IEEE TRANSACTIONS ON COMMUNICATIONS

MAY 1987 VOLUME COM-35 NUMBER 5 (ISSN 0090-6778)

A PUBLICATION OF THE IEEE COMMUNICATIONS SOCIETY

PAPERS

CORRESPONDENCE

IEEE COMMUNICATIONS SOCIETY

The field of interest of the IEEE Communications Society consists of all telecommunications including telephone, telegraphy, facsimile, and point-to-point television, by electromagnetic propagation including radio; wire; aerial, underground, coaxial, and submarine cables; waveguides, communication satellites, and lasers; in marine, aeronautical, space and fixed station services; repeaters, radio relaying, signal storage, and regeneration; telecommunication error detection and correction; multiplexing and carrier techniques; communication switching systems; data communications; and communication theory.

In addition to the above, this TRANSACTIONS contains papers pertaining to analog and digital signal processing and modulation, audio and video encoding techniques, the theory and design of transmitters, receivers, and repeaters for communications via optical and sonic media, the design and analysis of computer communication systems, and the development of communication software. Contributions of theory enhancing the understanding of communication systems and techniques are included, as are discussions of the social implications of the development of this communication technology. All members of the IEEE are eligible for membership in the Society upon payment of the annual Society membership fee of $12.00. Members may receive this TRANSACTIONS or the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS upon payment of an additional $8.00 ($20.00 total), or both publications upon payment of an additional $16.00 ($28.00 total). For information on joining, write to the IEEE at the address below.

(Information for Authors, ComSoc Board of Governors, Departments, and Committees appear on Cover III.)

IEEE TRANSACTIONS ON COMMUNICATIONS is published monthly by The Institute of Electrical and Electronics Engineers, Inc. Responsibility for the contents rests upon the authors and not upon the IEEE, the Society Council, or its members. **Headquarters**: 345 East 47th Street, New York, NY 10017-2394 U.S.A. **IEEE Service Center** (for orders, subscriptions, address changes, Region/Section/Student Services): 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331. **Telephones**: Headquarters 212-705 + extension: Information-7900, General Manager-7910, Controller-7748, Educational Services-7860, Publishing Services-7560, Standards-7960, Technical Services-7890. IEEE Service Center 201-981-0060. Professional Services: Washington Office 202-785-0017. **NY Telecopier**: 212-752-4929. **NY Telex**: 236-411 (international messages only). **IEEE Service Center** (for orders, subscriptions, address changes, Educational Activities, Region/Section/Student Services): 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331. **NJ Telephone**: 201-981-0060. **IEEE Washington Office** (for U.S. professional activities): 1111 19th Street, NW, Suite 608, Washington, DC 20036. **Washington Telephone**: 202-785-0017. **Price/Publication Information:** Individual copies: IEEE members $10.00 (first copy only), nonmembers $20.00 per copy. Available in microfiche and microfilm. Printed in U.S.A. Second-class postage paid at New York, NY and at additional mailing offices. **Postmaster:** Send address changes to IEEE TRANSACTIONS ON COMMUNICATIONS, IEEE, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331.

PAPERS TO BE PUBLISHED IN THE NEXT ISSUE
June 1987

Papers

Correspondence

Routing in the Manhattan Street Network

NICHOLAS F. MAXEMCHUK, SENIOR MEMBER, IEEE

Abstract—**The Manhattan Street Network is a regular, two-connected network, designed for packet communications in a local or metropolitan area. It operates as a slotted system, similar to conventional loop networks. Unlike loop networks, routing decisions must be made at every node in this network. In this paper, several distributed routing rules are investigated that take advantage of the regular structure of the network.**

In an operational network, irregularities occur in the structure because of the addressing mechanisms, adding single nodes, and failures. A fractional addressing scheme is described that makes it possible to add new rows or columns to the network without changing the addresses of existing nodes. A technique is described for adding one node at a time to the network, while changing only two existing links. Finally, two procedures are described that allow the network to adapt to node or link failures. The effect that irregularities have on routing mechanisms designed for a regular structure is investigated.

I. Introduction

THE Manhattan Street Network (MSN) [1], Fig. 1, is a two-connected, regular network with unidirectional links. The links are arranged in a structure that resembles the streets and avenues in Manhattan. The MSN topology is being applied to a local or metropolitan area packet communication system.

The nodes in the MSN are described in Section II. The structure of the nodes and the access strategy are similar to those in a slotted loop system. The principle difference between the MSN and a loop network is that there are two links arriving at and leaving each node instead of a single link, and a routing decision must be made for each packet transmitted at each node. An experimental network is being constructed with a 50 Mbit/s transmission rate on each link and 128 bit fixed sized packets. More than 750 000 routing decisions per second may have to be made at each node in this network. In this type of network, the routing rule must be simple.

In this paper, distributed routing rules for the MSN are investigated. Simple routing rules that use the regular structure of the network are compared to shortest path algorithms and random routing strategies. In the MSN, shown in Fig. 1, the number of rows and columns completely defines the network, and if these numbers are known, the shortest path between any pair of nodes can be determined. In addition, because of the cyclic structure of the MSN, routing is only dependent upon the relative location of the current node with respect to the destination, as defined in Section III-B, and the same routing rule can be used at every node. In Section IV-A, a distributed rule is described that finds the shortest path. In Sections IV-B and IV-C, two simplifications of the shortest path rule are described. The simplified rules do not always find the shortest path, and the effect that these rules have on the average path length is investigated in Section IV-E.

In Section III-A, a fractional addressing scheme is described. This addressing scheme has two advantages over the integer addressing scheme in Fig. 1.

Paper approved by the Editor for Wide Area Networks of the IEEE Communications Society. Manuscript received November 25, 1985; revised December 11, 1986.

The author is with AT&T Bell Laboratories, Murray Hill, NJ 07974.

IEEE Log Number 8613925.

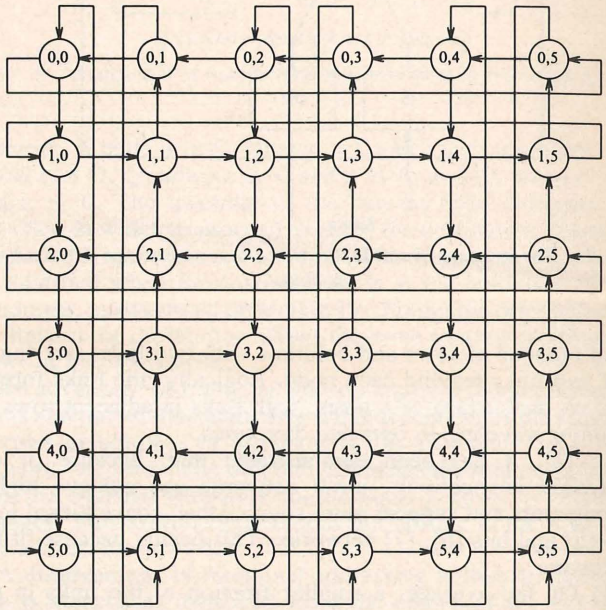

Fig. 1. 36-node MSN.

1) New rows or columns are added to the network without changing the addresses of existing nodes.

2) The distributed routing rules are independent of the number of rows or columns in the network.

A disadvantage of fractional addressing is that the distributed routing rules must operate without knowing the position of all of the rows and columns in the network and cannot always find the shortest path to the destination. The effect that this addressing scheme has on the average distance between nodes in investigated in Section IV-E.

The MSN is a regular structure with an even number of rows and columns and is not defined for an arbitrary number of nodes. A realistic network, in which nodes are added and other nodes or links fail, may be approximated by the MSN, but it is unlikely that it will exactly correspond to the regular structure. The routing rule that is selected for the MSN must operate in networks with irregularities. The effect that irregularities have on the routing rules depends upon the techniques used to add nodes and remove failed components. In Section V-A, a technique is described for adding one node at a time to the network. When this technique is used, only two links must be changed when a new node is added to the network. As nodes are added, a row or column may not have a full complement of nodes. The routing rules operate without knowing which rows or columns are incomplete. In Sections V-B and V-C, procedures are described that allow the network to adapt to node or link failures. The adaptations guarantee that the nodes continue to operate without losing packets at any of the surviving nodes. The routing rules operate without knowing which nodes or links have failed.

II. System Description

The MSN, Fig. 1, is a member of a class of multiply-connected, regular, mesh-configured networks. There is an

0090-6778/87/0500-0503$01.00

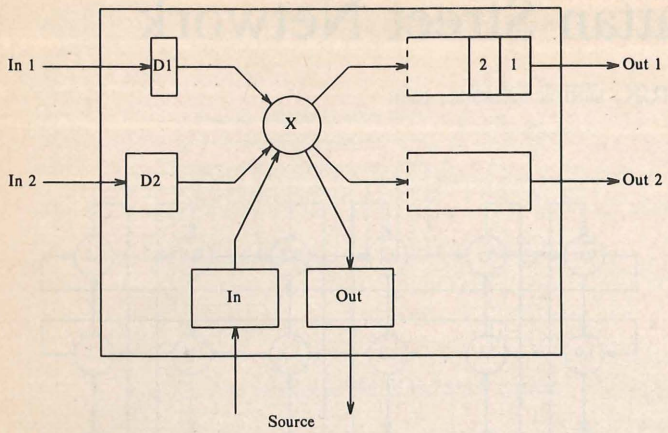

Fig. 2. The structure of a node in a two-connected network with fixed size packets.

even number of rows and columns with two links arriving at and two links leaving each node. Logically, the links form a grid on the surface of a torus, with links in adjacent rows or columns traveling in opposite directions.

In [2], it has been demonstrated that, because of the increased connectivity, mesh networks can achieve higher throughputs and support more sources than conventional loop [3]–[5] and bus [6], [7] networks. This occurs because of the following.

1) On the average, a smaller fraction of the links in the network are used to interconnect a source and destination.

2) Sources that communicate frequently can be clustered into communities of interest that do not interfere with one another.

In a network with several paths arriving at a node, messages from more than one incoming link may be destined for the same outgoing link. Data from several links can be concentrated onto one link by storing the data, forwarding them when the link is available, and establishing protocols to recover messages that are lost because of buffer overflows. In [2], a slotted system is described that does not require buffering on the output links and does not lose packets because of buffer overflows. The structure of a node in this network is shown in Fig. 2.

The packets in the slotted system are a fixed size. A node periodically transmits a packet from an input line, a packet from the source, or an empty packet on each output line. At each node, the packets from the input lines are delayed so that they arrive at the switch at the time that the node transmits a packet. The node switches each of the incoming packets not destined for the node to one of the output links. If the buffer for an output link is full, and two incoming packets are destined for this link, one of the packets is forced to take the other output link. This strategy guarantees that packets are never lost because of buffer overflow, even if the output buffer size at the nodes is reduced to zero; however, the larger the buffers at a node, the less likely it is that a packet must be misdirected. Packets from the source are only transmitted when there is an empty slot on an output link. The node controls the source so that packets do not arrive faster than they are transmitted, and the rate available to the source decreases when the network is busy. Packets that are misdirected take a longer path to their destination and prevent more new packets from entering the network. Therefore, there is a tradeoff between buffer size and the throughput of the network.

In the MSN, each time a packet is misdirected, the length of the path to the destination is increased by at most four links. In addition, there are many nodes for which either outgoing link provides the same path length to the destination, and when a packet may take either link, the probability any packet will have to be misdirected decreases. A recently published analysis and simulation of the MSN [8] indicates that the MSN operates reasonably efficiently without buffers on the outgoing links, and the experimental system that is being implemented does not have buffers.

III. Addressing Nodes in the Network

Each node in the network has a unique address that is called the node's absolute address. To simplify the routing rule, the absolute address of a node reflects the regular structure of the MSN. Because of the cyclic nature of the MSN, routing only depends on the relative position of the current node and the destination, called the current nodes relative address, and not on the absolute address of any node. Relative addresses allow the same routing rule to be used at each node.

A. Absolute Addresses

In Fig. 1, the rows in the MSN are sequentially numbered from 0 to $m - 1$, the columns are numbered from 0 to $n - 1$, and the absolute address of a node is its row and column. The odd-numbered rows have links in one direction and the even-numbered rows have links in the opposite direction. New rows or columns are added in pairs to preserve the alternating directions, and the address of an existing row or column changes when new elements are added in the middle of the network. To reduce the effect of changing addresses on the communications routines at the source, there must be a transformation between a logical address by which the source refers to the destination and a physical address that is the destination node's current row and column. The transformation need only be performed at the source node of a packet; however, a transformation table must be maintained at every node in the network, and a protocol must be developed to update the table as physical addresses change.

An alternative to changing the address of a node when new rows and columns are added is to plan for expansion by not using all of the addresses initially. For instance, in the initial implementation of the network, the rows may be numbered 0, 11, 22, $\cdots$ so that ten new rows can be added between each of the initial rows. By leaving an even number of integers between assigned rows or columns, the alternating direction can be retained as new rows are added. The spacing between rows can be decreased when nodes in the same community of interest are in adjacent rows, and it can be increased where new communities of interest may be inserted. This approach requires careful planning because the network growth must be predicted when the initial network is designed. The planning can be reduced by using fractional addresses rather than integer addresses. Fractional addresses allow an arbitrary number of pairs of rows to be added at any position in the network.

The fractional addressing scheme that has been selected is shown in Fig. 3. The first two rows or columns are labelel 0 and 1. Rows are added in pairs and are labeled as two fractions, 1/3 of the way between two other rows. For instance, two rows added between 0 and 1 are labeled 1/3 and 2/3 and two rows added between 2/3 and 1 are labeled 7/9 and 8/9. New rows that are added between 1 and 0 are considered to be between 1 and 2 so that they have different addresses from the rows between 0 and 1. For instance, two rows added between 1 and 0 are labeled 4/3 and 5/3 and two rows between 5/3 and 0 are labeled 16/9 and 17/9. Fractional addressing does not constrain the total number of rows that can be added to the network or the number of rows that can be added to a community of interest in a particular area of the network. In addition, the fractional addressing scheme selected guarantees that all rows with an even numerator have links in one direction and all rows with an odd numerator have links in the opposite direction, as in the integer addressed system.

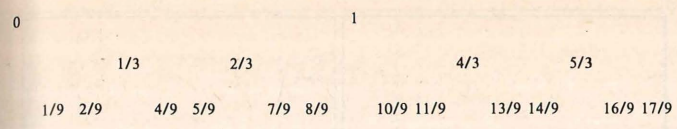

Fig. 3. Fractional addressing in the MSN.

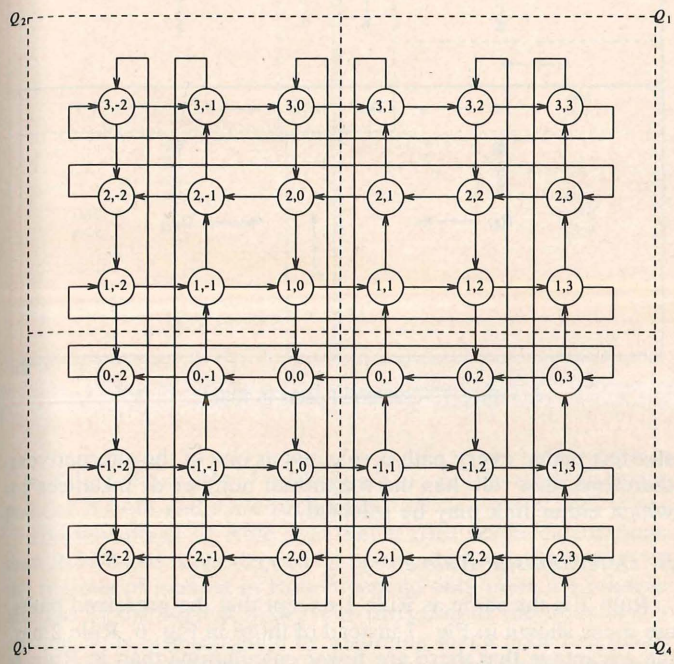

Fig. 4. Relative addresses in a 36-node MSN.

B. Relative Addresses

Because of the cyclic structure of the MSN, any node can be considered to be in the center of the network. The relative address (r, c) of a node with absolute address (r_{fr}, c_{fr}) with respect to the destination node with absolute address (r_{to}, c_{to}) is defined so that the destination node is approximately at the center of the network, has relative address (0, 0), and has both row and column links directed toward decreasing numbered nodes, as in Fig. 4.

The relative address in an $m \times n$ integer-addressed network is

$$r = \frac{m}{2} - \left\{ \left(\frac{m}{2} - D_c(r_{fr} - r_{to}) \right) \bmod m \right\}$$

$$c = \frac{n}{2} - \left\{ \left(\frac{n}{2} - D_r(c_{fr} - c_{to}) \right) \bmod n \right\}, \qquad (1)$$

and in a fractionally addressed network is

$$r = 1 - \{(1 - D_c(r_{fr} - r_{to})) \bmod 2\}$$

$$c = 1 - \{(1 - D_r(c_{fr} - c_{to})) \bmod 2\} \qquad (2)$$

where D_c and D_r are dependent upon the direction of the links at the destination node. In a network with the links in the even and odd rows and columns directed toward increasing or decreasing rows and columns, as in Fig. 1, $D_c = +1$ when c_{to} (the numerator of c_{to} in a fractionally addressed network) is even, $D_c = -1$ when c_{to} is odd, $D_r = +1$ when r_{to} is even, and $D_r = -1$ when r_{to} is odd.

The definition of the relative coordinates in (1) and (2) limits the relative address of the current node to $-(m/2) < r \leq m/2$, and $-(n/2) < c \leq n/2$ for an integer-addressed network and $-1 < r, c \leq 1$ for a fractionally addressed network. A node is in Q_1 when $r > 0$ and $c > 0$, Q_2 when $r > 0$ and $c \leq 0$, Q_3 when $r \leq 0$ and $c \leq 0$, and Q_4 when $r > 0$ and $c \leq 0$. The quadrant of the current node indicates the direction in which to proceed to get to the destination. Because the network has unidirectional links, this routing strategy must be modified when the current node is at the boundary of the quadrants, as discussed in Section IV. Fixing the orientation of the links at the destination allows the same routing decisions to apply at the boundaries.

Quadrant 3 or 4			Quadrant 1 or 2				
$-\frac{2}{3}$	$-\frac{1}{3}$	0	$\frac{1}{9}$	$\frac{2}{9}$	$\frac{1}{3}$	$\frac{2}{3}$	1

a) Actual assignment of rows to quadrants

Quadrant 3 or 4			Quadrant 1 or 2			
$-\frac{2}{3}$	$-\frac{1}{3}$	0	$\frac{1}{9}$	$\frac{2}{9}$	$\frac{1}{3}$	$\frac{2}{3}$

b) Expected assignment of rows to quadrants

Fig. 5. Assignment of rows to quadrants in a network with eight rows.

An advantage of fractional addressing over integer addressing is that the relative addresses are independent of the number of rows or columns in the network. In an integer-addressed network, the arithmetic unit that calculates the relative address must be changed whenever the number of rows or columns changes. This arithmetic unit does not change in the fractionally addressed network.

A disadvantage of fractional addressing is that the destination is sometimes displaced from the center of the network when relative addresses are calculated. For instance, in Fig. 5(a), a possible assignment of rows to quadrants is shown for a network with eight rows and only two of a possible 12 rows in the 1/9th addressing level. Five rows are assigned to quadrants 1 or 2 and only three to quadrants 3 or 4, and as a result, packets routed from nodes with a relative address $(1, X)$ may take a longer path to the destination. In Fig. 5(b), the assignment of rows to quadrants that places the destination in the center is shown.

It is evident from this example that new rows should be added uniformly, when possible, in order to calculate the quadrant correctly. However, the quadrant is most likely to be calculated incorrectly for the nodes that are furthest apart. In large networks, with many small communities of interest, expanding the network with nonuniform addresses that keep nodes in their communities of interest is preferable to forcing nodes to join distant parts of the network. If most packets remain within the community of interest and are not directed to the nodes that are furthest away, the distance between nodes that communicate frequently is kept small and the effect of nonuniform addresses is less than in a network with uniform traffic requirements.

IV. Distributed Routing Rules

In Sections IV-A, B, and C, three distributed routing rules are described that use the regular structure of the MSN to select a path to the destination. Rule 1 determines all shortest paths to the destination for integer addressed MSN's. Rules 2 and 3 reduce the number of calculations that are performed at each node, but occasionally take longer paths. Rules 1 and 2 are dependent upon the addresses of the adjacent nodes to which a node is connnected; rule 3 is not.

In complete, integer-addressed networks, the address of adjacent nodes is known. In fractionally addressed networks or in networks with partially full rows or columns, as described in Section V-A, the address of adjacent nodes is not known. If rule 1 or 2 is used, a technique must be used to determine the node to which each node is connected. In the experimental network, the nodes to which a node is connected is stored locally and changed manually when the network connectivity

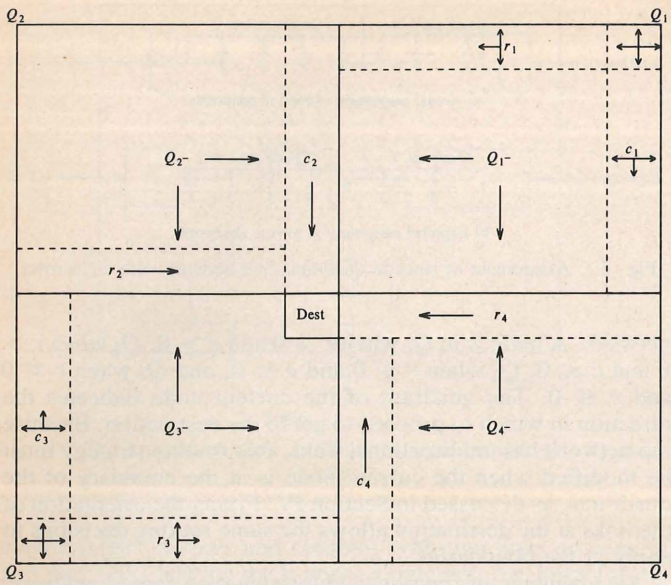

Fig. 6. Preferred paths in Rule 1.

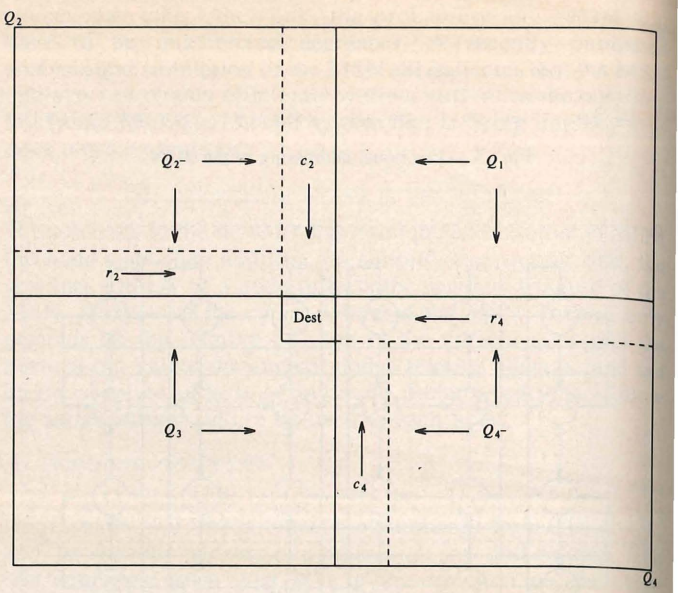

Fig. 7. Preferred paths in Rule 2.

changes. When the procedure described in Section V-A is used to add nodes, the connectivity for only two nodes changes when a new node is added to the network. Therefore, a manual rather than an automatic procedure is reasonable. When nodes or links fail and are bypassed automatically, as in Sections V-B and C, the connectivity information at a node is not changed, and it is incorrect.

These three rules are referred to as deterministic routing rules. In addition, two random routing rules are described in Section IV-D. Rule A is independent of the address of adjacent nodes, and rule B routes packets correctly when the destination is one node away. In Section IV-E, the path lengths resulting from the routing rules are compared in integer and fractionally addressed networks.

A. Deterministic Rule 1

The solid arrows in Fig. 6 show the preferred direction of travel from the relative positions in the network to the destination for the first routing rule. In this figure, r_1, r_2, r_3, and r_4 are rows at the edges of the quadrants, and c_1, c_2, c_3, and c_4 are columns at the edges of the quadrants. The first routing rule is as follows.

Rule 1:

- Select the preferred path if there is one preferred path from a node.
- Select either path if there are zero or two preferred paths from a node.

To implement the first rule, the relative addresses of the current node (r, c), the next node along the column (r_{nxt}, c), and the next node along the row (r, c_{nxt}) are calculated. The quadrant is determined from (r, c) as in Section III-B, the direction of the link along the row is determined from $c - c_{nxt}$, and the direction of the link along the column from $r - r_{nxt}$. A node in Q_4 is in row r_4 if $r = 0$, and a node in Q_2 is in column c_2 when $c = 0$. When a node in Q_2 is also in r_2, the link directed down is not preferred. These links can be determined because $r_{nxt} = 0$ and $c_{nxt} \neq 0$. Similarly, a link is directed to the left from c_4 if (r, c) is in Q_4, $c_{nxt} = 0$, and $r_{nxt} \neq 0$. Row r_1 is at the outside edge of the network. A node in Q_1 is in r_1 and has a preferred link that is not preferred in Q_1 if $c_{nxt} \neq 0$ and c_{nxt} is in Q_2. Similarly, when a node is in r_3, c_1, or c_3 and has a preferred link that is not preferred in the rest of the quadrant, an adjacent node is also in a different quadrant.

In the Appendix, it is shown that this routing rule selects the shortest path from any node to the destination in an integer addressed network. Furthermore, when there are several shortest paths, every path is selected as one of the alternatives; therefore, this rule has the maximum number of instances in which either link may be selected.

B. Deterministic Rule 2

Rule 2 is the same as Rule 1 except that the preferred paths are those shown in Fig. 7 instead of those in Fig. 6. Rule 2 has the advantage that there are fewer calculations than in Rule 1 because the special cases when nodes are in r_1, c_1, r_3, and c_3 are not determined. However, this rule has the disadvantage that nodes in these special rows and columns take a slightly longer path to the destination. In Rule 2, it is still necessary to know c_{nxt} and r_{nxt} in order to determine when a node is in r_2 or c_4.

The routing rule is simplified in this manner because it should have a relatively small effect on the average path length. The nodes that are affected are those that are furthest from the destination. Fewer packets are affected by changes in routing rules in these nodes than elsewhere in the network because packets from other nodes are not intentionally routed through these nodes to get to the destination, and in a network with communities of interest, nodes are more likely to communicate with nodes that are nearby. By contrast, a change in the routing rule in c_2 and r_4 would affect every packet headed for the destination. In addition, incorrect paths are not selected at all of the nodes in the affected rows and columns. From Fig. 6, preferred paths in the quadrants are also preferred paths in the special rows and columns at the edges of the network. Incorrect decisions may only be made at nodes where neither path is thought to be preferred and one of the paths is shorter. Furthermore, from Table V in the Appendix, the longer paths at the edge of the network are two greater than the preferred paths, while elsewhere in the network, they are four greater. The effect of longer paths on the average shortest path length is shown in Section IV-E.

C. Deterministic Rule 3

The solid arrows in Fig. 8 show the preferred paths and the dashed arrows show the alternate paths. The routing rule is as follows.

Rule 3:

- Select the preferred path if there is one preferred path from a node.
- Select the alternate path if there is no preferred path and one alternate path from the node.

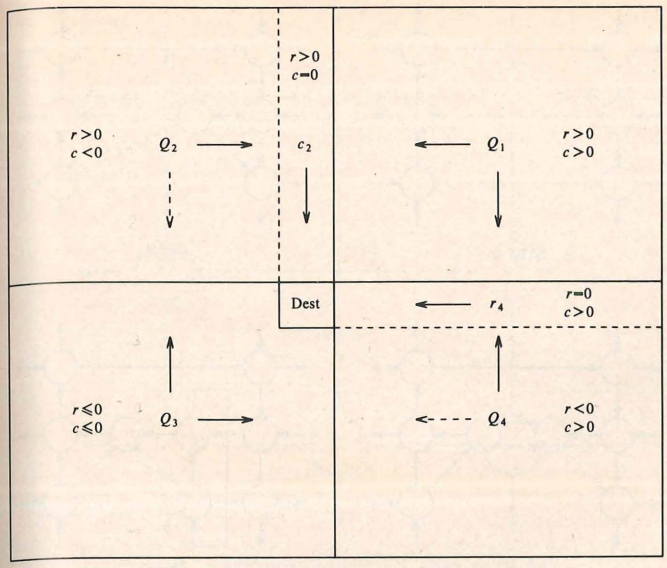

Fig. 8. Preferred and alternate paths in Rule 3.

• Select either path if neither path is a preferred or alternate path or if both paths are preferred.

The advantage of Rule 3 is that it uses fewer calculations than Rule 2 and is not dependent upon r_{nxt} or c_{nxt}. From Fig. 8, the regions of interest in Rule 3 depend only upon the relative address of the current node. The direction of the links at the current node is determined by assuming that a node has a link directed to the left when r is even and to the right when r is odd, and a link directed down when c is even and up when c is odd. The disadvantage of Rule 3 is that there are fewer instances in which either path from the node may be selected. In Q_2 and Q_4 where Rule 2 may select either path, Rule 3 is constrained to select one of the paths. This increases the number of times when two packets arriving at the node conflict. There are also instances in incomplete networks, in Section V-A, where Rule 3 cannot get to a specific destination while Rule 2 can.

In complete networks, Rules 2 and 3 result in the same distance from any node to the destination. Whenever there are one or two preferred paths in Rule 2, one of these paths is selected by Rule 3. When there are two preferred paths in Rule 3, both paths have the same distance to the destination. Therefore, the path length is the same for both rules.

D. Random Routing

Two random routing rules have been considered. Rule A is completely random, a packet selects either link with equal probability, and at each node, checks to see if it is at the destination. Rule B assumes that the two nodes to which a node is connected is known. At each node, if the destination is one node away, the packet is directed there; otherwise, a path is selected at random.

There are two advantages to using random routing rules rather than deterministic rules. First, they are extremely easy to implement. In the random routing rules, it is not necessary to calculate the relative address of a node, its quadrant, or the direction of the links emanating from the nodes. Second, random routing rules are extremely tolerant of network irregularities. If nodes are added or fail in a perverse manner, the network may bear little resemblance to the regular structure, and the deterministic rules may not work. The random rules provide an alternative to the deterministic rules when this occurs. These random routing rules were first investigated by Prosser [9] as a routing mechanism for survivable networks. The disadvantage of random routing rules is that they use more links to get between a source and destination, and this results in a smaller network throughput.

TABLE I
THE EFFICIENCY OF ROUTING RULES RELATIVE TO THE SHORTEST PATH ALGORITHM

Network	Short. Path	Efficiency of Routing Rules					
		Deterministic				Random	
		Integer Addr.		Fractional Addr.		A	B
		1	2, 3	1	2, 3		
4x4	2.93	1.00	1.00	.95	.94	.21	.79
4x6	3.30	1.00	.97	.97	.95	.14	.30
6x6	3.71	1.00	.97	1.00	.97	.10	.21
6x8	4.34	1.00	.98	.99	.97	.09	.17
8x8	5.02	1.00	1.00	1.00	.98	.07	.14
8x10	5.42	1.00	.99	1.00	.98	.06	.11
10x10	5.84	1.00	.99	1.00	.99	.05	.09
10x12	6.42	1.00	.99	1.00	.99	.05	.08
12x12	7.02	1.00	1.00	1.00	.99	.04	.07
12x14	7.45	1.00	1.00	1.00	.99	.04	.06
14x14	7.89	1.00	.99	1.00	.99	.03	.06

E. Comparison

A comparison of the deterministic routing rules in integer addressed and fractionally addressed networks is presented in Table I. The average distance between nodes for a routing rule is calculated by determining the average distance between each source and destination in the network. The efficiency of the routing rule is the average of the shortest distance between nodes over the average distance between nodes using the routing rule. In the comparisons, there is no contention, and a packet always takes the path specified by the routing rule. When the rule decides that both paths are equivalent, either path is selected with probability 0.5. Because of this random component, a packet does not always take the same length path from a source to the destination. To compensate for the random component, the efficiency is calculated by determining the average distance between each node several times and averaging the result. The number of times that the average distance is determined is varied so that the span of values representing a 95 percent confidence interval is less than 1 percent of the average value.

Table I shows that Rule 1 determines the shortest path in integer addressed networks, and that Rules 2 and 3 result in the same average distance between nodes. Rule 2 selects longer paths than Rule 1 when the relative location of a node is at the edge of the network, and this effect is also seen in the table. In the simulations, fractionally addressed rows and columns are added to the network two at a time in the order shown in Table II, which makes the depth or the row and column addresses as uniform as possible. It is evident that fractionally addressed rows can be added to large networks in a way that has a small effect on the average path length.

Random rules are inefficient. In the networks in Table I, the average path length using random routing can be 33 times longer than the path lengths resulting from the deterministic rules. It is inadvisable to use random routing when a network has some regularity to its structure. However, a hybrid random and deterministic rule can be used to obtain the efficiency of the deterministic rule in a regular network and the survivability of the random rule. For instance, a random component can be inserted in the routing rule after a packet has traversed a larger number of nodes than expected. The number of nodes a packet has traversed must be tracked in any practical network because when a node fails, packets destined for this node must be purged from the network.

V. Network Irregularities

In addition to getting packets quickly between nodes in complete, regular MSN's, the routing rules must continue to

TABLE II
THE ORDER IN WHICH ROWS AND COLUMNS ARE ADDED TO THE NETWORK

Number of Rows or Cols.	Address of Rows or Columns Added	
	0	1
4	1/3	2/3
6	4/3	5/3
8	1/9	2/9
10	10/9	11/9
12	4/9	5/9
14	13/9	14/9

function in irregular networks. The irregularities investigated in this section occur when the number of nodes that are added to the network are not sufficient to completely fill a row or column and when nodes or links fail and are deleted from the network.

The effect of the irregularities on the routing rules depends upon the procedures used to add and delete nodes and links from the network. In Section V-A, a procedure for adding one node at a time to a network is described. This procedure has the characteristic that only two existing links must be changed to add a new node. In Sections V-B and C, procedures for deleting failed nodes and failed links are described. These procedures are similar to those used in loop networks [10] and can be implemented automatically. The source is not informed when a packet cannot be delivered to the destination and not all failures are detected. Therefore, a higher level acknowledgment protocol is still required to guarantee that packets are delivered.

A. Adding Nodes One at a Time

A procedure is shown in Fig. 9 for adding one node at a time to an MSN. Each time a node is added, two links must be changed. The two links that will be changed when the next node is added are shown by dashed lines. When this procedure is followed, two complete new rows or columns are eventually added to the network.

Adding one node at a time makes the network less regular and affects the ability of the distributed routing rules to find the shortest path to a destination. The effect on a 6 × 6 network is shown in Table III. When 12 nodes are added, a 6 × 8 network is formed. In this table, the efficiency is calculated as in Section IV-E. The italicized numbers in parentheses indicate the fraction of source destination pairs that are unable to communicate.

There are several cases for which Rule 3 cannot find a path. The reason this routing rule fails is seen by examining Step 4 in Fig. 9. Assume that a packet at node A is destined for node B. Node A is an odd-numbered column in Q_4; therefore, in Rule 3, the link along the column is assumed to be directed upward and is selected. Unfortunately, the column is not complete and the packet ends up at node C. At node C, which is also in an odd-numbered column in Q_4, the upward-directed path is selected, and the packet arrives back at node A. At node A, the path to node C is again selected, and the packet is stuck in a loop. In Rules 1 and 2, at node A it is known that the next node along the column is node C and that both links are directed away from the destination. Therefore, at node C, either link is selected with probability 0.5 and the loop is avoided.

B. Node Failures

Loop systems have active components in the path at each node, and if one of these components fails, the loop is broken. When nodes fail, they are bypassed so that the remainder of the loop continues to operate. Loss of power at the node is a common failure because power is usually obtained from a local source. This type of failure is automatically corrected by using a relay to create a path around the node [10]. The relay is

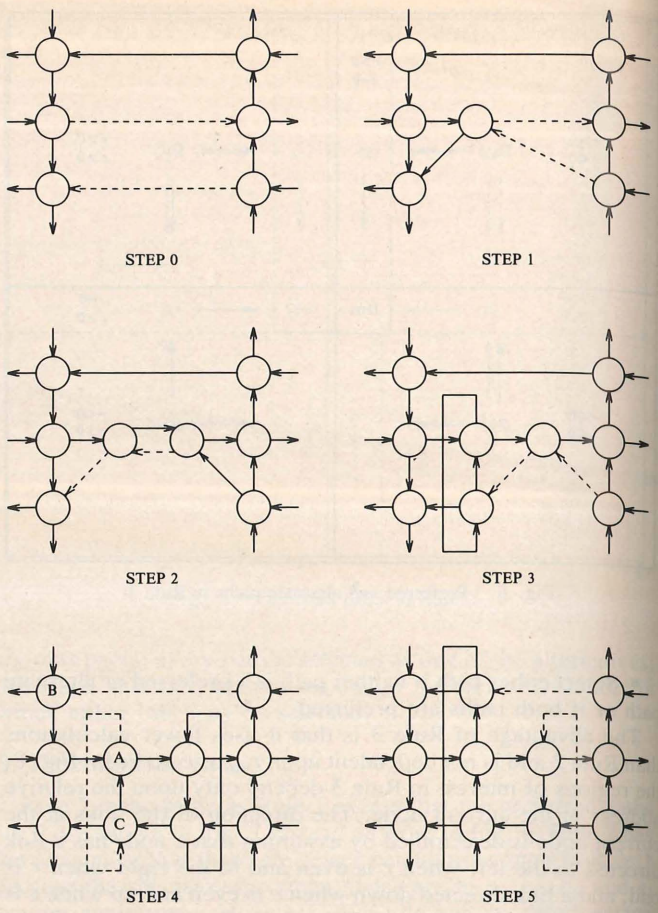

Fig. 9. Adding two columns to an existing network, one node at a time.

TABLE III
THE EFFICIENCY OF LOCAL ROUTING RULES RELATIVE TO THE SHORTEST PATH ALGORITHM AS SINGLE NODES ARE ADDED TO A 6 × 6 NETWORK

Network	Short. Path	Efficiency of Routing Rules				
		Deterministic			Random	
		1	2	3	A	B
6x6	3.71	1.00	.97	.97	.10	.20
add 1	3.77	.98	.95	.93	.10	.21
add 2	3.80	.95	.94	.93 *(.005)*	.10	.20
add 3	3.91	.92	.91	.91 *(.005)*	.10	.19
add 4	3.95	.92	.90	.91 *(.013)*	.10	.19
add 5	3.99	.89	.89	.88 *(.007)*	.09	.19
add 6	4.04	.88	.87	.87	.09	.18
add 7	4.07	.92	.90	.91	.09	.19
add 8	4.16	.95	.93	.93	.09	.18
add 9	4.18	.94	.93	.93	.09	.19
add 10	4.22	.94	.93	.93	.09	.18
add 11	4.26	.97	.95	.95	.09	.17
6x8	4.34	1.00	.97	.97	.09	.17

open when there is power and closes to bypass the node when power is lost. When nodes fail in the MSN, the system is not completely disabled as in a loop; however, packets that arrive at the node are lost and the node should be bypassed to prevent this from occurring. Loss of power at a node in the MSN can operate relays as in a loop system; however, there are two links entering and leaving each node and two relays must be used. The failure recovery procedure selected connects the row through and the column through, as shown in Fig. 10.

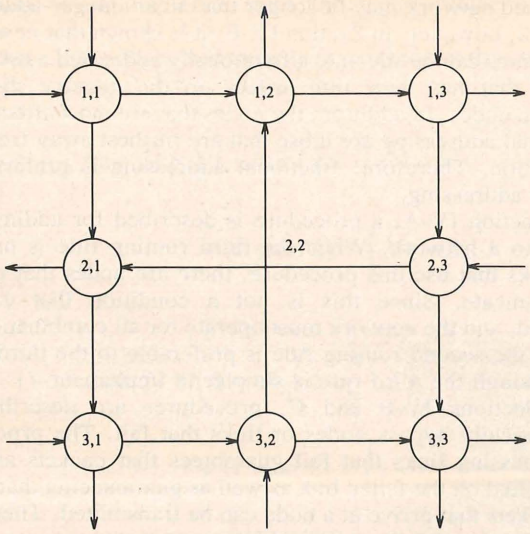

Fig. 10. Operation of the MSN when nodes fail.

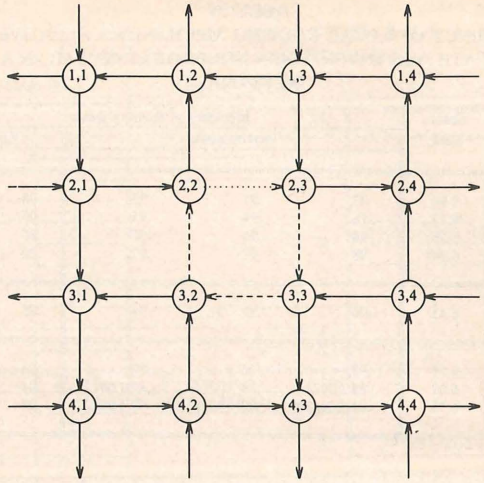

Fig. 11. Operation of the MSN when links fail.

C. Link Failures

In loop systems, a transmitter sends bits continuously, even when there is no information to send. This allows the receiver to retain bit synchronization between packets. Broken links and certain node failures are detected by the loss of signal. More subtle failures can be detected when the periodic start of slot does not occur in systems with fixed size slots or when there are violations of the pseudoternary modulation rules used in wire systems. Loop systems can be designed to bypass segments with failed links by constructing the loop as a series of subloops that start and end at a central location [10]. When the loss of signal is detected on a subloop, the subloop is bypassed and the signal from the previous subloop is switched to the next subloop. This allows a large part of the loop system to operate when links on one of the subloops fail.

The MSN is a slotted system with continuous transmission on each of the links; therefore, failures can be detected on the links arriving at a node as in a loop system. When a link has failed and is detected at the termination node, the origination node of the link must be informed. Otherwise, packets that are transmitted on the inoperable link will be lost, and if packets between a pair of nodes are always routed along that link, the pair of nodes will not be able to communicate. In addition, the implementation of the MSN described in Section II does not lose packets because the node can transmit as many packets as it receives. In order to preserve this characteristic, when a link leaving a node fails, data on one of the incoming links must stop so that the in-degree and out-degree of the node remains the same.

One way to prevent transmission on a link that has failed and to keep the in-degree equal to the out-degree at every node is to stop transmitting on a directed cycle of links that includes the link that has failed. The in-degree and out-degree of each node in the cycle is reduced by one and the link that has failed is not used. To minimize the effect that this strategy has on the throughput and connectivity of the network, the number of links in the cycle must be kept as small as possible and the cycle should not pass through any node twice.

A simple rule that meets these conditions most of the time is to stop transmitting on a row if a signal is not received on a column and to stop transmitting on a column when a signal is not received on a row. The operation of this rule is shown in Fig. 11. In this example, the dotted link from node 2,2 to node 2,3 fails. According to the rule, the dashed links are taken out of service. That is,

- Node 2,3 receives no signal on the row; therefore, it stops transmitting on the column
- Node 3,3 receives no signal on the column; therefore, it stops transmitting on the row
- Node 3,2 receives no signal on the row; therefore, it stops transmitting on the column
- Node 2,2 receives no signal on the column; therefore, it does not try to transmit on the failed link on the row.

When the failed link is restored, the cycle is returned to service by forcing transmission on this link.

When a single failure occurs in a complete MSN, this procedure removes four links, which is the minimum number of links in a cycle. Also, at most one link is removed at each node. Therefore, this simple rule has the desirable characteristics in this instance. However, when there are mutliple link and node failures or if the network has partially full rows or columns, these characteristics are not always obtained. For instance, if the link from 4,4 to 3,4 also fails, both links to node 3,3 will stop, and an operable node is removed from the network. This removal rule is simple, and it works well when there are a few removals. Since a network will be repaired when removals occur, it is unlikely that there will be many removals, and this simple rule is adequate.

D. Effect on Routing Rules

Simulations were conducted to determine the effect that failures have on the distributed routing rules, and the results are presented in Table IV. The fraction of nodes that are in the network, but cannot communicate using the distributed routing rules, are italicized. In the simulations, a random selection of nodes or links fail in a 10×12 network and the bypass or link removal rules are applied. Each experiment is repeated ten times. This results in the span of values in the 95 percent confidence interval for the average path length being less than 1 percent of the mean for node failures. The span of values in the 95 percent confidence intervals for link failures ranges from 1 to 5 percent of the mean. The efficiency is calculated as in Section IV-E.

When nodes are added to each network in Section V-A, it is assumed that the routing rules that use information about the next node know about the changes. This is reasonable because adding nodes is a planned activity. When links or nodes fail, the network is modified automatically, without operator intervention. After failures occur, if the node that each node is connected to is to be known, a protocol must be developed to distribute this information. In the simulations, the change in network connectivity after failures is not known, and the routing rules operate with incorrect information.

TABLE IV
THE EFFICIENCY OF LOCAL ROUTING MECHANISMS RELATIVE TO THE SHORTEST PATH ALGORITHM WHEN NODES OR LINKS FAIL IN A 10 × 12 NETWORK

Failures Size	Short. Path	Efficiency of Routing Rules				
		Deterministic			Random	
		1	2	3	A	B
8 Nodes	5.94	.93	.91	.92	.05	.09
4 Nodes	6.15	.96	.94	.95	.05	.08
2 Nodes	6.28	.98	.96	.96	.05	.08
1 Node	6.34	.98	.97	.97	.05	.08
None	6.42	1.00	.99	.99	.05	.08
1 Link	6.51	.93	.93	.92	.05	.08
2 Links	6.59	.89 *(.007)*	.88 *(.001)*	.87 *(.001)*	.05	.08
4 Links	6.75	.81 *(.014)*	.80 *(.001)*	.79 *(.007)*	.04	.07

The simulations show that the distributed routing rules operate reasonably efficiently when up to eight nodes fail and are bypassed. When up to four links fail and up to 16 links are removed from the network, the distributed routing rules still operate reasonably efficiently. However, a small fraction of the nodes in the network, shown by italicized numbers, cannot communicate until the network is repaired. In the first deterministic rule, a greater fraction of the nodes cannot communicate compared to the other two rules. The first rule uses information about the edges of the network to improve routing decisions. When this information is incorrect, routing failures can occur.

VI. CONCLUSIONS

The objective of this paper is to study simple mechanisms for routing packets in the MSN. The routing rules must not only operate in complete rectangular networks, but must also operate when single nodes are added and when failures occur.

Three distributed routing rules are described in Section IV that use the regular structure of the MSN to simplify routing. The first rule provides the shortest path between any source and destination in an integer addressed network. The second is simpler to implement than the first rule, but results in slightly longer paths. The third rule is the simplest to implement; it has the same path length as the second rule in complete networks, but it does not determine equal length shortest paths as well as the second rule. The nodes with longer path lengths in the second and third rules are those that are furthest from the destination. Networks are divided into communities of interest and nodes communicate more frequently with nodes that are nearby than with nodes that are further away. Therefore, the simpler routing rules are preferable to the first rule.

In Section III-A, a fractional addressing scheme is described that does not change network addresses when new rows or columns are added. In addition, in a fractionally addressed network, the routing rules are independent of the number of rows or columns in the network. This makes the arithmetic unit that calculates relative addresses simpler to implement in a fractionally addressed network than in an integer-addressed network. The path selected between nodes in a fractionally addressed network may be longer than in an integer-addressed network; however, in Section IV-E, it is shown that new rows or columns can be added to a fractionally addressed network in a way that has very little effect on the average distance between nodes. In addition, the nodes that are most affected by fractional addressing are those that are furthest away from the destination. Therefore, fractional addressing is preferred to integer addressing.

In Section IV-A, a procedure is described for adding new nodes to a network. When the third routing rule is used in networks that use this procedure, there are nodes that cannot communicate. Since this is not a condition that can be repaired, and the network must operate for all combinations of nodes, the second routing rule is preferable to the third rule, even though the third rule is simpler to implement.

In Sections IV-B and C, procedures are described to automatically bypass nodes or links that fail. The procedure for bypassing links that fail guarantees that packets are not transmitted on the failed link as well as guaranteeing that all of the packets that arrive at a node can be transmitted. There is a small fraction of the nodes that cannot communicate when multiple link failures occur. Unlike adding nodes, multiple failures should be repaired, and it is unlikely that the network will operate in this mode frequently. Therefore, this condition does not preclude using these failure recovery mechanisms.

APPENDIX

Theorem: The first routing rule, Section IV-A, selects all possible shortest paths to the destination in a complete, integer addressed MSN.

To prove this theorem,

1) a hypothetical distance function $d_{sp}(r, c)$ from every node to the destination is defined,

2) a path from (r, c) to the destination that has this distance is found,

3) it is shown that a shorter path does not exist,

4) it is shown that the routing rule can select every path with this distance, and does not select any paths with a larger distance.

The distance function that will be tested is

$$d_{sp}(r, c) = |r| + |c| + x_{sp}(r, c)$$

where

$$x_{sp}(r, c) = \left(2 * (r \bmod 2) * (c \bmod 2) * \left(1-\delta\left(r-\frac{m}{2}\right)\right) * \left(1-\delta\left(c-\frac{n}{2}\right)\right) * ((1-U(-r)) * (1-U(-c))\right)$$
$$+ (2 * (1-r \bmod 2) * (c \bmod 2)) * ((1-U(-r)) * U(-c))$$
$$+ \left(2+2 * (1-r \bmod 2) * (1-c \bmod 2) * \left(1-\delta\left(r+\frac{m}{2}-1\right)\right) * \left(1-\delta\left(c+\frac{n}{2}-1\right)\right) - 4 * \delta(r) * \delta(c)\right)$$
$$* (U(-r) * U(-c)) + (2 * (r \bmod 2) * (1-c \bmod 2)) * (U(-r) * (1-U(-c)))$$

and

$$U(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases} \quad \text{and} \quad \delta(x) = \begin{cases} 0 & \text{for } x \neq 0 \\ 1 & \text{for } x = 0. \end{cases}$$

The correction factor $x_{sp}(r, c)$ is included in the distance function to account for longer paths that must be taken because of the unidirectional links. It adds two to $d_{sp}(r, c)$ for nodes that only have links directed away from the destination. There are two additional links added to $d_{sp}(r, c)$ for all nodes in Q_3 because packets from this quadrant must pass the destination and return. In addition, there is a modification in $x_{sp}(r, c)$ at the edges of the first and third quadrant that is caused by the

destination being displaced from the exact center of the network.

A node (r, c) is connected to node (r_{nxt}, c) along the column and (r, c_{nxt}) along the row where

$$r_{nxt} = r - (1 - c \bmod 2) * \left(1 - m\delta\left(r + \frac{m}{2} - 1\right)\right) + (c \bmod 2) * \left(1 - m\delta\left(r - \frac{m}{2}\right)\right)$$

and

$$c_{nxt} = c - (1 - r \bmod 2) * \left(1 - n\delta\left(c + \frac{n}{2} - 1\right)\right) + (r \bmod 2) * \left(1 - n\delta\left(c - \frac{n}{2}\right)\right).$$

The change in this distance function between the (r, c) and (r_{nxt}, c) is

$$\Delta_c(r, c) = d_{sp}(r_{nxt}, c) - d_{sp}(r, c) = (|r_{nxt}| - |r|) + (x_{sp}(r_{nxt}, c) - x_{sp}(r, c))$$

and between (r, c) and (r, c_{nxt}) is

$$\Delta_r(r, c) = d_{sp}(r, c_{nxt}) - d_{sp}(r, c) = (|c_{nxt}| - |c|) + (x_{sp}(r, c_{nxt}) - x_{sp}(r, c)).$$

In an $m \times n$, integer-addressed network, the regions in Fig. 6 are

$$Q_1 = \left\{(r, c) : 1 \le r \le \frac{m}{2},\ 1 \le c \le \frac{n}{2}\right\}$$

$$Q_{1-} = \{Q_1 \cap \overline{(r_1 \cup c_1 \cup I_1)}\}$$

$$r_1 = \left\{(r, c) : r = \frac{m}{2},\ 1 \le c < \frac{n}{2}\right\}$$

$$c_1 = \left\{(r, c) : 1 \le r < \frac{m}{2},\ c = \frac{n}{2}\right\}$$

$$I_1 = \left(\frac{m}{2}, \frac{n}{2}\right)$$

$$Q_2 = \left\{(r, c) : 1 \le r \le \frac{m}{2},\ -\frac{n}{2} + 1 \le c \le 0\right\}$$

$$Q_{2-} = \{Q_2 \cap \overline{(r_2 \cup c_2)}\}$$

$$r_2 = \left\{(r, c) : r = 1,\ -\frac{n}{2} + 1 \le c < 0\right\}$$

$$c_2 = \left\{(r, c) : 1 \le r \le \frac{m}{2},\ c = 0\right\}$$

$$Q_3 = \left\{(r, c) : -\frac{m}{2} + 1 \le r \le 0,\ -\frac{n}{2} + 1 \le c \le 0 \cap \overline{r = c = 0}\right\}$$

$$Q_{3-} = \{Q_3 \cap \overline{(r_3 \cup c_3 \cup I_3)}\}$$

$$r_3 = \left\{(r, c) : r = -\frac{m}{2} + 1,\ -\frac{n}{2} + 1 < c \le 0\right\}$$

$$c_3 = \left\{(r, c) : -\frac{m}{2} + 1 < r \le 0,\ c = -\frac{n}{2} + 1\right\}$$

$$I_3 = \left(-\frac{m}{2} + 1,\ -\frac{n}{2} + 1\right)$$

$$Q_4 = \left\{(r, c) : -\frac{m}{2} + 1 \le r \le 0,\ 1 \le c \le \frac{n}{2}\right\}$$

$$Q_{4-} = \{Q_4 \cap \overline{(r_4 \cup c_4)}\}$$

$$r_4 = \left\{(r, c) : r = 0,\ 1 \le c \le \frac{n}{2}\right\}$$

$$c_4 = \left\{(r, c) : -\frac{m}{2} + 1 \le r < 0,\ c = 1\right\}.$$

TABLE V
COMPARISON OF ROUTING DISCUSSIONS FROM RULE 1 AND THE CHANGE IN DISTANCE FROM $d_{sp}(r, c)$

Current Node (r,c)			Column → (r_{nxt},c)		Row → (r,c_{nxt})	
	mod 2 r c	Region	Δ_c	R_c	Δ_r	R_r
Q_1	0 1	Q_{1-}	$3-2\delta(r-(m/2-1))$		-1	X
		c_1	+1			
		r_1+I_1	-1	X		
	0 0	Q_1				
	1 0	Q_{1-}			$3-2\delta(c-(n/2-1))$	
		r_1			+1	
		c_1+I_1			-1	X
	1 1	Q_{1-}				
		c_1	+1			
		I_1	-1	X		
		r_1			+1	
Q_2	1 1	Q_2+r_2	$3-2\delta(r-m/2)$		-1	X
	1 0	r_2	$3-2\delta(c+n/2-1)$			
		Q_{2-}	-1	X		
		c_2			$3-2\delta(r-m/2)$	
	0 0	Q_2+c_2			$3-\delta(c+n/2-1)$	
	0 1	Q_{2-}	-1	X		
Q_3	0 1	Q_{3-}	-1	X	$3-2\delta(c+n/2-2)$	
		r_3			+1	
		c_3+I_3			-1	X
	1 1	Q_3				
	1 0	Q_{3-}	$3-2\delta(r+m/2-2)$			
		c_3	+1			
		r_3+I_3	-1	X		
	0 0	Q_{3-}				
		r_3			+1	
		I_3			-1	X
		c_3	+1			
Q_4	1 1	Q_4+c_4	-1	X	$3-2\delta(c-n/2)$	
		c_4			$3-2\delta(r+m/2-1)$	
	0 1	Q_{4-}			-1	X
		r_4	$3-2\delta(c-n/2)$			
	0 0	Q_4+r_4	$3-2\delta(r+m/2-1)$			
	1 0	Q_{4-}	-1	X		

In Table V, $\Delta_c(r, c)$ and $\Delta_r(r, c)$ are listed for r and c even and odd in each of the regions. Those table entries that are not possible, such as r even in r_2, have been eliminated. An "X" in column R_c indicates that the first routing rule selects the link to (r_{nxt}, c) and an "X' in column R_r indicates that the routing rule selects the link to (r, c_{nxt}).

Table V shows the following.

1) $\Delta_r(r, c) = -1$ or $\Delta_c(r, c) = -1$ for all $(r, c) \neq (0, 0)$. Therefore, each node $(r, c) \neq (0, 0)$ is connected to at least one node that is one closer to the destination. It is possible to travel from any node to the destination in the number of steps specified by $d_{sp}(r, c)$ by selecting a link for which $\Delta_r(r, c) = -1$ or $\Delta_c(r, c) = -1$ at each node along the path. Therefore, $d_{sp}(r, c)$ is a valid distance function from every node (r, c) to the destination.

2) $\Delta_r(r, c) \ge -1$ and $\Delta_c(r, c) \ge -1$ for all $(r, c) \neq (0, 0)$. Therefore, $d_{sp}(r, c) =$ Minimum $(d_{sp}(r_{nxt}, c) + 1, d_{sp}(r,$

$c_{nxt}) + 1$) for all $(r, c) \neq 0$. Since it is not possible to select a link from any node that results in a shorter path to the destination, $d_{sp}(r, c)$ is the shortest path to the destination. This is the basis for many shortest path algorithms and is proven in [11, pp. 193–195].

3) $R_c = X$ or $R_r = X$ for every entry in the table. Therefore, Rule 1 selects at least one outgoing link at every node in the network.

4) $R_c = X$ if and only if $\Delta_c(r, c) = -1$ and $R_r = X$ if and only if $\Delta_r(r, c) = -1$ for all $(r, c) \neq (0, 0)$. Since $R_c = X$ only if $\Delta_c(r, c) = -1$ and $R_r = X$ only if $\Delta_r(r, c) = -1$, routing Rule 1 selects a shortest path to the destination. Since $R_c = X$ whenever $\Delta_c(r, c) = -1$ and $R_r = X$ whenever $\Delta_r(r, c) = -1$, routing Rule 1 can find every shortest path to the destination.

References

[1] N. F. Maxemchuk, "The Manhattan street network," in *Proc. GLOBECOM'85*, New Orleans, LA, Dec. 1986, pp. 255–261.

[2] ——, "Regular and mesh topologies in local and metropolitan area networks," *AT&T Tech. J.*, vol. 64, pp. 1659–1686, Sept. 1985.

[3] E. H. Steward, "A loop transmission system," in *Conf. Rec. Int. Conf. Commun.*, San Francisco, CA, June 1970, pp. 36-1, 36-9.

[4] B. K. Penney and A. A. Baghdadi, "Survey of computer communications loop networks: Part 1," *Comput. Commun.*, vol. 2, pp. 165–180, Aug. 1979.

[5] ——, "Survey of computer communications loop networks: Part 2," *Comput. Commun.*, vol. 2, pp. 224–241, Oct. 1979.

[6] N. Abramson, "The Aloha system—Another alternative for computer communications," in *Fall Joint Comput. Conf., AFIPS Conf. Proc.*, vol. 37, 1970, pp. 281–285.

[7] R. M. Metcalf and D. R. Boggs, "Ethernet: Distributed packet switching for local computer networks," *Commun. ACM*, vol. 19, pp. 395–404, July 1976.

[8] A. G. Greenberg and J. Goodman, "Sharp approximate analysis of adaptive routing in mesh networks," submitted to *Proc. Int. Sem. Teletraffic Anal. and Comput. Performance Eval.*, The Netherlands, June 1986.

[9] R. T. Prosser, "Routing procedures in communications networks—Part I: Random procedures," *IRE Trans. Commun. Syst.*, pp. 322–329, Dec. 1962.

[10] H. E. White and N. F. Maxemchuk, "An experimental TDM data loop exchange," in *Proc. ICC '74*.

[11] H. Frank and I. T. Frisch, *Communication, Transmission and Transportation Networks*. Reading, MA: Addison-Wesley, 1971.

★

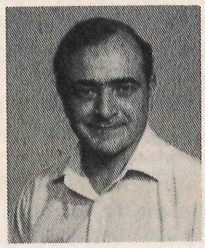

Nicholas F. Maxemchuk (M'72–SM'85) received the B.S.E.E. degree from the City College of New York, New York, NY, and the M.S.E.E. and Ph.D. degrees from the University of Pennsylvania, Philadelphia.

He is the Head of the Distributed Systems Research Department at AT&T Bell Laboratories, Murray Hill, NJ. He has been at Bell Laboratories since 1976, and his research interests include local and metropolitan area networks, protocols, speech editing, and picture processing. Prior to joining Bell Laboratories, he was at RCA Labs, Princeton, NJ, for eight years where his research interests included local area networks, error-correcting codes, and graphics compression. He has been on the adjunct faculty at Columbia University, New York, NY, where he was associated with the Center for Telecommunications Research, and the University of Pennsylvania, where he taught courses in computer communications networks.

Dr. Maxemchuk has served as the Editor for Data Communications for the IEEE TRANSACTIONS ON COMMUNICATIONS, as Guest Editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, and on the Program Committees of numerous conferences. He was awarded the RCA Laboratories Outstanding Achievement Award, the Bell Laboratories Distinguished Technical Staff Award, and the IEEE's Leonard G. Abraham Prize Paper Award.