

209,238(87)

PATENTS-US--A7209238⁶⁻²⁶

PATENTS-US--A7209238

DE89 010965

W-7405-ENG-36

Vance Faber
James W. Moore

INTERCONNECTION NETWORKS

209,238

INTERCONNECTION NETWORKS

Inventors: Vance Faber
P.O. Box 345
Los Alamos, New Mexico 87544

James W. Moore
P.O. Box 641
Los Alamos, New Mexico 87544

CITIZENS OF THE UNITED STATES

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

[Handwritten signature]

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

BACKGROUND OF THE INVENTION

5 This invention is generally related to interconnection networks for massively parallel computer systems and, more particularly, to vertex symmetric networks with simple routing schemes through these networks. This invention is the result of a contract with the Department of Energy (Contract No. W-7405-ENG-36).

10 Parallel computer systems are generally composed of many processors which communicate through an interconnection network. This network can be represented by a graph: a collection of nodes which represent the processors and edges which represent the wires which connect pairs of processors. It is highly desirable to be able to construct interconnection networks which have as many nodes as possible, given a fixed degree (the number of wire pairs connected to each node) and diameter (the maximum number of wires a message must traverse to get between any two processors). It is generally shown by W. G. Bridge et al., "On the Impossibility of Directed Moore Graphs," J. of Comb. Theory 29, 339-341 (1980), that a directed graph with degree d and diameter k can have at most $n_{\max} = d^k + d^{k+1} + \dots + d$ nodes. In M. Imase et al., "A Design for Directed Graphs With Minimum Diameter," IEEE Trans. Comp. C-32, 782-784 (1983) graphs are constructed which have n nodes, degree d , and diameter at most $k+1$ for all $n \leq n_{\max}$.

5 In a parallel processor system, it would be desirable for the network to have vertex symmetry, i.e., any node in the network looks essentially like any other node, wherein the routing algorithm is identical for each node in the network. Some of the graphs taught by the Imase reference have vertex symmetry and the desirable properties of vertex symmetry have been noted, see e.g., J. C. Bermond et al., "Strategies for Interconnection Networks: Some Methods from Graph Theory," *J. Parallel Distributed Computing* 3, 433-449 (1986).

10

15 It would also be desirable to easily expand parallel networks to add additional processors. Such a characteristic would enable a core network to be expanded or several core networks to be interconnected. It would be further desirable for message routing through the network to be accomplished using only simple algorithms for determining the routing path through the network. It is also desirable to minimize the number of routing steps which must be traversed along a routing path between a source processor and a destination processor.

20

These problems and others are addressed by the present invention, wherein a vertex symmetric routing network is provided for a plurality of processors to enable a shortest routing path through the network to be obtained.

25 Accordingly, it is an object of the present invention to provide networks having the smallest diameter known for a given degree.

30 Another object of the present invention is to require relatively simple routing algorithms at each processor for transmitting messages through the system.

Yet another object is to provide identical nodes throughout the network, i.e., each node has the same number of interconnection wires and the routing algorithm is identical for each node in the network.

Still another object of the present invention is to enable the subject networks to be expanded by only simple modifications to the networks.

Additional objects, advantages and novel features of the invention will be set forth in part in the description which follows, and in part will become apparent to those skilled in the art upon examination of the following or may be learned by practice of the invention. The objects and advantages of the invention may be realized and attained by means of the instrumentalities and combinations particularly pointed out in the appended claims.

SUMMARY OF THE INVENTION

To achieve the foregoing and other objects, and in accordance with the purposes of the present invention, as embodied and broadly described herein, the apparatus of this invention may comprise, in one characterization, a network for connecting a plurality of processors for cooperative interaction. The network is defined by a plurality of message transmission paths configured in a vertex symmetric graph selected from graphs $\Gamma_d(k)$ with degree d , diameter k , and $(d+1)!(d-k+1)!$ processors for each $d \geq k$ and $\Gamma_d(k-1)$ with degree $d-1$, diameter $k+1$, and $(d+1)!(d-k+1)!$ processors for each $d > k \geq 4$.

In another characterization of the invention, a network is provided for connecting a plurality of processors for cooperative interaction, where each of the processors has an address formed by one of the permutations from a predetermined sequence of letters chosen a selected number at a time, and an extended address formed by appending to the address the remaining ones of the predetermined sequence of letters. A plurality of transmission channels is provided from each

of the processors, where each processor has one less channel than the selected number of letters forming the sequence. Where a network $\Gamma_d(k,-1)$ is provided, no processor has a channel connected to form an edge in a direction δ_1 . Each of the channels has an identification number selected from the sequence of letters and connected from a first processor having a first extended address to a second processor having a second address formed from a second extended address defined by moving to the front of the first extended address the letter found in the position within the first extended address identified by the channel identification number. The second address is then formed by selecting the first elements of the second extended address corresponding to the selected number used to form the address permutations.

In yet another characterization of the present invention a process is provided for forming a routing path between a source processor address and a destination processor address. The plurality of processors are connected in a network defined by a vertex symmetric graph selected from graphs $\Gamma_d(k)$ with degree d , diameter k , and number of processors $(d+1)!/(d-k+1)!$, $d \geq k$ and $\Gamma_d(k,-1)$ with degree $d-1$, diameter $k+1$, and a number of processors $(d+1)!/(d-k+1)!$, $d \geq k \geq 4$. An address is assigned to each processor formed by one of the permutations of a sequence of letters chosen a selected number at a time and an extended address is formed by appending to the address the remaining ones of the sequence of letters. A plurality of transmission channels is then provided from each processor, with each channel having an identification number selected from the set of numbers one less than the number of letters in the

sequence. Each channel is connected from one processor having a first extended address to a second processor having a second address formed from a second extended address defined by moving to the front of the first 5 extended address the letter found in the position within the first extended address identified by the channel identification number. The second address is then formed by selecting the first selected number of letters in the second extended address, provided that the network formed 10 by said graph $\Gamma_d(k,-1)$ has no channel connections corresponding to edges of said graph $\Gamma_d(k)$ having a direction δ_1 .

BRIEF DESCRIPTION OF THE DRAWINGS

15 The accompanying drawings, which are incorporated in and form a part of the specification, illustrate the embodiments of the present invention and, together with the description, serve to explain the principles of the invention. In the drawings:

20 FIGURE 1 is a table showing the number of processors available in networks formed from graphs $\Gamma_d(k)$ according to the present invention.

FIGURE 2 is a table showing the number of processors available in networks formed from graphs $\Gamma_d(k,-1)$ according to the present invention.

25 FIGURE 3 is a vertex symmetric network defined by graph $\Gamma_3(3)$ with illustrative node addresses and channel connections.

FIGURE 4A is a flow chart for computing a full path routing.

30 FIGURE 4B is a flow chart for locating an initial caret position.

FIGURE 4C is a flow chart for PULL to form a new address from a current address.

FIGURE 4D is a flow diagram showing processing done at a node to forward a message.

FIGURE 4E is a flow diagram for special case routing from the diagonal band.

5 FIGURE 4F is a flow diagram for special case processing, CASE 1 and CASE 2.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

In accordance with the present invention, a network of processors is interconnected to provide the smallest maximum number of wire paths a message must traverse to get between any two processors for a given number of wire pairs connected to each processor. The network is configured to form vertex symmetric graphs $\Gamma_d(k)$ with degree d (the number of wire pairs connected to each processor, or node), diameter k (the maximum number of wires a message must traverse to get between any two processors), and $(d+1)!/(d-k+1)!$ nodes (processors) for each $d \geq k$ or graphs $\Gamma_d(k,-1)$ with degree $d-1$, diameter $k+1$, and $(d+1)!/(d-k+1)!$ nodes (processors) for each $d \geq k \geq 4$.

The present invention includes the following definitions and theorems which provide the necessary system conditions:

(1) Definition. Let G be a group, H a subgroup, and 25 Δ a subset. Suppose (i) $G = \langle \Delta \cup H \rangle$, (ii) $H\Delta H \subseteq \Delta H$, (iii) Δ is a set of distinct nonidentity coset representatives of H in G . Then the Cayley coset graph $\Gamma = (G, \Delta, H)$ can be formed with vertices $\{gH : g \in G\}$ and an edge (g_1H, g_2H) whenever $g_2H = g_1\delta H$ for some $\delta \in \Delta$.

(2) Fact. Suppose $g'_2H = g_2H$ and $g'_1H = g_1H$. Then there exists $\delta' \in \Delta$ such that $g'_2H = g'_1\delta'H$.

(3) Corollary. Γ is connected and the degree is $|\Delta|$.

(4) Definitions. Let G be any k -transitive group on $d+1$ letters; $T_k = \{0, 1, 2, \dots, k-1\}$. Choose $\delta_\alpha \in G$ 5 ($\alpha \neq 0$) so that

$$\delta_\alpha(0) = \alpha,$$

$$\delta_\alpha(j) = \begin{cases} j-1 & \text{for } 1 \leq j \leq \alpha \\ j & \text{for } j > \alpha \end{cases}, \quad 1 \leq j \leq k-1.$$

Let $\Delta = \{\delta_\alpha \mid \alpha \neq 0\}$ and $H = \{h \mid h_{T_k} = \text{identity}\}$.

Cosets of H in G are denoted by k -tuples $(g_0, g_1, \dots, g_{k-1})$ with all g_i distinct since each 10 coset gH is completely determined by its action on T_k .

(5) Theorem. $H(\Delta H) \subseteq \Delta H$.

(6) Corollary. The Cayley coset graph $\Gamma = (G, \Delta, H)$ can be formed with the number of nodes $(d+1)!/(d+1-k)!$ and degree d .

15 (7) Definition. Given $(x_0, x_1, \dots, x_{k-1})$, let $c(x_0, x_1, \dots, x_{k-1}) = |\{x \in T_k \mid x \neq x_i \text{ and } x > x_{k-1}\}|$, where $T_k = \{0, 1, 2, \dots, k-1\}$, $0 \leq i \leq k-1$.

20 (8) Lemma. If there is a path of length j , $j = 0, 1, \dots, k-1$ from $R = (0, 1, 2, \dots, k-1)$ to P then P has the form $(*) (x_0, x_1, \dots, x_{k-1})$ with $x_i \in T_k$ for $i \geq j$.

$$x_j < x_{j+1} < \dots < x_{k-1} \text{ and} \\ |\{x_i \mid i < j \text{ and } x_i \in T_k\}| \leq c(x_0, x_1, \dots, x_{k-1}).$$

(9) Theorem. There is then, a unique shortest path of length at most k from $R(0, 1, 2, \dots, k-1)$ to any vertex.

25 (10) Theorem. If G_1 and G_2 are k -transitive groups on $d+1$ letters, then the graphs $\Gamma(G_1, \Delta_1, H_1)$ and $\Gamma(G_2, \Delta_2, H_2)$ described in Corollary (6) are isomorphic graphs with $(d+1)!/(d+1-k)!$ vertices.

(11) Theorem. For each $k \leq d$, there exist vertex symmetric graphs $\Gamma_d(k)$ with degree d , diameter k , and $(d+1)!/(d+1-k)!$ nodes.

As provided by Theorem (9), the node interconnections herein described provide a unique shortest path between any two nodes. Since $\Gamma_d(k)$ is vertex transitive, all nodes look alike and it is sufficient to discuss only one source node, e.g. $R=(0,1,2,\dots,k-1)$, with a destination node $P=(x_0, x_1, \dots, x_{k-1})$. The shortest routing path occurs in two phases.

Phase I.

Computation of the smallest j such that STAR [defined as (*) in Lemma (8)] holds.

If $x_{k-1} \notin T_k$, set $j=k$ and terminate. Otherwise, compute CEE [defined as $c(x_0, x_1, \dots, x_{k-1})$ in Definition (7)], which is independent of j . Next, count $|\{x_i | x_i \notin T_k\}|$. If this is larger than CEE, set $j=k$ and terminate. Otherwise set j smallest so that $x_j < x_{j+1} < \dots < x_{k-1}$.

Phase II. Routing.

Routing is completely determined by the destination address components $(x_0, x_1, \dots, x_{j-1})$. Indeed, after ℓ steps, all that is needed is the current address $(y_0, y_1, \dots, y_{k-1})$ and the destination address truncated to $(x_0, x_1, \dots, x_{j-\ell-1})$. The message itself and further truncated destination address $(x_0, x_1, \dots, x_{j-\ell-2})$ is sent to the node whose address is formed by appending $x_{j-\ell-1}$ to the beginning of $(y_0, y_1, \dots, y_{k-1})$ and then (A) either removing $x_{j-\ell-1}$ from the list $(y_0, y_1, \dots, y_{k-1})$ if it occurs, or (B) removing y_{k-1} if $x_{j-\ell-1}$ does not occur in the list. When (A) happens, the output direction is δ_α , where α marks the position of $x_{j-\ell-1}$ in

the current address, and when (B) happens, then the output direction is determined by the choice of coset representative used to label the node. The convention herein used is that the δ_α for $\alpha \geq k$ point toward addresses whose leading digits increase with α .

5 Routing from a general source node $S=(s_0, \dots, s_{k-1})$ to a destination $P=(x_0, \dots, x_{k-1})$ is no more complex than routing from R. It is only necessary to redefine T_k to be the totally ordered set $\{s_0, s_1, \dots, s_{k-1}\}$ and to define CEE relative to T_k , where $CEE = |\{x \in T_k \mid x \neq x_i \text{ and } x \text{ follows } x_{k-1} \text{ in } T_k\}|$. The routing proceeds as outlined above.

10 It should be noted that there is always a route from R to P which takes exactly k steps (the "constant distance" route) unless the initial letter of the source address matches the last letter of the destination address. This route is formed by setting $j=k$ in Phase I regardless of its actual value and following the routing procedure in 15 Phase II.

20 It can be shown that in any of the graphs $\Gamma_d(k)$ with $d \geq k \geq 4$, it is possible to find a route between any two nodes which has length at most $k+1$ and does not use the direction δ_1 . By omitting δ_1 from $\Gamma_d(k)$, 25 the graphs $\Gamma_d(k-1)$ are formed with degree $d-1$, diameter $k+1$, and $(d+1)!/(d+1-k)!$ nodes, and this route is the constant distance route, except in three special cases. These cases are when (I) $x_{k-1}=0$, (II) $x_{k-1}=1$, (III) $x_{k-1} \neq 1$, $x_{k-2}=0$.

30 In these cases, there can be found a route which avoids δ_1 , as follows.

Case I.

If $x_{k-2} \neq 1$, the route is $(0, 1, \dots, k-1) \rightarrow (x_{k-2}, 0, 1, *, \dots, *) \rightarrow \dots \rightarrow (x_1, x_2, \dots, x_{k-2}, 0, *) \rightarrow$

$(x_0, x_1, \dots, x_{k-2}, 0)$. If $x_{k-2} = 1$, the route is
 $(0, 1, \dots, k-1) \rightarrow (x_0, 0, 1, *, \dots, *) \rightarrow (1, x_0, 0, *, \dots, *)$
 $\rightarrow \dots \rightarrow (*, \dots, *, x_{k-3}, 1, x_0, 0) \rightarrow$
 $(x_0, x_1, \dots, x_{k-3}, 1, 0)$.

5

Case II.

If $x_0 \neq 0$, the route is $(0, 1, \dots, k-1) \rightarrow$
 $(x_0, 0, 1, *, \dots, *) \rightarrow (1, x_0, 0, *, \dots, *) \rightarrow \dots \rightarrow$
 $(\dots, x_{k-3}, x_{k-2}, 1, x_0) \rightarrow (x_0, x_1, \dots, x_{k-2}, 1)$.
If $x_0 = 0$, the route is $(0, 1, 2, \dots, k-1) \rightarrow$
10 $(x_{k-2}, 0, 1, *, \dots, *) \rightarrow (x_{k-3}, x_{k-2}, 0, 1, *, \dots, *) \rightarrow$
 $\dots \rightarrow (x_1, x_2, \dots, x_{k-2}, 0, 1) \rightarrow$
 $(0, x_1, x_2, \dots, x_{k-2})$.

10

Case III.

Since $x_{k-1} \neq 1$, the route is $(0, 1, 2, \dots, k-1) \rightarrow$
15 $(x_{k-1}, 0, 1, *, \dots, *) \rightarrow (x_0, x_{k-1}, 0, *, \dots, *) \rightarrow$
 $(0, x_0, x_{k-1}, *, \dots, *) \rightarrow (x_{k-3}, 0, x_0, x_{k-1}, *, \dots, *) \rightarrow$
 $\dots \rightarrow (x_1, x_2, \dots, x_{k-3}, 0, x_0, x_{k-1}) \rightarrow$
 $(x_0, x_1, x_2, \dots, x_{k-3}, 0, x_{k-1})$.

15

In these cases, the direction δ_1 is avoided, while
20 increasing the diameter at most one step. Thus,
 $\Gamma_d(k, -1)$ denotes the graph $\Gamma_d(k)$ with the edges in
direction δ_1 removed, and $\Gamma_d(k, -1)$ is vertex
symmetric and its diameter is $k+1$.

20

The number of processors, or nodes, in each of the
25 vertex symmetric directed graphs, $\Gamma_d(k)$, is shown in
Figure 1 for degree d , and diameter, $k \leq d$. For the set
of graphs, $\Gamma_d(k, -1)$, which have degree $d-1$ and
diameter $k+1$, the number of connected processors is
 $(d+1)!/(d+1-k)!$, as shown in Figure 2 for $d \geq k \geq 4$.
30 Then, if the networks defined by Theorem (11) are combined
with $\Gamma_d(k, -1)$, for each $k \leq d+1$, there exist vertex
symmetric graphs Γ with degree d , diameter k , and
 $\max((d+1)!/(d+1-k)!, (d+2)!/(d-k+3)!)$ nodes. This can be
seen by comparing Figures 1 and 2.

It is believed that the networks constructed according to the above theory are maximally fault tolerant. Further, the networks are expandable so that $d+1$ copies of $\Gamma_d(k)$ can be connected to form $\Gamma_{d+1}(k+1)$.
5 Additionally, the number of virtual channels per real channel needed to break deadlocks is certainly no larger than the diameter.

To illustrate a network constructed according to the above principles, Figure 3 depicts a 24 node network
10 ($\Gamma_3(3)$) constructed from four 6 node networks ($\Gamma_2(2)$). Each node has three wire pairs connecting the node with the network and, at most, three paths must be traversed to connect any node to any other node in the network. An illustrative routing from source node (0 1 2)
15 to destination node (2 0 3) is as follows.

Phase I. The last letter (3) of the destination address is not in the source address, so there is no short route.

Phase II. The first intermediate node address is formed by appending the last letter (3) of the destination address to the beginning of the source address to form (3 0 1 2) and then truncating to the address length (3 0 1). The header number (2 0) is transmitted to node (3 0 1). The last number (0) of the header is then appended to the front of the present address to form the next intermediate address (0 3 1), where the appended letter is present and is therefore dropped from the present address. The last letter (0) of the header is dropped and the header (2) is transmitted to address (0 3 1). Again the last letter (2) of the header is appended to the front of the current address to form the destination address (2 0 3), where the
20
25
30

header letter is not present in the current address and the last letter is dropped. Thus, the following route is indicated: (0 1 2) \rightarrow (3 0 1) \rightarrow (0 3 1) \rightarrow (2 0 3).

To construct a network according to the present invention, the number of nodes, or processors, must be determined and each node is identified by a sequence of letters, usually integers, to form an address for the node. Each address is one of the permutations of a predetermined set of letters (the "set" T_k), chosen some fixed number at a time corresponding to the address length. For example, Figure 3 is constructed from the set {0 1 2 3} taken three at a time to form the addresses depicted, e.g. 012, 102, 320, etc., where no two address letters are identical. It should be noted that each node also has an extended address from the node address, where the set letters which are not in the node address are added on in some standard manner. The standard manner used herein is the same sequence as in the set. Thus, the extended address of node 012 is 0123; for node 301, it is 3012.

In the following discussion, a convention is also needed to identify the position of a letter in an address to determine the routing directions, or "channels", as noted in the discussion above on Phase II routing. As used herein, the zero position is the first letter in the address. Thus, in the string of letters 23140, the zero element is 2, the first element is 3, the second element is 1, etc.

Each node must now be connected to other nodes. The number of channels for each node is one less than the number of letters in the set. The address of the node connected to each channel is determined as follows: for channel K , $K \geq 1$, remove the letter in the K th position in

the extended address of the source node and place the letter at the front of the extended address of that node, starting the count at the zero position. Connect channel K to the node whose address is the resulting first set of 5 letters truncated to the number of letters in the source node.

Referring again to Figure 3, it is seen, for example, that the following nodes are connected to source node
213: Extended address - 2 1 3 0

Position	Node connection
1	1 2 3
2	3 2 1
3	0 2 1

15 By way of further example, given a source node address 1 4 5 2 3 and a letter set {0 1 2 3 4 5 6 7}, the extended source node address is 1 4 5 2 3 0 6 7. The node connected through channel 1 is determined by taking the letter in the first position, 4, and placing it in front 20 of the source node extended address to form 4 1 5 2 3 0 6 7, with the node address formed from the first five letters, or 4 1 5 2 3. Likewise channel 6 is connected to node 6 1 4 5 2.

Tracing a path from a source node to a destination 25 node requires only the source node address and the destination node address. Starting with the right hand letter of the destination address, that letter is found in the extended source address and placed at the front of the source address to find a first intermediate node. The 30 position of that letter in the source address determines the appropriate channel for transmission. The next right hand letter of the destination is then removed from the extended address of the first intermediate node and placed

at the front to form a second intermediate node address. Again, the channel corresponds to the position of the selected letter in the first intermediate node address.

Example: Source - 1 3 4 2 5
 5 Destination - 2 6 7 3 4
 Extended
 Address - 1 3 4 2 5 0 6 7
 Channel To
 2 4 1 3 2 5
 10 2 3 4 1 2 5
 7 7 3 4 1 2
 7 6 7 3 4 1
 6 2 6 7 3 4.

In this example, the full network is not known, only that the letter set has at least 8 elements {0 1 ... 7}. Communication between two nodes takes place in the sub-network which is "spanned" by the letters which make up the node addresses. Thus, "0" never appears in this routing. However, the letter "0" must appear in the extended address to enable channel numbers to be computed.

It will be appreciated that the routing between processors in a network formed according to the present invention can be done in a process which is the same at each processor, or node. In a preferred embodiment, this processing is done according to the flow charts shown in Figures 4A, 4B, 4C, 4D, 4E, and 4F. The flow charts can be implemented in relatively simple computer software by a person of ordinary programming skill on a selected computer system. The only information required is the current node address and the final destination address truncated as herein above described.

Referring now to Figures 4A and 4B, there are shown flow charts for routing a message along a full routing

path. PATH 20 is initialized for computing the full message route. Then the CARET 22 is located for determining the first header from which the initial channel and connecting node is determined for routing from 5 the source node. Once CARET 22 is located, ROUTING 24 is performed. PULL 26 (Figure 4C) removes the letter immediately to the left of the caret and places that letter at the front of the current node address to generate the next node address. The new node address is 10 saved and the caret is moved one place to the left 28 to form the header used with the new node address to determine a subsequent node address. It is then determined whether the caret has moved through the original header. If not, the program is returned 34 to 15 ROUTING 24 to determine subsequent node addresses. If yes, PATH 20 has been completed.

The routine for determining CARET 22, the initial caret position, is shown in Figure 4B. When CARET 22 is called, a SET 42 of letters is first created, formed from 20 the set of nonnegative integers up to the maximum address element. The destination address is compared 44 with the source address. If the last element of the destination address is not in the source address, the caret is placed to the right of the destination address and the program 25 returns 46 to ROUTING 24. If the last element is present, CEE is set 48 to the count of the number of source address elements after the last element of the destination address which are not also in the destination address. CEE can range from zero to one less than the number of letters in 30 the address.

To determine whether a short path is available, STAR is set 52 to the count of elements in the destination address which are not also in the source address. CEE and

STAR are compared 54. If CEE and STAR are not equal, no short path exists and the program sets the caret to the right of the destination address elements and the program returns 56 to ROUTING 24. If CEE and STAR are equal, the 5 initial caret location in the destination address is determined in steps 58-66.

The source address is first placed in extended format 58 and the destination address is looked up 62 in the extended source address. To "look up" an address in 10 another address is to determine the location of the second address letters in the first extended address to derive a number comprising the position sequence. A caret is moved to the left through the position sequence and the destination address 64. Each position sequence number is 15 compared 66 with the proceeding number and the caret is moved 64 to the left until a larger number is found. The program is returned to CARET 22 with the caret located in the destination address at the same location determined for the position sequence caret in process steps 64 and 66. ROUTING 24 is then initiated with the header 20 determined by the caret in the moved position.

For example, for source address 3 2 4 5 1 , destination address 0 5 1 3 2 , and set {012345678}, the extended source address is 3 2 4 5 1 0 6 7 8. Look up 25 62 determines the position sequence for the destination address to be 5 3 4 0 1. As the caret 64 moves left through the position sequence, comparison 66 determines that the first two index numbers are in decreasing order. The caret is also moved two places left in the destination 30 address to form header 0 5 1 and ROUTING 24 is started with this header to form the path:

	NODE	HEADER	CHANNEL
	3 2 4 5 1	0 5 1	4
To	1 3 2 4 5	0 5	4
To	5 1 3 2 4	0	5
5	To 0 5 1 3 2		

10 In ROUTING 24, PULL 26 is used to form the sequence of addresses (see above example). The current node address is first extended 74. The letter immediately to the right of the caret, i.e., the last letter of the current header, is removed 76 in the current extended node address and added 78 to the beginning of the remaining address. The amended extended address is truncated 82 to the original address length to form a new address which is returned to the new address 28 (Figure 1).

15 In another embodiment, the individual nodes can be provided with simple look up tables to determine the message transmission channel from the last element of the current header. If the position sequence of the letter set in the extended node address is stored, then the 20 letter located at the position determined by the last letter of the header is the channel for transmission of the message and the new header. The NODE processing 90 shown in Figure 4D first removes 92 the last element of the received header. The removed element indexes 94 into the stored position sequence which is stored at the node 25 to find the channel number at the location identified by the removed element. The node then transmits 96 to the node connected to the indexed channel.

30 For example, select set {01234567} and a node address 1 3 4 2 5. The extended address is 1 3 4 2 5 0 6 7 and the set position sequence is then 5 0 3 1 2 4 6 7. If a message header has the trailing element 4, the location sequence is indexed to the fourth position to find the transmission channel, channel 2.

As shown by comparing Figure 1 and Figure 2, the $\Gamma_d(k,-1)$ network includes a band around the diagonal with a larger number of processors at the expense of one additional routing step. Routing in this band is complicated since the networks in the band have two-way connections between pairs of nodes. By removing one of the two-way connections, i.e. removing connections between channels having direction δ_1 in the network graph $\Gamma_d(k)$, the degree of the node is reduced by one at the expense of increasing the diameter by one to form the network graph $\Gamma_d(k,-1)$. The normal routing process can not proceed, however, if the removed connections are used in the routing. Since the $\Gamma_d(k,-1)$ network connections become the same as the $\Gamma_d(k)$ connections after a few initial steps, the routing process must provide for initial movements which does not use the channels which are deleted. It should be noted that although the channels are actually removed from the network, no renumbering occurs in order to preserve the numbering schemes and the algorithms, i.e., the channel numbers exist but are never called.

There are three special cases with initial processing determinations shown in Figures 4E and 4F which are determined from the relationship of the destination processor address to the source processor address. A special case routing determination is first made for diagonal band routing PATH2 100 by comparing the last one or two elements of the destination node address with the first one or two elements of the source node address. It is first determined if the last element of the destination can also be found in the first two elements of the source. If so, the last element of the destination is

compared 114 to the first element of the source. If the elements are the same, CASE ONE 120 is executed; if unequal, CASE TWO 130 is executed (Figure 4F).

5 If the last two elements of the destination are not the first two elements of the source, the next to last element of the destination is compared 116 to the first element in the source. If the elements are not equal, the processing returns to PATH 20 (Figure 4A). If the elements are equal, CASE THREE 140 is executed. As 10 hereinbelow used, each of the PULL steps use the PULL 26 routine set out in Figure 4C. The last element of the destination is first PULLED 142 to the source and the first element of the destination is then PULLED 144 to the source to form first and second intermediate addresses, 15 respectively. The second intermediate address forms the initiation address for the common routing algorithm ROUTING 24 (Figure 1). The caret is moved 146 one place to the left in the destination address to form a header from the above initiation address and the program returns 20 to ROUTING 24 (Figure 1).

Referring now to Figure 4F, if CASE ONE 120 is indicated, the next to last element of the destination is then compared 122 to the second element of the source. If the elements are not equal, the process returns to PATH 25 20. Otherwise, the first element of the destination is PULLED 124 to the source and the second element of the original source is PULLED 126 to the address formed at step 124 to form first and second intermediate addresses, 30 respectively. Again, the second intermediate address is the initiation address. The caret is moved 128 two places to the left in the destination address to form a header from the above initiation address and the process is returned to ROUTING 24.

If CASE TWO 130 is indicated, the first element of the source is compared 132 with the first element of the destination. If the two elements are not equal, the process is returned to PATH 20. Otherwise, the first element of the destination is PULLED 134 to the original source and the second element of the original source is PULLED 136 to the address formed by step 134 to form first and second intermediate addresses, respectively. The caret is moved one place to the left in the destination address to form a header from the above second, or initiation, address and the process is returned to ROUTING 24.

The following examples illustrate application of the three special cases:

CASE TWO

	Source	1 5 3 0 4 2
	Destination	4 2 3 0 1 5
	Step	Action
5	112	Last element in D (5) is in S first two (5 1).
	114	Last element in D (5) is not S first (1). CASE TWO.
	132	First element of S (1) is not first element of D (4).
	134	PULL D first to S to form 4 1 5 3 0 2.
	136	PULL S second to 134 address to form 5 4 1 3 0 2.
10	138	Move caret one place in D to form header 4 2 3 0 1.
	24	Routing from address 5 4 1 3 0 2, header 4 2 3 0 1.

Thus, the routing is:

From	1 5 3 0 4 2
To	4 1 5 3 0 2 (first intermediate address)
15	To 5 4 1 3 0 2 (initiation address)
	To 1 5 4 3 0 2 (ROUTING 24 using header 4 2 3 0 1)
	To 0 1 5 4 3 2
	To 3 0 1 5 4 2
	To 2 3 0 1 5 4
20	To 4 2 3 0 1 5 Destination

CASE THREE

	Source	1 3 5 0 4 2
	Destination	4 2 3 0 1 5
	Step	Action
25	112	Last element D (5) not first two elements of S (1 3).
	116	Next to last D (1) same as S first (1).
	140	CASE THREE
	142	PULL D last to S to form 5 1 3 0 4 2.
	144	PULL D first to 142 address to form 4 5 1 3 0 2.
30	146	Move caret in D to form header 4 2 3 0 1.
	24	ROUTING from 4 5 1 3 0 2 with header 4 2 3 0 1.

Thus, the path is :

From 1 3 5 0 4 2
To 5 1 3 0 4 2 (first intermediate address)
To 4 5 1 3 0 2 (initiation address)
5 To 1 4 5 3 0 2 (ROUTING 24 using header 4 2 3 0 1)
To 0 1 4 5 3 2
To 3 0 1 4 5 2
To 2 3 0 1 4 5
To 4 2 3 0 1 5 Destination.

10 In accordance with the present invention as described above, a new interconnected network of processors has been formed in vertex symmetric directed graphs, $\Gamma_d(k)$, which have degree d , diameter $k \leq d$, and a number of nodes, or processors, $(d+1)!/(d+1-k)!$. Further, a process 15 for routing messages through the network defined by the graphs has been provided which is simple and which yields the unique shortest paths between any two nodes. As special cases, interconnected networks formed in vertex symmetric directed graphs, $\Gamma_d(k-1)$ for $d \geq k \geq 4$, 20 which have degree $d-1$, diameter $k+1$, and $(d+1)!/(d+1-k)!$ nodes, are defined by omitting one edge from each node in $\Gamma_d(k)$.

25 The foregoing description of the preferred embodiments of the invention have been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed, and obviously many modifications and variations are possible in light of the above teaching. The embodiments were chosen and described in order to best 30 explain the principles of the invention and its practical application to thereby enable others skilled in the art to best utilize the invention in various embodiments and with

various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto.

ABSTRACT OF THE INVENTION

A network of interconnected processors is formed from a vertex symmetric graph selected from graphs $\Gamma_d(k)$ with degree d , diameter k , and $(d+1)!/(d-k+1)!$ processors for each $d \geq k$ and $\Gamma_d(k-1)$ with degree $d-1$, diameter $k+1$, and $(d+1)!/(d-k+1)!$ processors for each $d \geq k \geq 4$. Each processor has an address formed by one of the permutations from a predetermined sequence of letters chosen a selected number of letters at a time, and an extended address formed by appending to the address the remaining ones of the predetermined sequence of letters. A plurality of transmission channels is provided from each of the processors, where each processor has one less channel than the selected number of letters forming the sequence. Where a network $\Gamma_d(k-1)$ is provided, no processor has a channel connected to form an edge in a direction δ_1 . Each of the channels has an identification number selected from the sequence of letters and connected from a first processor having a first extended address to a second processor having a second address formed from a second extended address defined by moving to the front of the first extended address the letter found in the position within the first extended address defined by the channel identification number. The second address is then formed by selecting the first elements of the second extended address corresponding to the selected number used to form the address permutations.

$$\Gamma_d(k)$$

$d \setminus k$	2	3	4	5	6	7	8	9	10	11	12
2	6										
3	12	24									
4	20	60	120								
5	30	120	360	720							
6	42	210	840	2 520	5 040						
7	56	336	1 680	6 720	20 160	40 320					
8	72	504	3 024	15 120	60 480	181 440	362 880				
9	90	720	5 040	30 240	151 200	604 800	1 814 400	3 628 800			
10	110	990	7 920	55 440	332 640	1 663 200	6 652 800	19 958 400	39 916 800		
11	132	1 320	11 880	95 040	665 280	3 991 680	19 958 400	79 833 600	239 500 800	479 001 600	
12	156	1 716	17 160	154 440	1 235 520	8 648 640	51 891 840	259 459 200	1 037 836 800	3 113 510 400	6 227 020 800
13	182	2 184	24 024	240 240	2 162 160	17 297 280	121 080 960	726 485 760	3 632 428 800	14 529 715 200	43 589 145 600

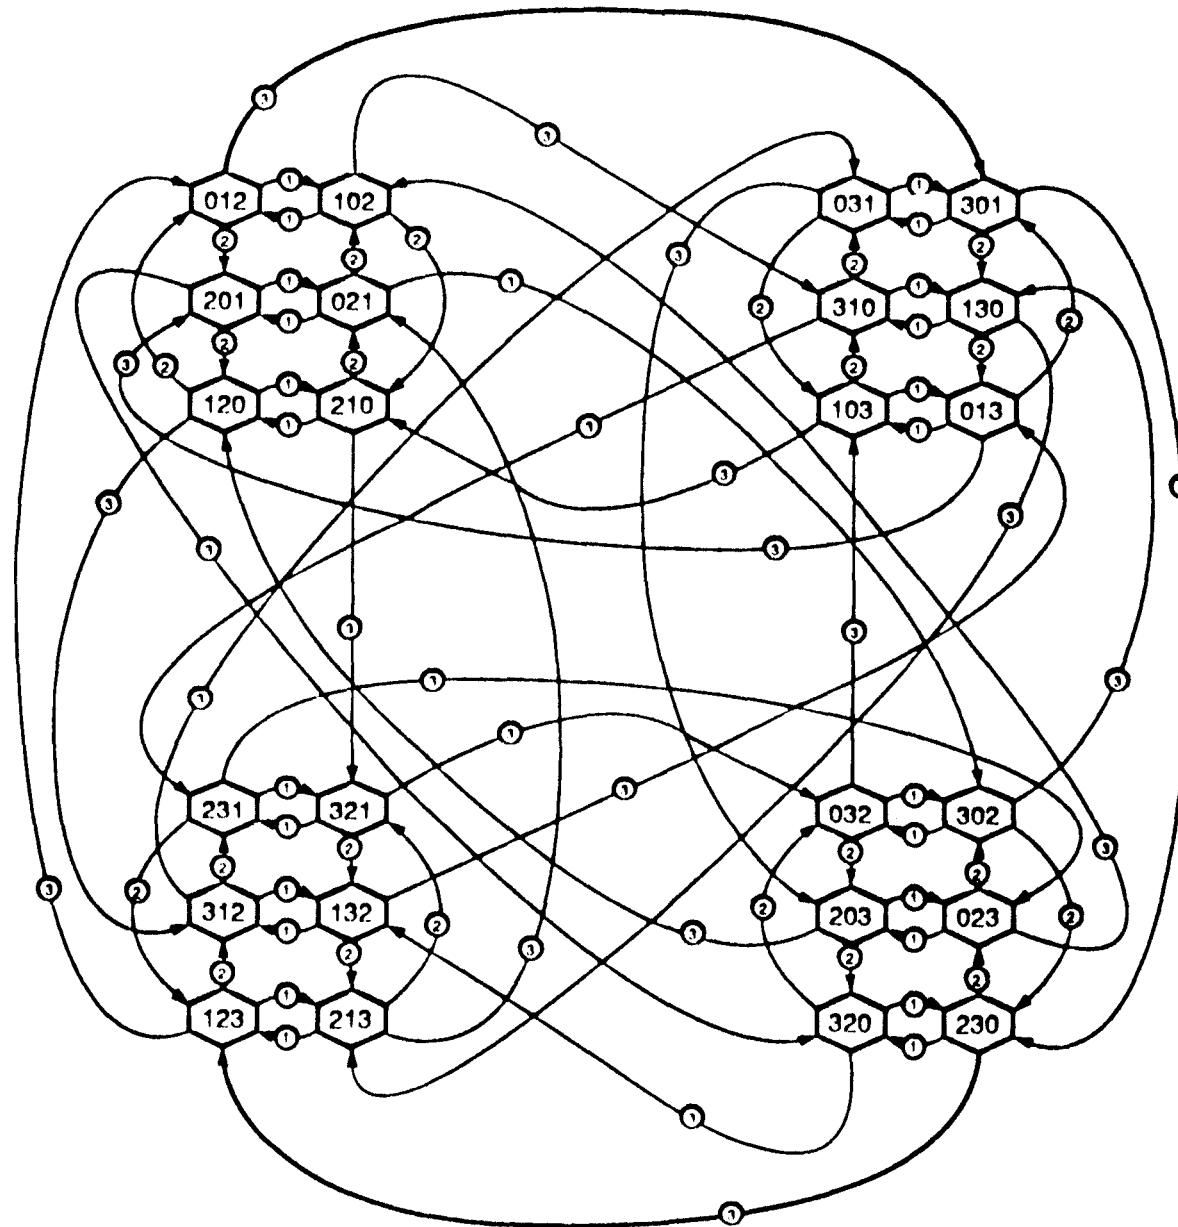
Fig. 1

$$\Gamma_d(k, -1)$$

$\frac{k+1}{d-1}$	5	6	7	8	9	10	11	12	13
3	120								
4	360	720							
5	840	2 520	5 040						
6	1680	6 720	20 160	40 320					
7	3 024	15 120	60 480	181 440	362 880				
8	5 040	30 240	151 200	604 800	1 814 400	3 628 800			
9	7 920	55 440	332 640	1 663 200	6 652 800	19 958 400	39 916 800		
10	11 880	95 040	665 280	3 991 680	19 958 400	79 833 600	239 500 800	479 001 600	
11	17 160	154 440	1 235 520	8 648 640	51 891 840	259 459 200	1 037 836 800	3 113 510 400	6 227 020 800
12	24 024	240 240	2 162 160	17 297 280	121 080 960	726 485 760	3 632 428 800	14 529 715 200	43 589 145 600

Fig. 2

Fig. 3



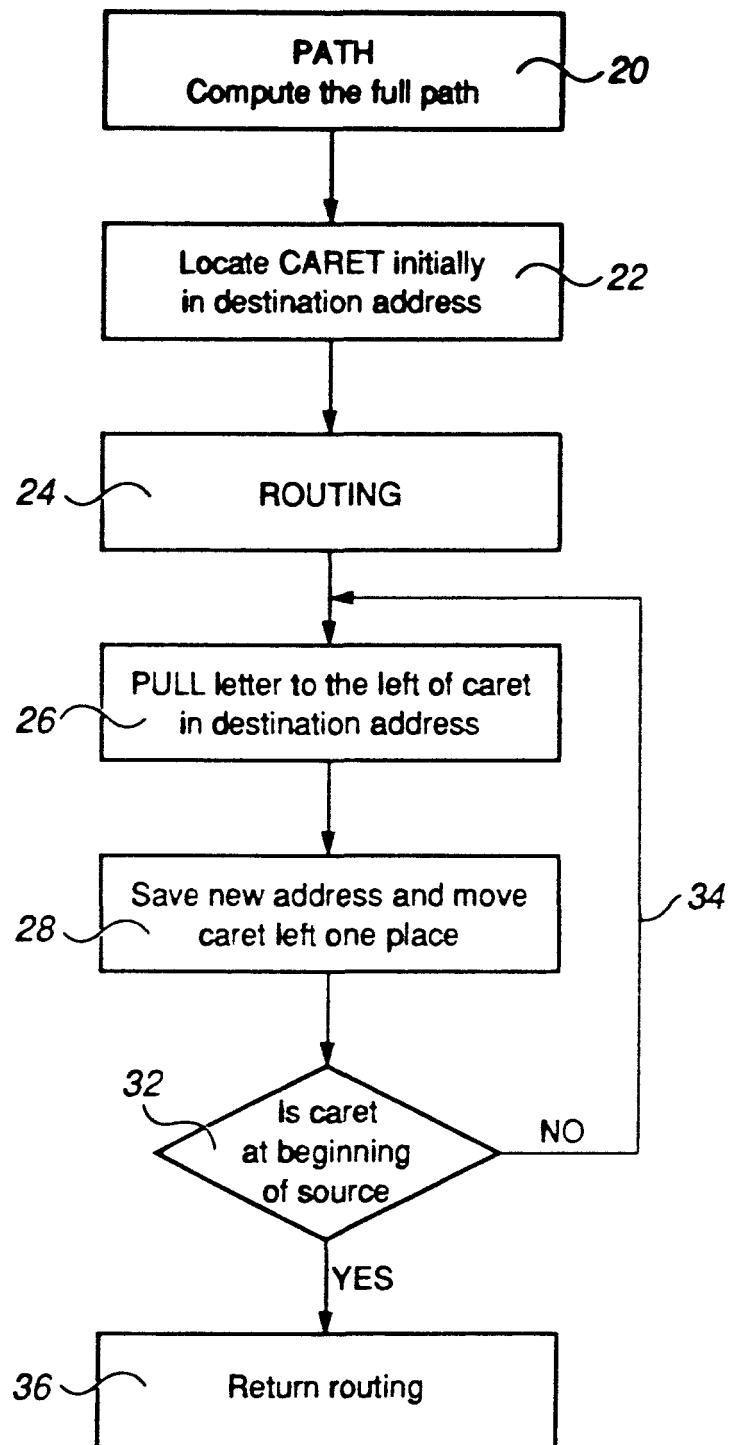
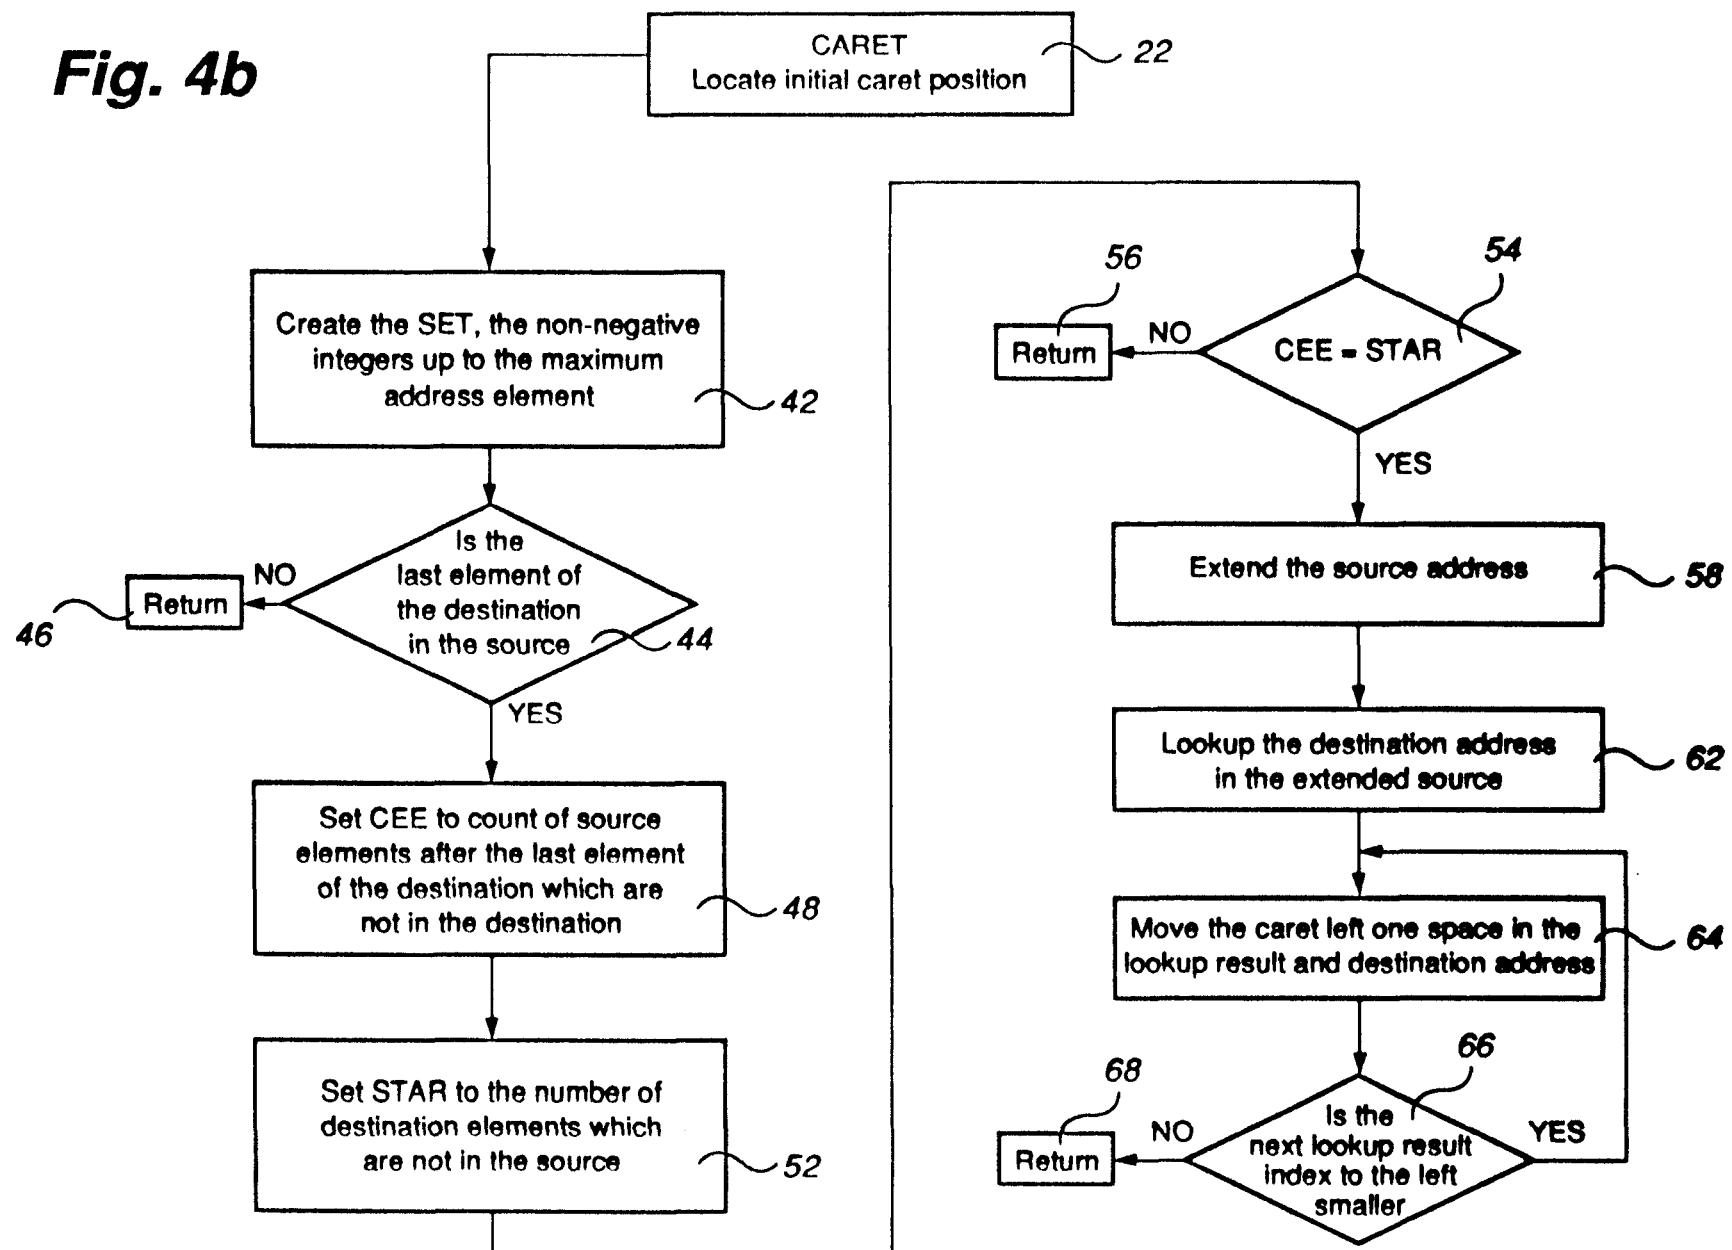


Fig. 4a

Fig. 4b



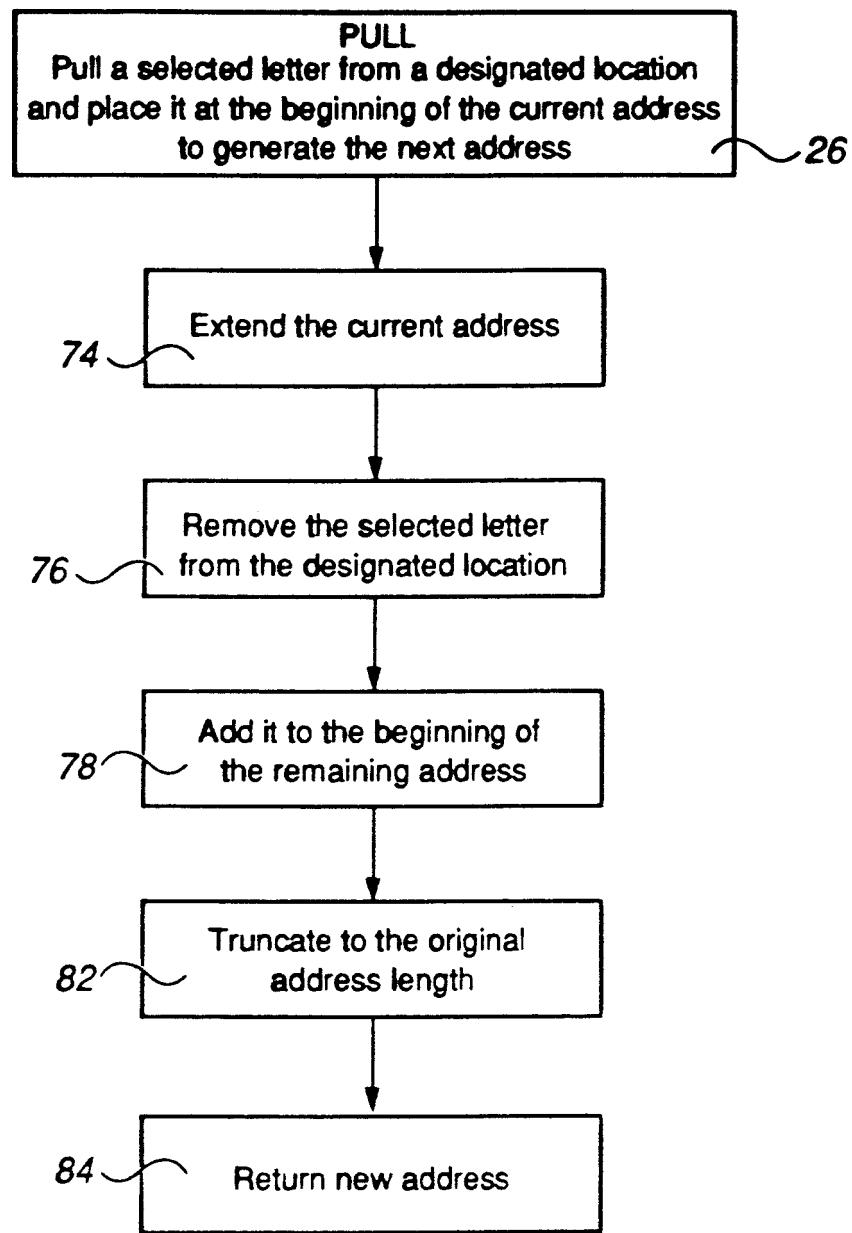


Fig. 4c

NODE

Processing done at a node to forward a message

The actual node address is not used by the node processor.
A coded form is used to simplify the node processing. The coded form is the channel number associated with each element of the node address.
This enables simple indexing to be done rather than table lookup.

90

Remove the last element of the destination address

92

Use this element as the index into the coded source node address to find the channel number

94

Send the message to new source node
on the selected channel!

96

Fig. 4d

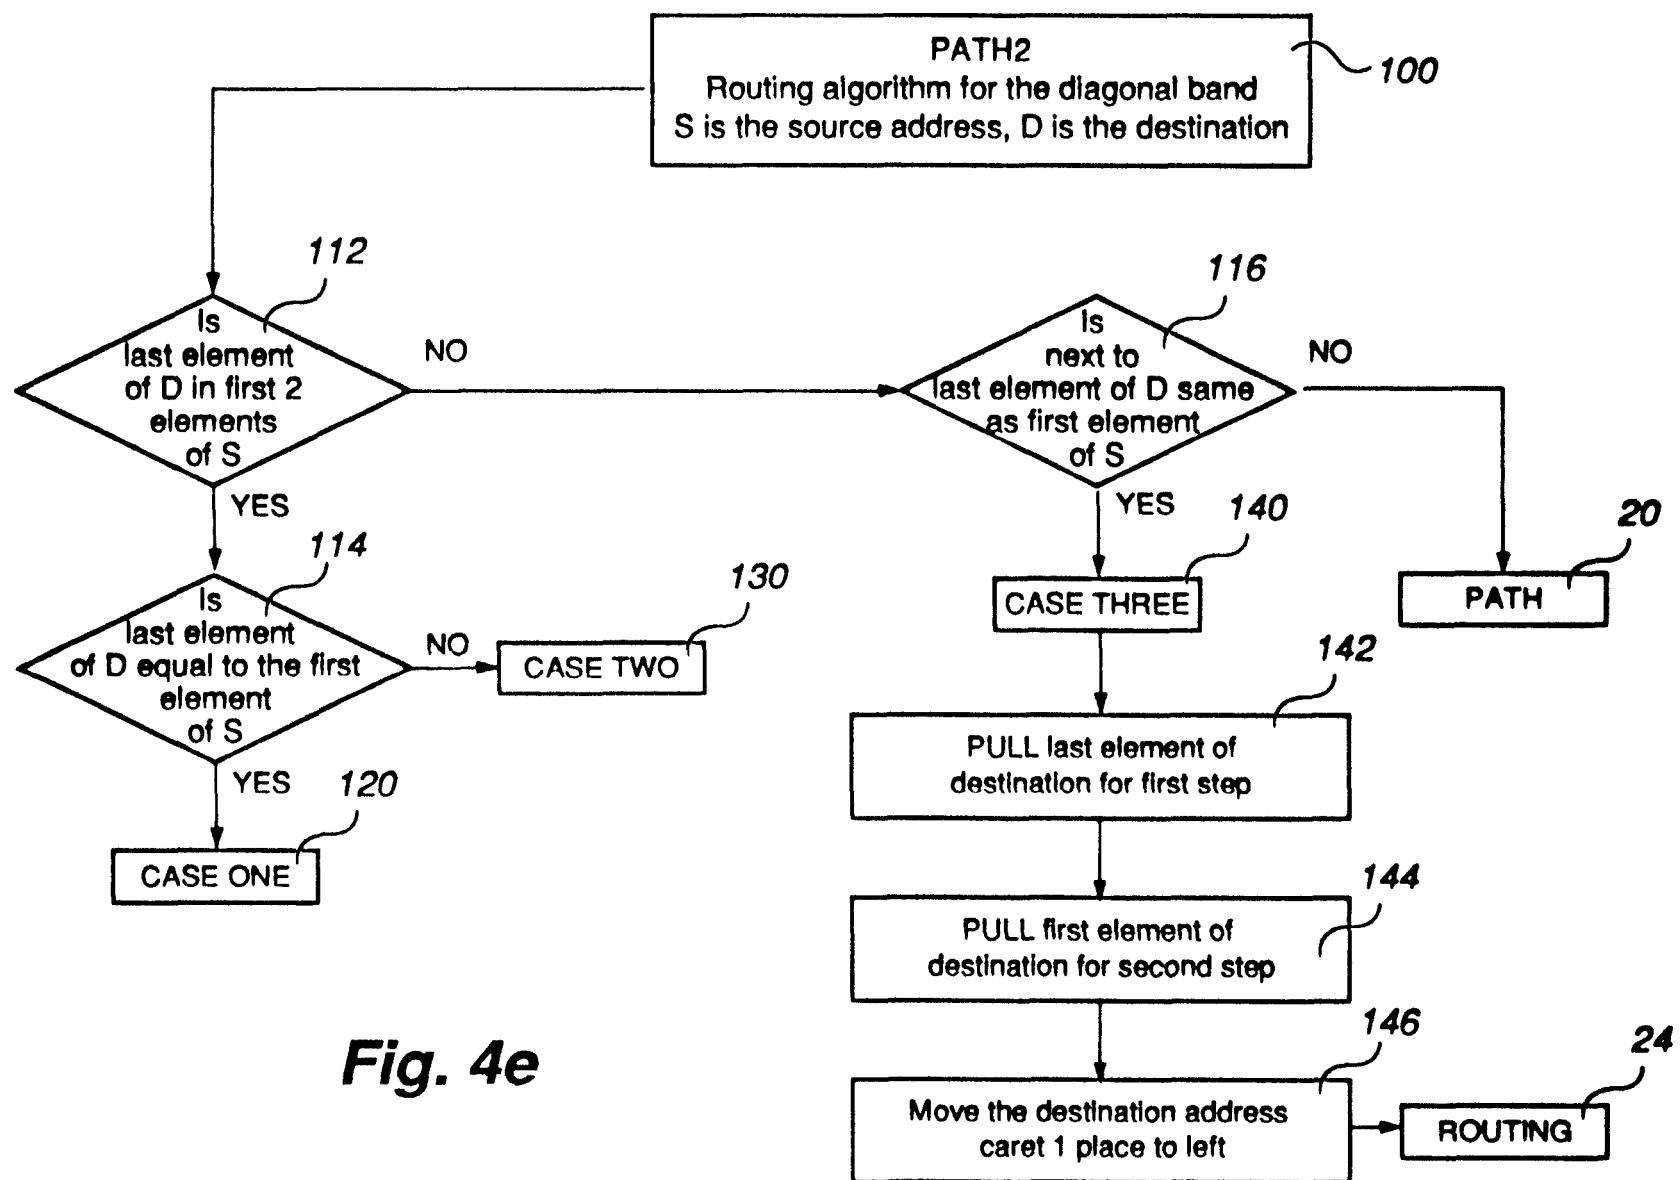


Fig. 4e

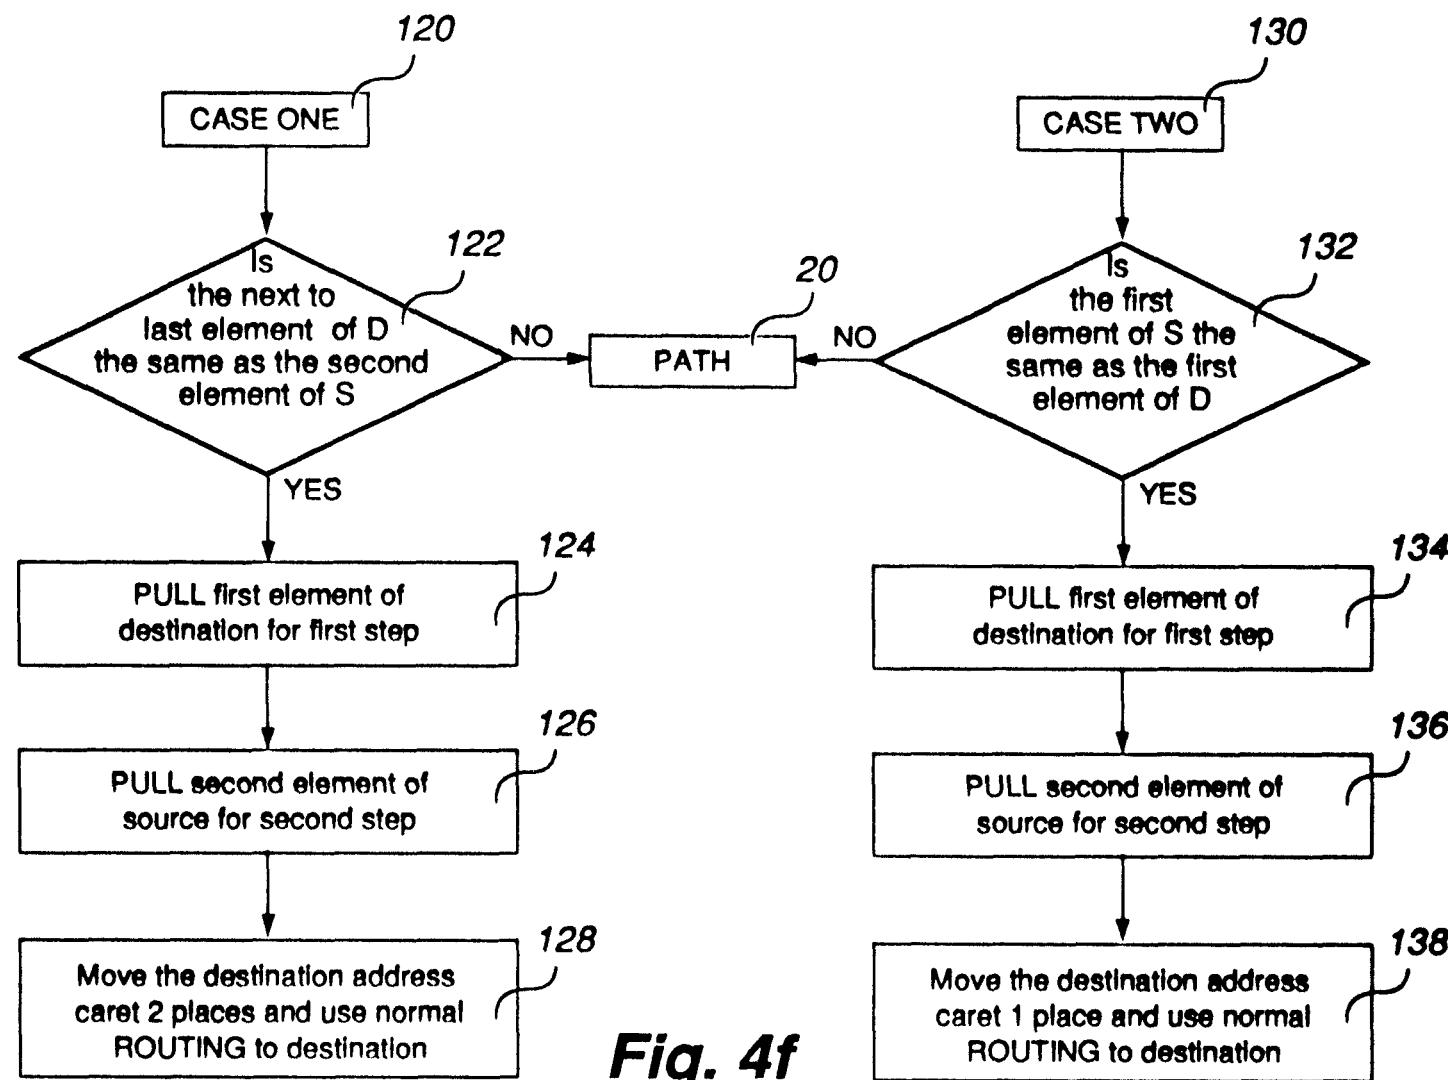


Fig. 4f