

Address-Independent Routing for Local Networks

A routing methodology is introduced which permits messages to be propagated throughout a network without recourse to destination or origin addresses. Two classes of networks, bidirectional trees and augmented rings, are analyzed from this point of view. An optimality property is proved for the bidirectional tree, and three types of address-independent routing strategies are derived. It is shown that augmented loops, a class of structures incorporating redundant links, may be rerouted to compensate for the failure of any single node or link.

1. Introduction

A local network (LN) is a means for interconnecting a set of proximate digital devices. There are numerous examples of such networks, often typified by an overall simplicity in interconnection and protocol structure. Many LN's are based on ring, bus, or star topologies. But are these the only alternatives to full circuit or packet switches? The topic of this paper is an approach to the characterization of alternative structures from a graph-theoretic point of view.

Interest in interconnection structures is increasing due to new technology and application areas. It is likely, for example, that central processors, secondary storage modules, and various other devices will communicate via what might be termed input/output local networks utilizing transmission lines operating at high speeds [1]. A desirable feature in such systems is that messages or packets be subject to a minimum amount of switching or forwarding delays. It is not clear that ring or bus structures at one extreme, or full switching at the other, are always the optimal alternatives.

This paper develops an approach to network characterization based on a notion of limited addressing. More precisely, the structures are assumed to be such that the route taken by a message from a given intermediate node is either independent of any address information contained in the message, or is perhaps a function of a set of addresses considerably smaller than the number of attached devices. Attention is restricted to networks where the route traversed by a given message is independent of such factors as congestion. The

structure of such LN's may be regarded as being composed of three layers. The first consists of the set of links connecting the nodes. The second layer comprises the routing strategy which determines the links traversed by a given message. Finally, there is the transmission protocol, controlling such factors as access, error recovery, and flow control. The paper concentrates on the first two layers, namely, the connectivity and routing, of the simplest networks of this type. In these systems, the choice of paths taken by a message from a given intermediate node is strictly a function of the link on which it reaches this node: routing is address-independent. The network nodes in this case may incorporate control logic, or they may simply be connectors and amplifiers, in which case the system may be viewed as a generalized bus structure.

Loops or rings, buses, and stars [1-6] are examples of address-independent LN's. Address independence, however, also permits configurations of substantially greater generality. Examples considered here are tree structures and loop-like systems. A principal result is that the bidirectional tree, a structure with $2(N - 1)$ links, where N is the number of nodes, has the property of requiring the minimum number of links for an address-independently routed configuration where the routing strategy is not that of the ring. The routing possibilities of such trees are explored in detail, and four distinct ones derived.

Also considered are a class of ring-like systems which incorporate redundant paths. Here the routing problem

© Copyright 1983 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

encompasses modifications required to compensate for additions and/or failures of network nodes or links.

It is interesting to note that, from the point of view developed here, buses may be considered equivalent to stars. Fibernet [7] is an example of a network where this notion appears most natural; the configuration is that of a star with a central reflecting node, but the protocol used is that of a bus (Ethernet [3]). Parenthetically, it is noted that buses are examples of trees of depth one, with a routing strategy which cannot be carried over to the general case.

The discussion is limited almost exclusively to graph-theoretic aspects of the problem. This point of view limits a comparison of the merits of the various structures to such issues as the required number of links, average and maximum path length, and reconfigurability for purposes of recovery.

2. A model

Let $\{V_i\}$, $i = 0, \dots, N - 1$, denote the set of nodes in the network and $\{E_{ij}\}$ a set of directed edges, where E_{ij} represents a path from V_i to V_j which does not enter any intermediate nodes. For simplicity, it is assumed that there is at most one edge directed from any V_i to each V_j . The $\{V_i\}$ and $\{E_{ij}\}$ jointly define a directed graph or digraph [8] G which describes the network connectivity. Figure 1 illustrates this notion for a loop and a star network.

From this point of view, a bus network may be regarded as a star network with a passive central reflecting node representing the bus structure. This model captures both the possibility of contention between the attached nodes and the lack of forwarding by these nodes. Fibernet [7] is an LN configured this way; it uses the Ethernet bus protocol.

Given the graph G , a routing table may be used to describe the set of links to be traversed by a message M_{ij} whose origin is V_i and destination is V_j . Perhaps the simplest routing strategy is one where forwarding decisions at each node V_k , $k \neq i$, are independent of i and j . This means that a node V_k , upon receiving a message M_{ij} along a link E_{rk} , need not determine the destination or origin addresses before forwarding M_{ij} along a set of links $\{E_{kq}\}$. That is, messages incoming along a link E_{rk} are forwarded along a set of links $\{E_{kq}\}$. If there is no buffering at the nodes, an implication is in general that at most one node can communicate at any given time, since otherwise messages could interfere. An exception to this rule occurs if the network is disjoint, that is, if it can be viewed as two or more separate networks.

The following is a set of requirements for the above type of routing, which will be termed *address-independent*:




Figure 1 Examples of network graphs: (a) a loop or ring; (b) a star or bus.

- i. A message originating at a node V_i is transmitted along a set of links or edges $\{E_{ij}\}^0$.
- ii. A message entering a node V_j on link E_{ij} is forwarded along a set of edges $\{E_{jk}\}$. If the set of edges $\{E_{jk}\}_i = \emptyset$, then V_j will be termed a *stop node* for edge E_{ij} .
- iii. There is no forwarding along the original edges $\{E_{ij}\}^0$. That is, the members of this set do not appear in $\{E_{ij}\}_k$ for any k given that node V_i is communicating (i.e., originating a message).
- iv. Propagation of all messages stops after they traverse a finite number of links.

For most networks it is also desirable to have

- v. The network be strongly connected: a message from any node V_i reaches all other nodes.

Another restriction, one which operates in practice, is that at most one version of a message is forwarded along a given link:

- vi. Let $\{V_k\}_i$ be the set of nodes traversed by a message originating at node V_i . Then an edge E_{pq} appears in at most one set $\{E_{pq}\}_k$ for $V_k \in \{V_k\}_i$, $\forall i$.

Restriction (vi) avoids the situation where a decision must be made concerning which version of a message to forward if a message enters V_i along more than a single path.

Another property of interest is that associated with the possibility of acknowledgment. For example, if a message is forwarded from V_i to V_k , it may be desirable to have

propagation in the return direction. This property is discussed in more detail below in connection with tree topologies.

Address-independent routing can be viewed as a simple form of broadcast mechanism. A given message M_{ij} reaches the same set of nodes as M_{ik} , $k \neq i$. Broadcasting may also be done in an address-dependent manner: stop nodes for M_{ij} , for example, may be made a function of i . This permits routing strategies more general than those possible under address independence, as discussed later in connection with tree structures. Address dependence may also be used to avoid broadcasting; here use of the destination address allows a message to be restricted to a subset of the network, thus permitting more than a single node to transmit at a given time. A simple example of a system with address dependence is a ring where the destination node V_j rather than the transmitting node V_i removes messages M_{ij} from the network.

The degree of address dependence worth incorporating in an LN hinges on several factors, including the technology to be used, the number of nodes attached, and the speed of the transmission medium. Fiber optics, for example, permits the implementation of various types of reflecting nodes which may be viewed as yielding address-independent forwarding. The number of attached devices, the traffic, and transmission bandwidths affect the delays that can be tolerated at intermediate nodes, as well as the requirement for multiple or nonblocking paths.

3. Matrix representation of address-independent routing

Address-independent routing schemes and their properties may be framed in terms of *routing matrices* whose entries denote the forwarding to be done for each edge or link in the network graph. Formulation of such matrices in a convenient form requires, however, a renumbering of the links. The following is a convention for this purpose. Suppose the M links $\{E_{jk}\}$ are represented by a set $\{F_i\}$, $i = 0, 1, \dots, M - 1$, where, for example, $F_0, F_1, \dots, F_{i(1)}$ correspond, respectively, to members of $\{E_{0j}\}$, $j = 0, 1, \dots, N - 1$; $F_{i(1)+1}, F_{i(1)+2}, \dots, F_{i(1)+i(2)}$ correspond to the set $\{E_{1j}\}$, and in general $j \geq k$ if $F_j \in \{E_{mv}\}$ and $F_k \in \{E_{vq}\}$ with $v \geq m$. Let $\{F_i\}_k$ denote the set $\{E_{kj}\}$, $j = 0, 1, \dots, N - 1$. A forwarding or transition matrix $Q = \{q_{ij}\}$ may then be defined where $q_{ij} = 1$ if

1. F_i is an edge entering a node V_k from which F_j emerges, and
2. A message received along F_i is forwarded along F_j when V_k is not the transmitting node, with $q_{ij} = 0$ otherwise.

Let $H_i = \{h_j(i)\}$ be the initial transmission vector for node V_i , where $h_j(i) = 1$ if F_j is an edge leaving node V_i that is

used to transmit a message initiated by V_i and $h_j(i) = 0$ otherwise. The transmission matrix Q may be modified in accordance with restriction (iii) in order to obtain the routes when V_i is transmitting. Let $Q(i)$ be the resulting matrix. Then $q_{jk}(i) = q_{jk}$ if $h_k(i) = 0$, and $q_{jk}(i) = 0$ if $h_k(i) \neq 0$ due to condition (iii) of Section 2.

H_i denotes the first set of edges traversed during transmission from the originating node V_i , that is, the branches from depth 0 in the path tree from node V_i . $H_i Q^{n-1}(i)$ denotes the branches from depth $n - 1$ in the tree. Let

$$V_j \cap [H_i + H_i Q(i) + \dots + H_i Q^{n-1}(i)] \triangleq R(n), \quad (1)$$

where $R = 1$ if a message from V_i reaches V_j along a path requiring n edges, and $R = 0$ otherwise. Restrictions (iv) and (v) then correspond, respectively, to

$$H_i Q_i^M = 0, \forall i \quad (2)$$

and

$$|R(m - 1)| = N, \quad (3)$$

that is, all nodes are reached by paths requiring no more than $M - 1$ edges.

As an example, consider the ring of Fig. 1, where $F_0 = E_{01}$, $F_1 = E_{12}$, and $F_2 = E_{20}$:

$$Q = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}. \quad (4)$$

Suppose node V_1 emits a message. Then

$$H(1) = [1 \ 0 \ 0], \quad (5)$$

$$Q(1) = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad (6)$$

satisfying Eqs. (2) and (3).

4. Minimum number of links for non-loop routing

Note that the minimum number of directed links required for a network connecting N nodes is N . Here the only routing option is that of the ring, that is, a route which corresponds to a cycle in the directed graph representing the network. Rings may also be configured so that a given node is visited more than once, as shown in Fig. 2. Here the number of links $M > N$.

For some applications, it may be desirable, for reasons beyond the scope of this paper, to avoid ring routing. What then are the alternatives if one wishes to preserve address independence? It is shown below that the minimum number

of links required to connect N nodes in this case is $M \geq 2(N - 1)$.

Proposition 1

Suppose the number of directed links in an LN is $M < 2(N - 1)$, where M is the number of links and N the number of nodes. Suppose further that the network is strongly connected and routing is address independent. Then the routing is such that each message M_{ij} traverses a cycle in G containing all nodes.

Proof

$M \leq 2N - 3$, so that there are at least three nodes $\{V_i^*\}$ which have a single entering link.

It is first shown that either all $V_j \in \{V_i^*\}$ are on a routing loop, or at least one such node, V_a^* , forwards messages received on its entering link on outgoing links which are disjoint from those it uses to transmit messages M_{aj} , $\forall \{V_i^*\}$. In other words, the absence of such a routing loop implies that some node V_a^* forwards no messages from $V_j \in \{V_i^*\}$ to any other members of this set. Suppose this is not true. Then a message to V_a^* from, say, V_b^* , reaches V_c^* , is forwarded to, say, V_d^* , and eventually goes around a loop, which it repeats indefinitely. Note that outgoing links from V_a^* may be partitioned into two classes: those used to forward messages entering V_a^* on its single link and those it uses only for messages M_{aj} , $\forall j$. The latter links will be termed *private*. If no routing loop exists which connects all $V_j \in \{V_i^*\}$, then a new network G^1 may be obtained as follows: V_a^* is eliminated along with its private links. Other links E_{aj} are moved to V_k , the node from which the single link entering V_a^* emanates. The routing tables for G^1 are then such that the sequence of nodes visited by any message M_{ij} , $i \neq a$, is the same as in G , except that V_a^* does not appear. The new graph is fully connected and routing is address independent. Note that $M^* \leq 2(N - 1) - 3$, since at least two links were eliminated along with V_a^* .

The new graph G^1 also has at least three nodes each with a single entering link. These are either connected via a routing loop or a new graph G^2 may be obtained in the same manner. Eventually, either a loop is found or a graph with three nodes and three links is obtained. This is necessarily connected by a ring.

It is now shown that if a network is obtained which has a ring of the above form, then all remaining nodes are on this ring.

The ring contains at least three nodes with a single entering link. All messages M_{ik} entering one of these nodes are forwarded on the ring. Thus a message M_{ik} is propagated indefinitely unless V_i is also on the loop, $\forall i$.




Figure 2 A generalized ring.

Finally, it is shown that all nodes in the original network lie on a ring. Let V_a^* be a node to be added to G^L in a process reversing the above elimination procedure. The resulting network G^{L-1} still contains a ring. To see why this is so, consider links which were affected by the removal of V_a^* to obtain G^L from G^{L-1} . These include the link entering V_a^* , private links from V_a^* , and a set of links from V_a^* which were relocated. The routing table for G^L is such that messages are forwarded simultaneously on the relocated links. Thus, at most one can be part of the loop in G^L , and adding V_a^* to obtain G^{L-1} preserves the ring. This ring connects all nodes in G^{L-1} except possibly V_a^* . But there are at least two nodes in G^{L-1} besides V_a^* which have a single entering link. Thus V_a^* is also on the ring. $\square$

Section 5 discusses various properties of the bidirectional tree, a structure for which $M = 2(N - 1)$, that is, a structure of the minimum complexity for which non-ring routing is feasible.

5. Tree networks

An example of a network with $M = 2(N - 1)$ links, the minimum for non-ring address-independent routing, is a bidirectional tree or *bitree*. Figure 3 illustrates a bitree with five nodes. Stars and, as noted above, buses are from a graph-theoretic point of view examples of a bidirectional tree of depth one.

Tree structures arise naturally if there are no redundant paths; in this case the connectivity from any node V_i to other nodes in the LN is a tree of directed links. Matching each link with one in the reverse direction yields a complete network with the advantage of symmetry.




Figure 3 A bidirectional tree.

Another property of interest is that associated with clock synchronization. Suppose, for example, that each node has a clock, synchronized to some global master. If there are no redundant links, timing information is propagated along a tree. Nodes which are neighbors on this tree are in close synchronism. A bitree ensures that all transmissions are to nodes which are neighbors from the viewpoint of clocking.

The question of routing is now considered. The following proposition gives a limit to the degrees of freedom in routing for such structures.

Proposition 2

Consider an address-independently routed network with a set of nodes $\{V_i\}$. Let $\{V_i\}_L$ be the subset of nodes such that each member has a single entering and a single exiting link. Then $\{V_i\}_L$ has at most one member which is not a stop node for its entering link unless the routing strategy is that of a ring.

Proof

Consider two nodes, $V_j, V_k \in \{V_i\}_L$, and suppose that neither is a stop node for its incoming link. A message originating at V_i will reach V_k and return to V_j ; i.e., there is a cycle or loop in the routing strategy. Thus at most one node in $\{V_i\}_L$ forwards messages received on its incoming link. $\square$

In a bidirectional tree, the leaves comprise the set of nodes with a single entering and exiting link. If the routing strategy is not that of a ring, then Proposition 2 indicates that at most one of these is not a stop node. For convenience, this node, if it exists, will be termed the *root*.

Consider the problem of non-ring routing from some intermediate node V_j . It is desirable that V_j follow some overall forwarding pattern. A message received along a link E_{ij} can be forwarded on, for example,

- All links E_{jr} , $r = 1, 2, \dots$
- All links E_{jr} except E_{ji} .
- The link E_{jr} leading to the root.
- All links E_{jr} except that leading to the root.

The routing may vary depending on whether V_i is above (rootwards) or below V_j .

Proposition 3

Let G be a bitree with a root V_i . A message reaching a node V_j along a downward link is not transmitted upward from V_j if routing is not that of the ring and is address independent.

Proof

V_j is the root of a subtree G^1 , which may be viewed as a single supernode with its own internal routing. It follows from Proposition 2 that this supernode is a stop for its single entering link. $\square$

Propositions 2 and 3 determine which of the above possibilities (a) through (d) may be combined into address-independent routes. If V_i is rootward (above) V_j , then only (b) and (d) are possible; they are also equivalent. If V_i is below V_j , then (a), (b), and (c) are possible. This yields a total of four address-independent routing methods, including that of the ring.

Routing of the form (a), corresponding to an address-independent star network, cannot be used for trees of depth greater than one because of unlimited message propagation.

Four routing possibilities are as follows:

Type 1: (loop or ring routing). A type 1 routing matrix for the tree shown in Fig. 3 (there are several possibilities, depending on the order of tree branches) is

$$Q_{\text{type 1}} = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad (7)$$

Type 2: Here a message from a node V_i is propagated rootward; i.e., routing (c) is used. After reaching the root, routing (d) is used; that is, the message is propagated on all downward links. For a network proposal using this form of routing, see Reference [5]. The type 2 routing matrix for the

network shown in Fig. 3 is given by

$$Q_{\text{type 2}} = \begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}. \quad (8)$$

A network design incorporating type 2 routing is described in Reference [6].

Type 3: Here the forwarding at each node is (b): a message touches each node just once. The routing matrix for the example of Fig. 3 is given by

$$Q_{\text{type 3}} = \begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}. \quad (9)$$

Type 4: Here a message traversing an upward link is forwarded using strategy (a). A message traversing a downward link is forwarded using strategy (d). Note that a given node may receive a particular message more than once on a given link. The type 4 routing matrix is given by

$$Q_{\text{type 4}} = \begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \end{bmatrix}. \quad (10)$$

It is interesting to consider which of the routing possibilities is compatible with a network in which the nodes have no

internal buffering. Note that type 4 routing cannot be used in this case, since a given message can traverse a given link more than once, a violation of property (vi) of Section 2. Types 2 and 3 routing are compatible with bus protocols; for systems using such protocols, these may be regarded as a means of reducing the maximum distance required for message travel.

Two parameters of interest in a network are the longest path and the maximum number of links traversed by a given message. The latter parameter is related to the probability of transmission error. The occurrence of such errors in systems without forward error correction generally requires retransmission even in the case where the destination node receives the message correctly. The longest path determines the propagation delays and is especially important for network reliability due to such factors as timing jitter propagation.

The longest paths and maximum number of links traveled are given below for the three routing options. D represents the maximum depth of the tree and N the number of nodes. It is interesting to note that type 3 routing minimizes both parameters.

Type 1: Longest path: $2(N - 1)$. Total travel: $2(N - 1)$,
 Type 2: Longest path: $2D$. Total travel: $D + (N - 1)$,
 Type 3: Longest path: $\leq 2D$. Total travel: $(N - 1)$,
 Type 4: Longest path $\leq 2D$. Total travel $\leq D(N + 1)$.

Additional routing possibilities exist if address dependence is permitted. For example, dependence on the origin address permits option (a), the bus forwarding strategy, to be used on the first visit of a message to an intermediate node. Destination dependence permits the restriction of a message to a subset of the tree containing the path between origin and destination. A simple version of the latter is to have the address of, say, V_j as viewed from V_i composed of a combination of tree level and identifier, with the tree level describing how far the message must be propagated toward the root.

6. Loop-like systems

Consider the graph shown in Fig. 4. This is a ring with additional links which permit bypassing a failing node or link. The original loop links (e.g., link 0 from node 0 to node 1) will be termed the *primary* links and the others (e.g., link 1 from node 0 to node 2) will be termed *skip* links. More formally, there are N nodes V_i , $i = 0, 1, \dots, N - 1$. (Note that $N > 2$.) From each V_i there is a link to $V_{j(i)}$ and $V_{k(i)}$, where $j(i) = (i + 1) \bmod N$ and $k(i) = (i + 2) \bmod N$. The link to $j(i)$ is a primary link and the other a skip link. This structure will be termed an *augmented loop*.

Augmented loops are examples of systems with redundant paths. The routing problem is more general here in that, in




Figure 4 An augmented loop network.

addition to routing under normal conditions, there are questions concerning routing modification to recover from node or link failures.

Address-independent routing of the form described in Section 2 requires that any set of directed links which forms a loop (as determined by the routing matrix Q) contains either a stop node or the originating node.

Perhaps the simplest routing conforming to this condition is that where there are no stop nodes for primary links and all nodes are stops for skip links. Here skip links serve as auxiliary paths for checking the correctness of transmission. This will be termed *standard ring routing*. The forwarding matrix Q for the example of Fig. 4 is given by

$$Q_{st} = \begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (11)$$

An alternative is to have no stops for the skip links and every node a stop for the primary links.

Consider the problem of rerouting when the i th node or j th link fails. Here routing must bypass the failing element while otherwise maintaining complete connectivity. It is shown below that the augmented loop may be routed in an address-independent manner so as to bypass any single failing node or link.

Suppose a skip link fails. Here no rerouting is necessary. If the i th node fails, then the j th node, $j = (i \oplus 1)$ may forward messages from the skip link associated with the $(i \ominus 1)$ node, where $\oplus$ and $\ominus$ represent mod N addition and subtraction, respectively.

The case of failure of a primary link is somewhat more complex. For messages entering V_i on the primary link, let $P^i \rightarrow P_i$ and $P^i \rightarrow S_i$ represent their forwarding from V_i on the primary and skip links, respectively, and let $S^i \rightarrow P_i$, $S^i \rightarrow S_i$ represent such forwarding of messages entering on the skip link. A set of routing tables which serve to bypass the primary link from a node V_i for N odd and even are given below.

Bypass for N odd ($N > 2$)

Here all nodes may be configured as stop nodes for primary links, with no nodes being stops for skip links: $S^j \rightarrow S_j$, $P^j \rightarrow \emptyset$, $\forall j$. The matrix Q for the example of Fig. 4 is given by

$$Q = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \end{bmatrix}. \quad (12)$$

Bypass for N even ($N > 2$)

1. $S^j \rightarrow P_j$, $j = i \oplus 1, i \ominus 2$.
2. $P^j \rightarrow P_j$, $j = i \oplus 2, i \oplus 3, \dots, i \ominus 3$.
3. $P^j \rightarrow S_j$, $j = 9 \ominus 1, i \ominus 2$.
4. $S^j \rightarrow S_j$, $j = 1, i \oplus 2, i \oplus 4, \dots, i \ominus 4$.

Figure 5 shows an example with $N = 4$ and $i = 0$.

Another question of interest is how to add or delete nodes while maintaining full communication. Suppose it is desired to add a node between V_i and V_{i+1} . This may be done by the following sequence of steps:

1. Reroute as though the primary link from V_i has failed. Disconnect this link and connect the primary link to and from the new node.
2. Reroute to the standard loop arrangement with the protocol adjusted to failure of the skip links. Reconnect the skip links to incorporate the new node.
3. Return to standard routing.

Deletion of a node may be done in an analogous fashion.

7. Conclusion

Conditions were given for address-independent routing in networks where routes are not a function of the system state.

It was shown that a bidirectional tree yields the minimum number of links if a subset of these links do not form a routing loop or ring. Four general routing strategies were derived for the bidirectional tree, and their properties were briefly discussed. Issues associated with systems which have redundant paths were considered via the analysis of augmented loop configurations. It was shown that, by suitable rearrangement of routes, it is possible to recover from the failure of any single node or link while preserving addressing independence in the forwarding mechanism.

References

1. K. J. Thurber, "Architecture and Strategies for Local Networks," in *Advances in Computers*, Vol. 20, M. C. Youits, Ed., Academic Press, Inc., New York, 1982, pp. 83-114.
2. D. D. Clark, K. T. Pogram, and D. P. Reed, "An Introduction to Local Area Networks," *Proc. IEEE* **66**, 11, 1497-1517 (November 1978).
3. J. A. M. Metcalf and D. R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," *Commun. ACM* **19**, 7, 395-404 (July 1976).
4. N. Abramson, "The Aloha System," *Technical Report No. B72-1*, University of Hawaii, Honolulu, HI, January 1972.
5. D. C. Loomis, "Ring Communication Protocols," *Technical Report 20*, Department of Information and Computer Science, University of California, Irvine, CA.
6. F. P. Closs and R. P. Lee, "A Multistar Broadcast Network for Local Area Communication," in *Local Networks for Computer Communications*, A. West and P. Janson, Eds., North-Holland Publishing Co., New York, 1981, pp. 61-80.
7. C. W. Marshall, *Applied Graph Theory*, Wiley-Interscience, New York, 1971.
8. E. G. Rawson and R. M. Metcalfe, "Fibernet: Multimode Fibers for Local Computer Networks," *IEEE Trans. Commun.* **COM-26**, 7, 983-990 (July 1978).

Received April 5, 1983; revised May 11, 1983




Figure 5 Bypass routing for N even.

Peter A. Franaszek IBM Research Division, P.O. Box 218, Yorktown Heights, New York 10598. Dr. Franaszek is a member of the Computer Sciences Department at the Thomas J. Watson Research Center. His interests include analytical problems associated with storage hierarchies, computer organization, magnetic recording, and digital communications. He received his B.Sc. degree from Brown University, Providence, Rhode Island, in 1962, and the M.A. and Ph.D. degrees from Princeton University in 1964 and 1965. During the academic year 1973-1974, he was on sabbatical leave at Stanford University as a Consulting Associate Professor of Electrical Engineering and Computer Science. Prior to joining IBM in 1968, he was a member of the technical staff at Bell Telephone Laboratories. Dr. Franaszek is a member of the American Association for the Advancement of Science, Sigma Xi, and Tau Beta Pi.