectors, preferably four. The preferred three-dimensional communication interconnect is a tetrahedral lattice having a regular, isotropic, three-dimensional topology in which each module connects to its four physically closest neighbors. The structure of the tetrahedral interconnect is isomorphic with a diamond lattice structure.

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	FR	France	MR	Mauritania
AU	Australia	GA	Gabon	MW	Malawi
BB	Barbados	GB	United Kingdom	NL	Netherlands
BE	Belgium	GN	Guinea	NO	Norway
BF	Burkina Faso	GR	Greece	NZ	New Zealand
BG	Bulgaria	HU	Hungary	PL	Poland
BJ	Benin	IE	Ireland	PT	Portugal
BR	Brazil	IT	Italy	RO	Romania
CA	Canada	JP	Japan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SK	Slovak Republic
CI	Côte d'Ivoire	LI	Liechtenstein	SN	Senegal
CM	Cameroon	LK	Sri Lanka	SU	Soviet Union
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	MC	Monaco	TG	Togo
DE	Germany	MG	Madagascar	UA	Ukraine
DK	Denmark	ML	Mali	US	United States of America
ES	Spain	MN	Mongolia	VN	Viet Nam
FI	Finland				

-1-

LOW-NEIGHBORHOOD THREE-DIMENSIONAL INTERCONNECTBackground of the Invention

This invention relates generally to the field of computer architecture. Typically, computers include an interconnected set of computational modules, each performing a specific task or function. These modules are assembled into a full system by connecting the appropriate wires that lead in and out of each through a bus. The bus allows various computer chips or boards to communicate and exchange data. This standardized approach has allowed system designers to physically arrange components in any manner they require without worrying about the details of interconnections.

However, the bus architecture seriously constrains the efficiency of communication among modules. This is a key limitation in today's quest for faster, more efficient computers. Since buses serialize all communications, they create a bottleneck. To make matters worse, the total capacity of the intercommunication remains roughly constant as the system size and the demand for the exchange of information increases. Moreover, the time taken by each transaction must accommodate the physical length of the bus, a crucial factor that limits the communication bandwidth in all but the shortest buses. These deficiencies reduce the viability of the bus as a basis for communication in modern computer systems.

To increase the efficiency in communication among modules, modern multiprocessor architectures have been

-2-

designed with various different topologies other than the bus. Such topologies include one-dimensional, two-dimensional and three-dimensional topologies.

#### Summary of the Invention

5       The size of future computer systems mandates the use of three-dimensional topologies for efficient utilization of space and communication bandwidth. For maximum flexibility, these topologies should have incremental expandability, arbitrary scalability and support fast  
10 static and dynamic communication.

      The present invention provides a modular and high-performance communications substrate that standardizes the mechanical, electrical and logical interconnection among modules arranged in a three-dimensional mesh. This  
15 mesh is constructed of replicated near-neighbor links whose mechanical and electrical characteristics are rigidly constrained and highly optimized. These constraints allow performance parameters for non-local communications that compete favorably with those of the  
20 ad-hoc interconnects found in contemporary supercomputers.

      In accordance with a preferred embodiment of the present invention, there is provided a three-dimensional electric circuit assembly. The assembly includes a  
25 plurality of modules to support respective circuits. Each module has at least three sets of connectors along fewer than six flat surfaces thereof. The connectors join the circuits on the substrate to adjacent modules. Adjacent modules are oriented relative to each other such  
30 that adjacent modules interconnect in a three-dimensional array.

      The modules may take any form including circuit boards, semiconductor chips, multichip modules and housings of any shape. Circuit board modules with

-3-

vertical top and bottom connectors are particularly suited to the assembly.

The preferred shape of the three-dimensional array is a tetrahedral lattice, wherein each of the joined  
5 modules communicate with four other modules. The preferred embodiment is isomorphic to a diamond lattice. The diamond lattice is particularly advantageous in that it is isotropic. Accordingly, each module has connectors for connecting to adjacent modules along four axes.  
10 Adjacent modules are connected along a single common axis but are rotated relative to each other such that connections from adjacent modules are made in opposite directions.

The assembly may be made from a plurality of flat  
15 modules with only vertical interconnections. Accordingly, each flat surface of a module may support a pair of like, parallel, extended connector sets. The connector sets on opposite surfaces of the substrates are complementary to each other and rotated 90° relative to  
20 each other. Modules are joined to adjacent modules such that the modules define a three-dimensional tetrahedral lattice.

In another embodiment, the three-dimensional circuit assembly is constructed from triangular shaped modules  
25 wherein each module communicates with three adjacent modules. The connectors may be formed on a top triangular surface, a bottom triangular surface and an edge of the module. Each connector extends along an axis transverse to an axis of connection, and the transverse  
30 axes are displaced 120° relative to each other.

The substrate of each module may support a plurality of circuit chips, and auxiliary circuit boards may be mounted to the substrates.

In the lattice which is isomorphic with the diamond  
35 lattice, each module is physically identical relative to

-4-

connectors, but is oriented in one of two directions. Each module is connected only to modules of the opposite orientation. Each of the two orientations may also be identified with one of two-phases of a two-phase clock, so that no neighboring modules operate on the same clock phase in intermodule communications.

In a lattice structure which is isomorphic with the diamond lattice and which comprises flat overlapping modules, the modules may be spaced horizontally within levels of the lattice, thereby allowing paths for coolant flow and space for auxiliary boards. The resultant serpentine flow paths may then insure turbulent flow as coolant fluid, air or liquid, is directed therethrough.

The three-dimensional assemblies of the present invention may be addressed by a four element vector. Each vector element represents four non-orthogonal directions that a particular message being sent from a substrate may travel. The four element addressing system ensures that no two modules are assigned the same address and that no two addresses are assigned the same module.

The present invention further includes a preferred method of providing a three-dimensional circuit assembly. First, a plurality of modules are provided. Then circuits are placed over the plurality of modules. Next, adjacent modules are connected along module axes such that the connected modules are rotated relative to each other to form a three-dimensional array.

While the present invention will hereinafter be described in connection with a preferred embodiment and method of use, it will be understood that it is not intended to limit the invention to this embodiment. Instead, it is intended to cover all alternatives, modifications, and equivalents as may be included within the spirit and scope of the present invention as defined by the appended claims.

-5-

Brief Description of the Drawing

Figure 1 shows a three-dimensional communication interconnect of the prior art.

Figure 2 shows a crystal structure of a diamond lattice to which one embodiment of the present invention is isomorphic.

Figure 3 is a rectilinear array remapping the atoms of the diamond lattice.

Figure 4a shows an exploded view of 3D tetrahedral interconnect with four connector sets at each node to connect to four neighbors (4N).

Figure 4b shows interconnection of several of the modules of Fig. 4a with connectors illustrated.

Figures 5a and 5b show detailed views of modules embodied in the 3D-4N mesh.

Figure 6 shows a 3D-3N communications interconnect.

Figures 7a and 7b show a front and backview of a drum embodied in the 3D-3N communications interconnect.

Figure 8 shows the module of Figures 7a and 7b in plan and side views.

Detailed Description of the Invention

Communication networks are characterized by their connectivity. Classification of networks depend on the number of nodes reachable in  $n$  time steps from a central node within an infinite network. If the number of reachable nodes grows linearly (e.g., a factor times the number of time steps plus a constant), the network is classified as one-dimensional and designated as  $O(n)$ . Similarly, if the number of reachable nodes grows quadratically, the network is classified as two-dimensional and designated as  $O(n^2)$ . Thus, if the number of distinct nodes grows as  $n^3$ , then the network is

-6-

classified as three-dimensional and is designated as  $O(n^3)$ .

Each node in the network is connected to its neighbors through communications ports or connector sets.

5 A node is characterized by the number of ports or the number of neighbors connected to the ports. A node having three ports is considered to be less complex than a node having six ports.

A description of a communications network can be  
10 formulated from the connectivity and the number of neighbors. For example, a series of two-port nodes strung together in a one-dimensional topology would be called a 1D-2N network; a collection of four-port nodes connected in a two-dimensional Cartesian mesh would be  
15 called a 2D-4N network. For certain combinations, this description is complete (e.g., there is only one 1D-2N network); for others, it is incomplete.

A scalable, three-dimensional, 6-neighbor (3D-6N) communication interconnect or mesh 10 is illustrated in  
20 Figure 1. The mesh includes substrates or modules or nodes 12 that are connected together to form a three-dimensional mesh. Figure 1 depicts each module as a cube whose faces provide signal contacts as well as power and coolant to six immediate neighbors.

25 The present invention discloses various topologies that are more efficient and have a lower node complexity than the mesh shown in Figure 1. The preferred topology of the present invention is a 3D-4N network.

The preferred 3D-4N network is based upon the  
30 crystal structure of a diamond lattice (i.e., crystalline carbon). The diamond lattice is incorporated in the 3D-4N mesh of the present invention and is shown in Figure 2. In a diamond lattice, each carbon atom has four links or axes that bond with four neighboring atoms. An



-7-

aggregate of four atoms form a unit cell which, when repeated, constitutes the complete crystalline lattice. In the 3D-4N, circuit supporting modules correspond to the carbon atoms with each being linked to four neighboring modules. We designate each module in the diamond lattice shaped 3D-4N mesh with a numeral (0,1,2,3,4) that is representative of their position level within the mesh. Substrates having a level 0 are positioned at the bottom of the mesh, substrates of level 1 are positioned above level 0 substrates, level 2 substrates are positioned above level 1 substrates, then level 3 is the next highest level, and level 4 occupies the top level. Level 4 can then be labeled level 0 for the next group of the repeated lattice structure. The levels of the mesh are joined together by connecting each substrate's axes to substrates at adjacent levels. For example, each level 0 (4) substrate is linked to the axes of two level 1 and two level 3 substrates. Each level 1 substrate is linked to two level 0 and two level 2 substrates. Each level 2 substrate is linked by its four axes to two level 1 and two level 3 substrates. After linking all of the substrates together, the 3D-4N mesh has the form of a diamond lattice.

This 3D-4N tetrahedral mesh can be further conceptualized by using an orthogonal array of mechanically homogeneous nodes that are remapped by projection onto a plane from the diamond lattice shown in Figure 2. Figure 3 shows a rectilinear array of the remapped atoms. The nodes or substrates in Figure 3 are numbered to show the relative heights and levels of corresponding atoms in the diamond lattice cell. Links between the nodes of Figure 2 are represented in Figure 3 as arrows pointing towards nodes positioned at higher layers. The arrows between each respective node represent a horizontal and vertical link. For instance,

-8-

the arrow between a pair of nodes labeled 3 and 4 represent a link horizontally and vertically to a node in the level 4, which is directly above a level 0 node and the level 0 node of the next set of levels.

5       The 3D-4N mesh of the present invention is derived from the rectilinear array of Figure 3 by reducing the internode spacing and offsetting alternate layers in the x and y directions, allowing contacts on overlapping horizontal surfaces to effect communications in the x  
10       direction at the level 1/2 and level 3/0 boundaries, and in the y direction at the level 0/1 and level 2/3 boundaries.

Figure 4a shows the 3D-4N mesh 14 having five levels of modules that correspond to the tetrahedral lattice in  
15       Figure 2 and to the rectilinear array of Figure 3. The 3D-4N tetrahedral interconnect includes a plurality of flat surfaced modules or modules or nodes 16 which each support respective circuits 18 (Figure 4b). The lowest substrate level of mesh 14 is represented by a level 0.  
20       The next highest substrate level is represented by level 1, then level 2, then level 3, and the highest level is represented by a level 4.

The modules 16 are positioned in the mesh 14 so that there is an overlap along each of the plurality of  
25       substrates at levels 0-4. Vertical connectors and flat substrates are used to create the 3D-4N tetrahedral interconnect 14. The substrates overlap one another horizontally to allow connector mating. This is possible because only two of the four edges on each side of a  
30       substrate are used for connection, allowing uniformly sized substrates to be connected vertically with horizontal overlap but without horizontal interference. The result is a 3D-4N mesh that spirals up from the bottom level 0 to the top level 4.

-9-

The mesh is configured such that adjacent substrates are rotated in connector orientation, preferably about 90°, relative to each other. Thus, communication from any one substrate is generated in four different directions. For example, a module designated as a level 1 could vertically communicate data upward to two substrates designated as a level 2 and vertically communicate data downward to two substrates designated as level 0.

Three example axes of interconnection are illustrated at 15 in Figure 4a. These axes correspond to the links of Figure 3, with arrows pointing upwardly and to the links of Figure 2. Corresponding nodes of Figures 2, 3 and 4a are indicated by like two-digit numbers preceded by figure number. For example, node 201 in Figure 2 corresponds to node 301 in Figure 3 and node 401 in Figure 4a. It can be seen from the figures that node 401 connects downwardly in the x direction. It would also connect upwardly in the y direction with substrates not shown. The adjacent substrates 403 and 405 on the lower level 3 can connect downwardly in the y direction to substrates 407, 409 and 411, 413, respectively, and upwardly in the x direction. The level 2 substrates then connect downwardly in the x direction to, for example, modules 415 and 417. Finally, the spiral is complete with connections downwardly in the y direction to module 419.

It can be seen that, though vertical connectors are used, the effective axes of connection are not along the z axes of the figures. For example, the connections from substrate 405 are in the (z, -x), (z, x), (-z, y) and (-z, -y) directions. The axes of connection in the adjacent levels, both levels 2 and 4, are in the (-z, x), (-z, -x), (z, -y) and (z, y) directions. From the above listings, it can be seen that adjacent modules have

-10-

connections in opposite directions along the same axes. For example, the  $(-z, x)$  direction of level 2 modules is directly opposite to the  $(z, -x)$  direction of level 3 and so on. As discussed below, there are two orientations of  
5 modules characterized by the directions of connection.

Within the mesh of modules are unoccupied spaces that form a plurality of serpentine channels 25 that convey cooling fluid throughout the mesh. As fluid is forced vertically through the channels, a turbulent flow  
10 of air is created within the mesh. The turbulent air flows through the many cross-over passages to evenly cool all heat dissipating substrates 16 within the mesh 14.

The plurality of air channels can be closed to any degree to permit the mesh 14 to be consolidated into a  
15 denser interconnection if desired. Consolidation is done by moving the substrates closer to each other within each level, thus increasing the horizontal overlap between level. A dense interconnection is desirable for attaining a faster and more efficient computer network.  
20 The 3D-4N mesh of the present invention has the capability to be indefinitely dense, but such a structure would have limited cooling capability.

The diamond lattice 3D-4N network 14 has an isotropic topology, which means that any logical subtree  
25 observed from any port of any substrate in an infinite network is isomorphic (substantially identical) to all other subtrees observed from that or any other substrate in the same network.

A further characteristic of this diamond lattice 3D-  
30 4N network is the mechanical homogeneity of the assembly where all substrates are mechanically identical at their interfaces. Thus, any node may be placed in any position within the network and maintain the mechanical integrity of the network. Although the nodes are homogeneous they

-11-

may be merely mechanically compatible and perform different logical functions.

Figures 4b shows a more detailed view of the modules 16 within the 3D-4N mesh of Figure 4a. Each module includes circuits 18, sockets 19 and pins 21. The sockets 19 are placed along opposite sides of each flat surfaced substrate for joining modules above. Pins 21 are placed underneath each of the flat surfaced substrates at opposite sides for connecting modules below. The pins and sockets of each module are complementary to the sockets and pins of each adjacent modules. Thus, the plurality of substrates are joined together through respective sockets and pins.

Figure 4b shows modules each communicating with adjacent modules via their respective connector sets. For example, a module designated as level 2 communicates with the circuits of modules designated as levels 1 and level 3. This configuration allows information from modules to propagate in four directions.

Each module is connected to its neighbors only in the vertical dimension, meaning that individual boards may be removed from, or added to, the top of the machine in the fashion of bricks. At each step of assembly or disassembly, the mating or unmating force of only half of a single board's connectors must be applied, and alignment need only be accomplished on the scale of a single node's width.

The mesh of the present invention obviates the problem associated with the 6-neighbor Cartesian three-dimensional system (3D-6N) of Figure 1, which is built from very large circuit boards stacked atop one another via a number of connectors with a total of thousands of pins. To repair an interior fault, a mechanical access operation must take place which usually means splitting

-12-

the entire machine along a two-dimensional plane intersecting the defective part.

Figure 5a shows integrated circuits 18 and 19 placed on the top surface of a module. The circuits allow each  
5 node to process data and communicate with other circuits on adjacent nodes through their respective connector sets. For example, chip 18a may provide the intermodule communications control and switching circuitry while chip 18b provides local node processing.

10 In an embodiment with minimal horizontal overlap, only one quarter of the three dimensional volume is occupied. This leaves a vertical space for auxiliary boards such as "daughterboards" to be plugged atop each substrate. Also, this space could be used for  
15 distributed peripherals such as disk drives. Figure 5b shows a daughterboard 27 standing over the substrate. On top of the daughterboard are four integrated circuits 31. The multi-chip module allows each node to increase its communication and processing capabilities.

20 The plurality of modules facilitate communication in the mesh through the use of cross-point switches and registers which are supported by each respective module. The cross-point switches and registers provide a parallel connection among the plurality of modules.

25 The hardware complexity of the 3D-4N interconnect 14 is characterized by the cross point switches which have a high performance switching node of C channels, each of width W, that grows roughly with  $O(C^2W)$ . Networks such as the 3D-6N neighbor mesh utilize 6 neighbor switches,  
30 with one channel for each direction of travel along each orthogonal dimension. By utilizing a 4 neighbor topology instead of 6, the tetrahedral interconnect 14 realizes a reduction of switch complexity. Keeping the number of

-13-

communications pins per node (C W) constant, we have

$$C_4 = \frac{2}{3} C_6 \text{ and } W_4 = \frac{3}{2} W_6$$

The 4 and 6 sub-units denote 4 neighbor and 6 neighbor networks, respectively. This yields a switch complexity  
5 for the 4 neighbor switch which is 2/3 (66%) as large as that of the total pin equivalent 6 neighbor switch. If an additional channel is used for communications to the local node, then a 5 neighbor is compared to a 7 neighbor switch, and the savings are somewhat reduced to 5/7  
10 (75%), but are still substantial.

In a mesh, a significant portion of the signal propagation delay of a cross-point switch drops linearly with a reduction of its capacitive load. This load corresponds roughly to the number of switched channels.  
15 Thus, there is a significant decrease of individual switching node latency (or, equivalently, an increase in network clock frequency) in 4 neighbor switches versus 6 neighbor switches due to the lower number of channels.

The 3D-4D tetrahedral interconnect 14 as shown in  
20 Figure 4a is isotropic. An anisotropic 3D-4N network may be constructed by removing two of the six channels connecting each node of a 6N 3D mesh. The configuration would appear as in Figure 1, with each module having connections on its top and bottom surfaces as well as two  
25 opposite lateral faces. The lateral connections would alternate in X and Y directions with successive vertical levels.

Unfortunately, if channel widths are uniform, such a network has unequal bandwidth along the three

-14-

directions of travel, making intelligent routing and load distribution a complex matter. In order to balance network load, it becomes necessary to relocate communicating processes further apart along the

5 directions of high network bandwidth and closer together along the directions of low network bandwidth. Alternatively, one may vary the number of pins dedicated (i.e., the width) per channel in each dimension. Overall, the anisotropic network is inferior to an

10 isotropic network of equivalent channel and pin count.

As discussed above relative to Figure 4a, adjacent substrates connect in opposite directions. It is seen that two node orientations are present, as distinguished by the X, Y and Z (depth) offsets of their connections to

15 neighboring nodes:

Table 1

Connection Offset to Neighbor

	<u>Orientation 0</u>			<u>Orientation 1</u>		
	<u>X</u>	<u>Y</u>	<u>Z</u>	<u>X</u>	<u>Y</u>	<u>Z</u>
20	+1	0	-1	-1	0	+1
	-1	0	-1	+1	0	+1
	0	+1	+1	0	-1	-1
	0	-1	+1	0	+1	-1

Node orientations of the present invention are

25 labeled as '0' and '1'. Note that the available link offsets of each node orientation are negative of one another, i.e., orientation '0' nodes only connect to orientation '1' nodes, and vice versa.

Each node-to-node connection involves an increment

30 or decrement in either the X or Y directions (but not both) and always an increment or decrement in the Z (depth) direction.



-15-

In the Cartesian 3D-6N network of Figure 1, messages usually travel along "Manhattan" routes, with each step being made in the direction of one of the orthogonal axes. The tetrahedral interconnect 14, despite being a  
5 three-dimensional topology, has 4 non-orthogonal directions of travel, corresponding to the 4 link offsets which are shown in Table 1. Thus, it is difficult to describe the node coordinates and routing directions of the tetrahedral interconnect in a three coordinate  
10 system.

A simplified addressing scheme is obtained by utilizing a 4 coordinate or 4-dimensional addressing system. In the 4-D addressing system, the direction of each link axis in the diamond lattice corresponds to one  
15 of the 4 dimensions.

As mentioned above, there are two node orientations in the tetrahedral interconnect. At any node, a message may only travel on any one of the 4 dimensions in a single direction. Having traveled one link in that  
20 direction, the message arrives at a node of the opposite orientation, and may not continue along that same dimension in the same direction. Instead, travel must proceed in the opposite direction, along any of the 4 dimensions.

25 For the sake of uniformity, traveling from any orientation '0' node to any orientation '1' node has an offset value of +1, while traveling from any orientation '1' node to any orientation '0' node has an offset value of -1. The mapping of 3-D position offsets [X,Y,Z] to  
30 4-D address offsets [A,B,C,D] is indicated in Table 2.

-16-

Table 2Offset to Neighbor: Orientation 0 Node

<u>3-D Position</u>			<u>4-D Position</u>				
<u>X</u>	<u>Y</u>	<u>Z</u>	<u>A</u>	<u>B</u>	<u>C</u>	<u>D</u>	
5	+1	0	-1	+1	0	0	0
	-1	0	-1	0	+1	0	0
	0	+1	+1	0	0	+1	0
	0	-1	+1	0	0	0	+1

Offset to Neighbor: Orientation 1 Node

10	<u>3-D Position</u>			<u>4-D Position</u>			
	<u>X</u>	<u>Y</u>	<u>Z</u>	<u>A</u>	<u>B</u>	<u>C</u>	<u>D</u>
	-1	0	+1	-1	0	0	0
	+1	0	+1	0	-1	0	0
	0	-1	-1	0	0	-1	0
15	0	+1	-1	0	0	0	-1

The center network node is labeled [0,0,0,0] and arbitrarily declared to be orientation '0'. A message may travel from the center node [0,0,0,0] to either [1,0,0,0], [0,1,0,0], [0,0,1,0] or [0,0,0,1]. Once at this second node (an orientation '1'), the direction of travel may be on any of the 4 dimensions, but is restricted to be in the opposite direction, i.e., an address increment of -1. For example, if the second node was [1,0,0,0], the message may now travel to a third node at either [0,0,0,0], [1,-1,0,0], [1,0,-1,0] or [1,0,0,-1].

Since continued travel must invoke alternating address coordinate increments of +1 and -1, the sum of the 4D address coordinates is constrained to be 0 on all orientation '0' nodes (those being an even number of hops away from the central node), and +1 on all orientation '1' nodes (those having an odd number of hops away from the center). Thus, the coordinate sum is exactly equal

-17-

to the node orientation, and is also equal to the parity of the routing distance from the center.

By induction, one can show that any address whose 4 coordinates sum to either 0 or 1 is a valid node address.

5 The sum-to-0 or sum-to-1 constraints correspond to forcing the node coordinates to lie on one of two 4-dimensional hyperplanes, with each hyperplane corresponding to the orientation of the node. Since links only connect orientation '0' to orientation '1' nodes (and vice versa), messages cannot travel directly  
10 within each hyperplane. Instead, they must zig-zag between the two.

Since the present invention uses a 4-D address, there may be concern as to how well defined the  
15 addressing system is. The 4-D address space is well defined if:

1. No two nodes may be assigned the same address.
2. No two addresses may be assigned the same node.

Since we know that physical node locations in 3-space are  
20 unique, verifying these two properties equates to demonstrating that the mapping from 3-D locations to 4-D addresses is one-to-one.

If we label a center node [0,0,0,0], and use the offsets indicated in the table above, a node's 3-D  
25 position ([X,Y,Z]) may be calculated from its 4-D address ([A,B,C,D]) as follows:

$$\begin{bmatrix} X \\ Y \\ Z \\ T \end{bmatrix} = \begin{bmatrix} A & +1 & -1 & 0 & 0 \\ B & 0 & 0 & +1 & -1 \\ C & -1 & -1 & +1 & +1 \\ D & +1 & +1 & +1 & +1 \end{bmatrix}$$

30

Note that a fourth quantity T has been added to the 3-D physical coordinate which indicates the orientation (either 0 or 1) of the node, and is given by the sum of its 4-D address.

-18-

From the above equation, a 4-D address [A,B,C,D,] is seen to map to only one 3-D location and orientation [X,Y,Z,T].

The right hand matrix has a unique inverse, so that:

5

$$\begin{bmatrix} A \\ B \\ C \\ D \end{bmatrix} = \begin{bmatrix} X & +1/2 & 0 & -1/4 & +1/4 \\ Y & -1/2 & 0 & -1/4 & +1/4 \\ Z & 0 & +1/2 & +1/4 & +1/4 \\ T & 0 & -1/2 & +1/4 & +1/4 \end{bmatrix}$$

10 Thus, a given 3-D location (and orientation, which is necessarily constrained to have only one value) maps to only one 4-D address. The mapping is therefore one-to-one.

The 4-dimensional addressing scheme allows routing  
15 within the tetrahedral interconnect in a Manhattan-style fashion, albeit, with 3 out of every 4 streets (links) missing. The missing links limit the number of available routes, but the computation of route-distance is still the same. One simply sums the address differences over  
20 all 4 coordinates to obtain the minimum number of links that must be traversed.

At each step along a minimal distance route, a choice of at most 3 directions of travel is possible (if one is not at the destination, the difference between the  
25 current and desired address can have at most 3 non-zero coordinates of the same sign). If one is at an orientation '0' node, the next hop will add 1 to one of the dimensions, thus one has the choice of any of the dimensions whose target coordinate is greater than the  
30 corresponding coordinate of the current node. For some minimal routes (i.e. [0,0,0,0] to [0,0,10,-10]) there is no choice whatsoever.

Deadlock-free adaptive routing schemes that work for Cartesian 3D-4N meshes also work for the 3D-4N

-19-

tetrahedral interconnect, provided that the 4-D addressing scheme is used (which provide Manhattan routing) and that each "missing" link in 4-D is marked as unavailable to the router. The 3D-4N tetrahedral  
5 interconnect has straight through channels which only last a single "hop" before a turn is required. The new deadlock-free routing scheme of the present invention significantly improves performance and load balancing in the tetrahedral interconnect.

10 One example of such a routing scheme involves the grouping of pairs of links as the means to route from one node of a given type to another node of the same type. Let us assume arbitrarily that the beginning node is in orientation 1, and that a route to another orientation 1  
15 node is required. This means that an even number of hops must be used, i.e. first subtract 1, and then add 1 to pairs of address coordinates until the destination is reached (if the 4-D address has 4 bins filled with varying numbers of tokens, this corresponds to removing  
20 a token from one bin and depositing it in another).

At any node, the choice of which pair of coordinates to decrement and then increment involves a choice of 1 of 12 possibilities. This corresponds to 2 polarities of double-hop travel along 6 different unsigned directions.  
25 This is not to suggest a 6-dimensional addressing scheme, which is easily shown to be not well defined (i.e., there are two distinct paths between two nodes with different 6-D address offsets). The 6 directions are more useful for direct routing than the 4 Manhattan address  
30 dimensions because travel may continue in the same direction for successive double-hops. Thus, a routing strategy that utilizes strict ordering of travel directions to avoid deadlock (and adaptive modifications of such) can now be used directly, albeit, with six  
35 considered directions of travel (instead of the usual

-20-

three) at each node. In a non-adaptive system, at each intermediate node, one would route as far as possible along each of the six double-hop directions until the destination was reached.

5        Since each link in the double-hop routing scheme is used for three different signed directions of travel, a buffering of three words in each signed link direction is required.

Figure 6 shows another three-dimensional mesh (3D-  
10 3N) 24 constructed from 3-neighbor nodes or substrates (3D-3N) formed from physically triangular solids. Each node in this diagram comprises a flat solid whose cross-section is an equilateral triangle. Three nodes are stacked vertically to form a drum 28. Drums are  
15 assembled into vertical columns which are then close-packed in the two horizontal dimensions.

Each node 26 communicates with vertically adjacent neighbors through connectors 30 offset by  $120^\circ$ , so that nodes in a column are rotated with respect to one  
20 another. The third communication port of each node is on one of the rectangular lateral faces and mates with a similar port of a laterally adjacent node. One connector is placed on the upper face, another on the lower face, and the last on one of the side faces.

25        Figures 7a and 7b show a front and back view of the drum. Each drum 28 communicates with five neighboring drums, wherein connections with the three lateral neighbors are at different levels within the drum. Lateral communication between neighboring drums typically  
30 requires vertical communication to reach a layer with the appropriate lateral connection. The connectors on the vertical faces form a spiral with each face having a connector.

In order for a node group to tessellate in 3-space,  
35 a connector is placed on each collective face. The

-21-

condition that upper and lower connectors be on distinct faces of the triangular node ensures that the nodes' orientation spirals through a vertical column. Therefore, the drum 28 has a minimum node group that  
5 satisfies the connector-on-each-face condition.

If the drums shown in Figures 7a and 7b is interpreted with the views shown in Figure 8, the drums, if pushed together from the right and left, will mate properly (recall the edge connector is rotationally to  
10 the right of the orientation mark). Thus, the column may have an arbitrary height, with adjacent columns rotated 120° relative to each other.

The 3D-3N mesh 24 of the present invention may have rectangular nodes instead of triangular nodes. One  
15 possibility is a 3D-3N mesh constructed with rectangular nodes is composed of four blocks rather than the three since there are four vertical faces. Another possibility is to create drums of 3 nodes out of rectangular nodes by not putting connectors on one of the sides of the  
20 rectangle. The drums may still be put together in a way that preserves the topology of the triangular-node system.

While the invention has been particularly shown and described in conjunction with a preferred embodiment  
25 thereof, it will be understood that many alternatives, modifications, and variations will be apparent to those skilled in the art without departing from the spirit and scope of the invention as defined by the appended claims. For example, though connectors are shown mounted to  
30 preferred vertical and horizontal surfaces, they may, for example, be mounted on beveled edges such that they are directed at an angle. The terms vertical and horizontal are simply relative terms and do not limit the orientation of the array.

-22-

CLAIMS

1. A three-dimensional electronic circuit assembly comprising: a plurality of flat surfaced modules supporting respective  
5 circuits, each module having at least three sets of connectors along fewer than six flat surfaces thereof for joining the circuits on the modules to adjacent modules, adjacent modules being oriented relative to each other  
10 such that adjacent modules interconnect in a three-dimensional array.
2. A three-dimensional electronic circuit assembly according to Claim 1, wherein each module communicates with three other modules.
- 15 3. A three-dimensional electronic circuit assembly according to Claim 1, wherein each module communicates with four other modules.
4. A three-dimensional electronic circuit assembly according to Claim 3, wherein the  
20 modules are addressed within the array by a four element vector.
5. A three-dimensional electronic circuit assembly according to Claim 1 wherein the three-dimensional array is an isotropic  
25 lattice.
6. A three-dimensional electronic circuit assembly as claimed in Claim 1 wherein a plurality of circuit chips are supported by each module.



-23-

7. A three-dimensional electronic circuit assembly according to Claim 1 further comprising auxiliary circuit boards mounted to the substrates.
- 5 8. A three-dimensional electronic circuit assembly according to Claim 1 wherein the connectors provide vertical connections.
9. A three-dimensional electronic circuit assembly as claimed in Claim 1 wherein the  
10 three-dimensional array forms a lattice structure which is isomorphic with a diamond lattice.
10. A three-dimensional electronic circuit assembly according to Claim 1 wherein each  
15 module is of one of two orientations, each orientation having different directions of interconnection, each module being connected only to modules of the opposite orientation.
11. A three-dimensional electronic circuit  
20 assembly according to Claim 10 further comprising a two-phase clock input for controlling communications between modules, modules of opposite orientation operating on opposite phases of the clock.
- 25 12. A three-dimensional circuit assembly comprising a plurality of modules having flat surfaces, each surface supporting a pair of like, parallel, extended connector sets, the connector sets on opposite surfaces of the  
30 modules being complementary to each other and

-24-

rotated 90° relative to each other, modules being joined to adjacent modules such that the substrates define a three-dimensional tetrahedral lattice.

- 5 13. A three-dimensional circuit assembly according to Claim 12, wherein the plurality of modules are spaced horizontally within levels such that there are serpentine vertical flow paths between the plurality of modules.
- 10 14. A three-dimensional electronic circuit assembly according to Claim 12 wherein the substrates are addressed within the array by a four element vector.
- 15 15. A three-dimensional electronic circuit assembly as claimed in Claim 12 wherein a plurality of circuit chips are supported by substrates within modules.
- 20 16. A three-dimensional electronic circuit assembly according to Claim 12 further comprising auxiliary circuit boards mounted to substrates within modules.
- 25 17. A three-dimensional electronic circuit assembly according to Claim 12 wherein each module is in one of two orientations, each orientation having different directions of interconnection, each module being connected only to modules of the opposite orientation.

-25-

18. A three-dimensional electronic circuit assembly according to Claim 17 further comprising a two-phase clock input for controlling communications between modules, modules of opposite orientations outputting signals with opposite phases of the clock.
19. A three-dimensional electronic circuit assembly comprising a plurality of modules supporting respective circuits, each module having connectors for connecting to adjacent modules of a three-dimensional array of modules along four axes, adjacent modules being connected along a single common axis but being rotated relative to each other such that connections from adjacent modules are made in opposite directions.
20. A three-dimensional electronic circuit assembly according to Claim 19, wherein the modules are addressed within the array by a four element vector.
21. A three-dimensional electronic circuit assembly according to Claim 19 further comprising a two-phase clock input for controlling communications between modules, modules of opposite orientation operating on opposite phases of the clock.
22. A three-dimensional electronic circuit assembly according to Claim 19 wherein the connectors provide vertical connections.

-26-

23. A three-dimensional electronic circuit assembly comprising a plurality of modules interconnected in a three-dimensional lattice structure which is isomorphic with a diamond lattice.
- 5
24. A three-dimensional electronic circuit assembly comprising a plurality of modules supporting respective circuits, each module having connectors for connecting to four adjacent modules of a three-dimensional array of modules, the modules supporting communication circuitry which addresses each module of the array by a unique four dimensional vector.
- 10
25. A three-dimension electronic circuit assembly according to Claim 24, wherein the circuits of each module are connected vertically.
- 15
26. A three-dimensional electronic circuit assembly comprising a plurality of modules supporting respective circuits, each module having four connector sets for connecting the circuits of adjacent modules, the modules being rotated relative to each other and each being joined to four adjacent modules.
- 20
27. A three-dimensional electronic circuit assembly comprising a plurality of modules supporting respective circuits, each module having three connector sets for connecting the circuits of adjacent modules, the modules being rotated relative to each other and each being joined to three adjacent modules.
- 25
- 30

-27-

28. A three-dimensional electronic circuit assembly according to Claim 27, wherein the modules are triangular.
29. A three-dimensional electronic circuit assembly according to Claim 26 wherein the connectors are formed on a top triangular surface, a bottom triangular surface, and an edge of the module, each extending along an axis transverse to an axis of connection, the transverse axes being displaced 120° relative to each other.
30. A method of providing a three-dimensional electronic circuit assembly, including the steps of:
- providing a plurality of modules;
  - supporting circuits on the plurality of modules; and
  - connecting adjacent modules along fewer than six module axes such that the connected modules are rotated relative to each other to form a three-dimensional array.
31. A method according to Claim 30, further including the step of forcing cooling fluid through serpentine channels between the modules.
32. A method according to Claim 30, further including the step of addressing the modules with a four element vector.

-28-

33. A method according to Claim 30, wherein each of the connected modules communicates with three other modules.
- 5 34. A method according to Claim 30, wherein the connected modules each communicate with four other modules.
35. A method according to Claim 30 wherein the three-dimensional array is isotropic.
- 10 36. A method according to Claim 30 further comprising mounting multiple electronic circuit chips to each module.
37. A method according to Claim 30 further comprising mounting an auxiliary circuit board to each module.
- 15 38. A method according to Claim 30 wherein the modules are connected through vertical connections.
39. A method according to Claim 30 wherein the three-dimensional array is isomorphic with a diamond lattice.
- 20 40. A method according to Claim 30 wherein the substrates within the three-dimensional array assume one of two orientations, the method further comprising clocking a module to communicate in alternate phases of a clock signal according to orientation.
- 25




Fig. 1  
PRIOR ART

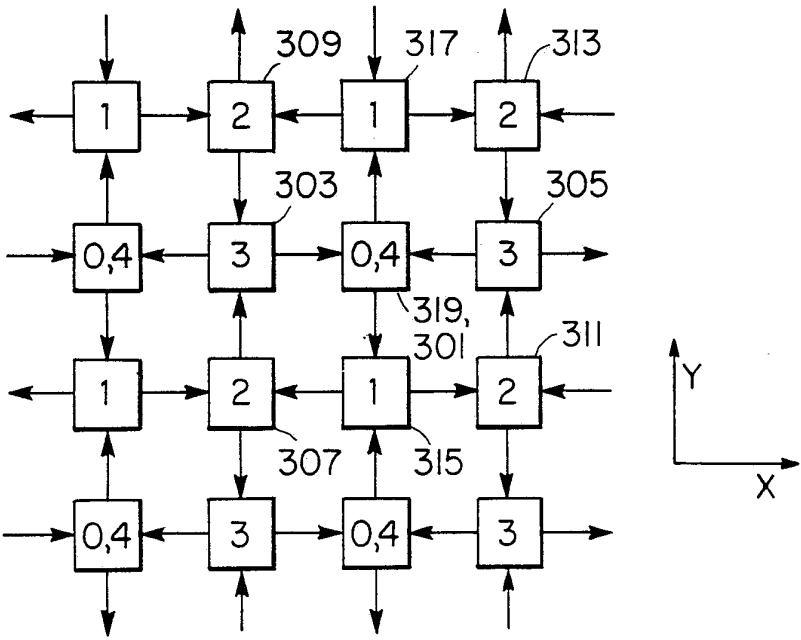


Fig. 3

2/6







Fig. 2



3/6



*Fig. 5b*



4/6








Fig. 4b

5/6

*Fig 6**Fig. 7a**Fig. 7b*

6/6

*Fig. 8*

## INTERNATIONAL SEARCH REPORT

PCT/US 93/02416

International Application No

<b>I. CLASSIFICATION OF SUBJECT MATTER</b> (if several classification symbols apply, indicate all) <sup>6</sup>		
According to International Patent Classification (IPC) or to both National Classification and IPC		
Int.Cl. 5 G06F15/80		
<b>II. FIELDS SEARCHED</b>		
Minimum Documentation Searched <sup>7</sup>		
Classification System	Classification Symbols	
Int.Cl. 5	G06F	
Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched <sup>8</sup>		
<b>III. DOCUMENTS CONSIDERED TO BE RELEVANT<sup>9</sup></b>		
Category <sup>10</sup>	Citation of Document, <sup>11</sup> with indication, where appropriate, of the relevant passages <sup>12</sup>	Relevant to Claim No. <sup>13</sup>
X	WO,A,9 119 258 (WAVETRACER, INC.) 12 December 1991	1-3, 12-13, 19-21, 24-31
A	see page 1, line 19 - line 31	4-11, 14-18, 22-23, 32-40
	see page 3, line 1 - line 23 see page 4, line 12 - page 6, line 36 see figures 1-2 ---	-/--
<p><sup>10</sup> Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"&amp;" document member of the same patent family</p>		
<b>IV. CERTIFICATION</b>		
Date of the Actual Completion of the International Search	Date of Mailing of this International Search Report	
07 JULY 1993	26. 07. 93	
International Searching Authority	Signature of Authorized Officer	
EUROPEAN PATENT OFFICE	NGUYEN XUAN HIEP C.	

III. DOCUMENTS CONSIDERED TO BE RELEVANT (CONTINUED FROM THE SECOND SHEET)		
Category <sup>a</sup>	Citation of Document, with indication, where appropriate, of the relevant passages	Relevant to Claim No.
X	JOURNAL OF VLSI SIGNAL PROCESSING vol. 2, no. 2, October 1990, DORDRECHT NL pages 79 - 87 , XP229128 M.J. LITTLE ET AL. 'The 3-D Computer'	1-2, 12-13, 19,20, 24, 26-27,30
A	see page 79, right column, paragraph 2 - page 81, right column; figure 1	3-11, 14-18, 21-23, 25,28, 31-40
A	US,A,4 742 552 (ANDREWS) 3 May 1988  see column 1, line 7 - column 2, line 8 see column 6, line 34 - column 7, line 15 see column 9, line 22 - line 42	1-2, 12-13, 19-20, 24, 26-27,30
A	DE,A,3 616 821 (W. HILBERG) 19 November 1987 see column 1, line 32 - column 2, line 13 see column 3, line 54 - column 4, line 40 see column 5, line 64 - column 8, line 21 see figures 16-17	1-40

**ANNEX TO THE INTERNATIONAL SEARCH REPORT  
ON INTERNATIONAL PATENT APPLICATION NO.**

US 9302416  
SA 71691

This annex lists the patent family members relating to the patent documents cited in the above-mentioned international search report.  
The members are as contained in the European Patent Office EDP file on  
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

07/07/93

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO-A-9119258	12-12-91	US-A- 5133073	21-07-92
		AU-A- 8077291	31-12-91
		EP-A- 0532663	24-03-93
-----			
US-A-4742552	03-05-88	None	
-----			
DE-A-3616821	19-11-87	None	
-----			