NBSIR 82-2588

HYBRID GRIDNET
Packet and Circuit Switching in a Single Network

U.S. DEPARTMENT OF COMMERCE
National Bureau of Standards
Institute for Computer Sciences and Technology
Center for Computer Systems Engineering
Washington, DC 20234

October 1982

Task Code RF
Work Unit 00014

Sponsored by the
[Defe]nse Nuclear Agency
[Was]hington, DC 20305

NBSIR 82-2588

HYBRID GRIDNET PACKET AND CIRCUIT SWITCHING IN A SINGLE NETWORK

R. T. Moore

U.S. DEPARTMENT OF COMMERCE
National Bureau of Standards
Institute for Computer Sciences and Technology
Center for Computer Systems Engineering
Washington, DC

October 1982

Task Code RF
Work Unit 00014

Sponsored by the
Defense Nuclear Agency
Washington, DC 20305

U.S. DEPARTMENT OF COMMERCE, Malcolm Baldrige, *Secretary*
NATIONAL BUREAU OF STANDARDS, Ernest Ambler, *Director*

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF FLOW CHARTS

HYBRID GRIDNET
Packet and Circuit Switching in a Single Network

R. T. Moore

This report describes a concept for overlaying a GRIDNET with additional channels and switching facilities that are used to establish point-to-point circuits on a demand basis. Switched connections are established following the exchange of appropriate frames between the stations that are involved. These exchanges use the regular GRIDNET packet switching facilities to provide essential supervisory and control functions including collision avoidance by means of circuit reservation in advance of connection. Switched circuits are automatically disconnected whenever no traffic is observed on them for a designated interval of time.

Alternate routing of a switched circuit is provided by the alternate routing algorithms that guide GRIDNET packets. These algorithms do not distinguish between candidate links that are unavailable because they are busy or because they are inoperable. In either case, an alternate route to the destination is selected as long as one exists.

Key words: Alternate routing; circuit switching; communications networks; distributed control; integrated switching; packet switching; survivability.

INTRODUCTION

The GRIDNET concept is intended to provide a highly survivable communications capability between many thousands of stations occupying positions on many hundreds of interconnected loops. It is a candidate for use in terrestrial applications where the data communications channels are provided by physical media, preferably optical fibers, and where capability to provide continuing communications with surviving nodes is vital despite disablement or destruction of any portions of the network that do not cause it to become fragmented. High survivability is obtained by providing many alternate routes that may be used for the transmission of a message from its source to its destination, and by the use of distributed processing in such a way that route selection and communications control do not depend upon any single node or link. Each loop within a GRIDNET is connected to up to four adjacent loops by stations which we have called gateway stations. Routing decisions are made by these gateway stations on the basis of limited knowledge of the operability status of other loops in the local neighborhood, rather than global knowledge of the operability status of the entire network. As a consequence, the fraction of the available bandwidth that is used to disseminate information about operability status is fixed and not dependent upon network size. Each of the gateway stations on a loop is also capable of performing the master station communications control functions required by the link level protocol. In the event that the gateway performing this activity should fail, one of the other gateways on the loop immediately takes over the master functions. These GRIDNET features insure that a message will be routed from its source to its destination as long as any operable path exists between these points. The penalty that is exacted for this capability is that a message may not take the shortest route to its destination that could be selected by using global network operability information.

The distributed control and alternate routing capabilities of GRIDNET are essential attributes of a survivable data communications network, however; its packet switched mode of operation does not readily support certain classes of service. Simulation results suggest that GRIDNET message packets might experience end to end delays ranging from a few tenths of a second, up to a second or more, depending on network size and loading. Services such as digital voice or video cannot easily accommodate large variations in interpacket arrival times. They are most conveniently handled by end to end circuits, either switched or dedicated, where the transmission delays are either fixed or are varying at a predictable rate.

A GRIDNET could be structured so as to provide a circuit switching capability by overlaying the network with additional channels that share the same cables and station facilities used by GRIDNET. Switching facilities and the necessary regenerative repeaters are also installed in the stations. The switches are actuated following the exchange of GRIDNET messages between the stations that are involved. These cause sequences of these added channels to be linked together to form a connected circuit between designated end points. Once established, a circuit remains connected for as long as it is required as evidenced by the continued flow of traffic. The number of these switched circuits that can be concurrently active is a complex function of the density of overlaid channels, network size, traffic loading, mean switched circuit length expressed in terms of the number of loops traversed, and the geographical distribution of the sources and destinations that are served by the switched service.

In the remainder of this paper, the basic GRIDNET configuration is reviewed and its operation is described. Then, the structure of the additional channels and switching arrangements that are required to provide a Hybrid GRIDNET that is capable of providing integrated switching service (27) are described together with the changes to the operating protocols that are needed to support the supplemental circuit switching capabilities. An example, together with an annotated Flow Chart, is provided to illustrate the operation of the system.

REVIEW OF GRIDNET

In 1979, a novel data communication concept called CROSSFIRE was proposed (1) for an application requiring a high degree of security and integrity. In this system, a primary station controls communications with a number of associated secondary stations using a bit-oriented, link level, protocol. Communication takes place over two loops, each carrying the same data. One loop transfers the data in a clockwise direction and the other carries the same data in a counterclockwise direction. Communication is between any of the several secondary stations and the single primary station that supervises and controls the flow of traffic on the loop. The secondary stations receive identical data that arrives at slightly different times on each of the two loops, and they regenerate the binary signals before forwarding each bit stream to the next adjacent station in the proper direction. This function is performed on all data received by a secondary station, even data addressed to it, without regard to its source. The primary station accepts, but does not regenerate and retransmit, data received on each loop from the secondary stations. Each time a primary station receives a correct copy of a message that it has previously transmitted on one of the loops it serves to confirm that the loops and all of the enroute repeaters are functioning correctly. As a result of the redundant transmission paths, the delivery of a packet to its destination on the loops cannot be prevented by cutting both loops or disabling a node at any single geographic location. The occurrence of such a cut or break can be immediately detected by the primary station since it will cease to receive one or both of the delayed copies of its outgoing transmission. Then, by polling each secondary station on the loop, and determining whether their response is on the clockwise or the counterclockwise loop, the location of the break can be localized to a region between two secondary stations. This region may consist of a single link, or it may include a node and the links that are immediately adjacent to that node.

In addition, by conducting continuous polling during the intervals between the transmission of normal traffic, the loops are never idle for a significant period of time. This aids in the almost immediate detection of any attempt to cut or jam a loop. Further protection against the injection of spurious or unauthorized data is provided by comparing the contents of the data streams that are received over the clockwise and counterclockwise loops on a bit by bit basis. This is in addition to the use of the normal cyclical redundancy frame check sequence. Collectively, these procedures make it virtually certain that any fault, outage, or adversary's action will be quickly detected.

These CROSSFIRE advantages can be extended to large networks by interconnecting a number of CROSSFIRE type systems together using gateway stations at their intersections. Multiple interconnections of the loops and adaptive routing using distributed processing provide the potential for establishing alternate routes between distant pairs of stations despite simultaneous interruptions to the continuity of multiple loops. This conceptual approach has been termed GRIDNET (2).

Network Configuration

In GRIDNET, loops are interconnected in the regular hexagonal array pattern shown in Fig. 1, where stations have not been shown and where loops are shown as smooth, uniform shapes. Gateway stations are located at each point where two loops touch.

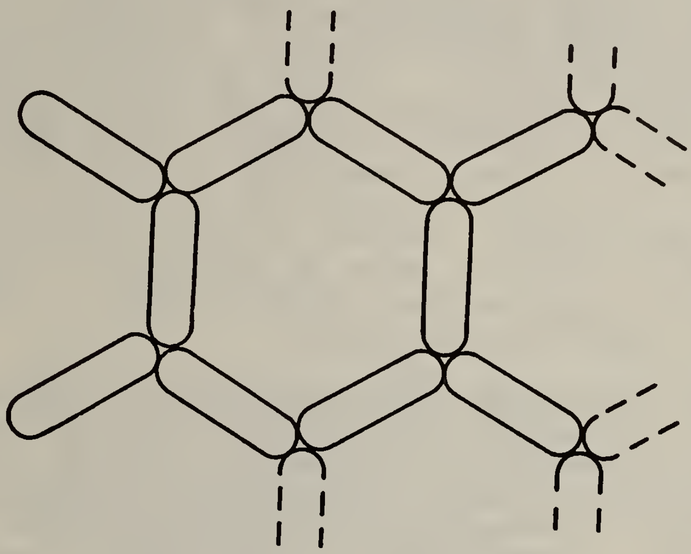

Figure 1. Loops Connected to Form GRIDNET

If this structure is flattened slightly, and the loops are collapsed and depicted as straight lines, the configuration of Fig. 2 results and this graphic representation provides the basis for an orderly addressing scheme.

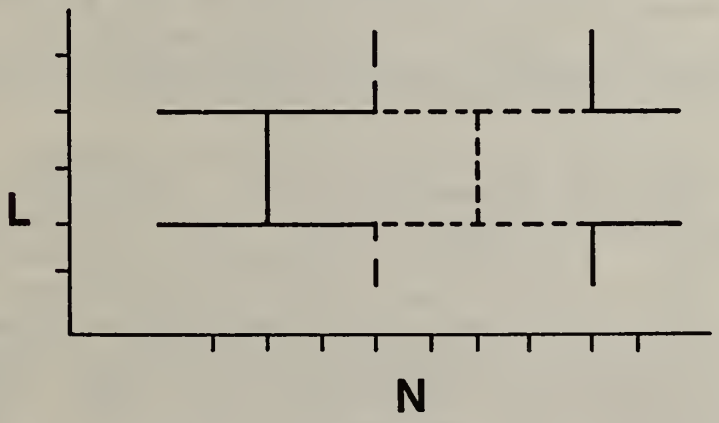

Figure 2. Graphic Representation of GRIDNET

Each GRIDNET address has three parts designated S, L, and N. The S address component represents the "station number" on a loop. The gateway station in the first quadrant on a loop is assigned the value of one for its S address component. The other stations are assigned increasingly higher S component addresses in accordance with their location on the loop as measured in a clockwise direction looking at their graphical representation on the array. The four lowest station numbers are reserved for assignment to the (up to four) gateway stations. The secondary stations are then numbered starting with the value 5. For a 12 station loop, with two secondaries between each pair of gateways, the station number sequence in the clockwise direction is 1, 5, 6, 2, 7, 8, 3, 9, 10, 4, 11, 12. The S component of the address will range from one to some maximum value as defined for the system. Typically, this maximum value will be less than 30. The L component of the address represents a "level" counting from the bottom upward on a topological graph of the system. In Fig. 2, the L values for the loops appear along the Y axis of the array. The final address component, N, represents a "loop number". These are shown along the X axis of the figure.

The L and N address components are not assigned values beginning with the origin of the coordinate system, and this provides room for future expansion of the network in any direction. In a similar fashion, all L,N values need not be occupied provided that lack of occupancy does not fragment the network. The origin of both L and N are such that the loops that appear as vertical interconnectors in Fig. 2 have even values for both L and N, while the horizontal interconnector loops have odd values for these address components. All of the address components must have integer values greater than zero. These address components are used in developing routing information for messages that flow through the network. The L and N portions of the address of a message are used to determine to which enroute loop it should be forwarded in order to reduce the distance to its destination L,N. When a message reaches its destination L,N, then the S address component designates the station on that loop to which delivery is made.

Each loop that is "interior" to the network and fully connected has four gateway stations. Peripheral loops that are only partially connected have fewer gateway stations, normally two. A gateway is logically two stations, one on each loop. Each of the two logical stations has a different address but they can communicate directly with each other by exchanging ownership of buffers containing information that is to be transferred between the two loops. All link control, flow control and routing functions are performed by the gateways. They also perform certain of the network management functions. Each gateway keeps track of the operational status of all of the stations on its home loop and adjacent loops and of the ability of each of the second adjacent loops to communicate with their foreign loops in order to accomplish traffic routing.

The graphic representation of GRIDNET shown in Fig. 2 provides a basis for identifying three types of loops based on the orientation of their connections to adjacent loops. The Type I, II, and III loops are shown in Fig. 3 with their respective gate numbers. The connecting adjacent loops are given as relative L and N values. The connecting loop gateway station number is also shown. These gateway station numbers are used in the routing procedure.

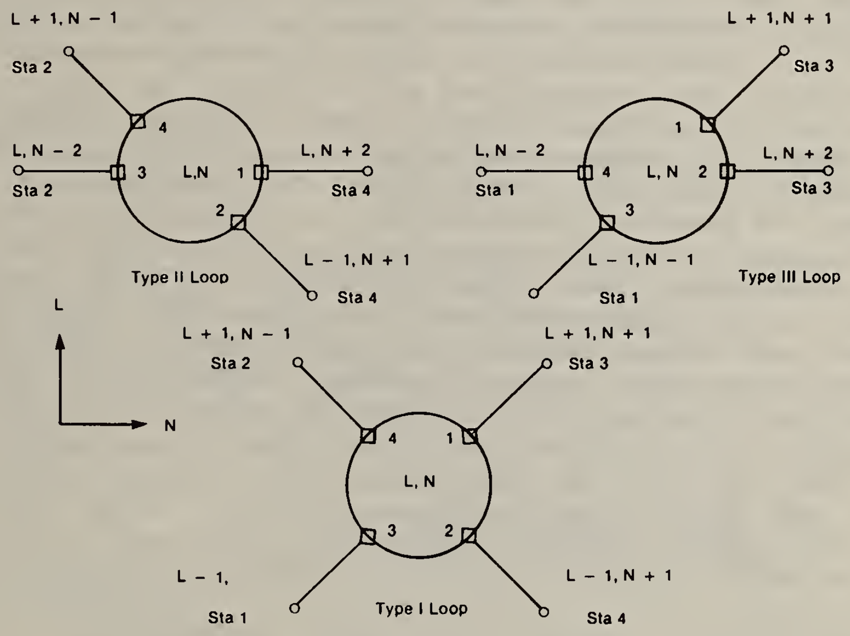

Figure 3. Loop Types and Gateway Station Numbers

Link Level Protocol

The link level protocol proposed for use in GRIDNET is a bit oriented protocol that employs a frame structure that is similar to the one defined in ANSI X3.66-1979, American National Standard for Advanced Data Communication Control Procedures (ADCCP). There is the same arrangement of fields, Flag (F), Address (A), Control (C), Information (I), Frame Check Sequence (FCS), and Flag (F), but the coding and interpretation of the control field has been modified to permit the use of unnumbered commands and responses. Only the S component of a GRIDNET station address is used in the A field of a frame while the full address (L, N, and S) of the source and destination is carried in a designation location within the I field along with a list of gateways defining a proposed route between these points. The latter forms a portion of what is termed an append list.

As in X3.66, a primary station issues commands to the secondary stations on a loop, and the secondary stations respond to these commands. The primary station is thus responsible for controlling all activity on the loop and for the integrity of these activities. Unlike X3.66, however, the role of primary station is rotated among the gateway stations on the loop. This role is transferred from one gateway to another whenever, a) a preset time interval has expired, or b) lack of polling or other activity on the loop indicates that the incumbent primary station has failed, or c) upon granting the request of another gateway station to become primary in order to expedite the delivery of traffic that it may be holding. When a new gateway assumes the status of primary station, it polls all other gateways to get information about their operability status, and their ability to communicate with their immediately adjacent and their second adjacent loops. This information is used to update status tables that are used in the calculation or modification of message routing information.

In X3.66, frames containing information fields are sequentially numbered with a modulus equal to eight for the basic control field format. This permits a station to have up to seven outstanding, unacknowledged frames at any given time, but it also requires that each station maintain certain state variables. Every secondary station maintains a send variable on the I frames that it transmits to and a receive variable on the I frames that it correctly receives from the primary station. Each primary station maintains a corresponding pair of variables for every secondary station with which it communicates. The maintenance and continual resetting of these state variables with each transfer of the role of primary status between the gateway stations of

GRIDNET is an undesirable load of overhead traffic. It can be avoided by operating in a "select/hold" mode where the primary station is responsible for insuring that acknowledgements are exchanged with the secondary station on a frame by frame basis. This mode of operation is particularly appropriate to the dual loop, link level configuration of GRIDNET since the response to a frame is at least partly predicated upon the comparison of FCS and frame contents as received on both the clockwise and counterclockwise loops. With acknowledgement (either positive or negative) on a frame by frame basis, there can never be more than a single outstanding, unacknowledged frame. Therefore, there is no need for frame numbers and the maintenance of state variables to accommodate them. If a primary station is unable to obtain an acknowledgement for a frame, that is an immediate indication of trouble. The frame is redirected along an alternate route if one is available. If there is no available alternate route, the message is returned to its origin to indicate that it could not be delivered. The primary station also generates and dispatches a message to the network maintenance organization requesting service.

In GRIDNET, the Poll/Final (P/F) bit defined in X3.66 is used for flow control. The primary station always sets the P/F bit to one when it expects a response from the secondary and sets it to zero when it does not want a response from the secondary. An example of the latter condition occurs when the primary station acknowledges a final information frame received from the secondary station. The secondary station always sets the P/F bit to one when it has no requirement to transmit another frame, and sets it to zero as a request to the primary for permission to transmit another frame.

Commands and Responses

The commands and responses that are implemented in GRIDNET are identified below. Where their definition, use or interpretaion differs from that given in X3.66, a definition is given below.

1. Exchange Identification (XID).

   The XID command/response has the bit pattern 1111(P/F)101.

The XID command is used to cause the addressed station to report its identification and status. The status report is contained in the information field of the XID response frame.

If the responding station is not a gateway the information field consists of a single octet. The first four bits of this octet are zeros (reserved for future codings). The fifth bit is a one if the frame was received only on the clockwise loop, and the sixth bit is a one if it was only received on the counterclockwise loop. The seventh bit is a one if the station has traffic waiting to be transmitted, otherwise it is a zero. The eighth bit is a one if the station is unable to accept a frame containing an information field (e.g., full buffers), otherwise it is a zero.

If the responding station is a gateway, the information field contains four octets. The first three bits of the first octet are zeros (reserved for future codings, one of which is used in the Hybrid GRIDNET). The fourth bit is a one if the adjacent foreign loop associated with that gateway is inoperable. The fifth through eighth bits convey the same meaning as for non-gateway secondaries. The second third and fourth octets are coded representations of the status of the stations on the adjacent and second adjacent loops that are potentially accessible through the responding gateway. The P/F bit is always set to one in the XID command. In the response frame, the P/F bit is zero only when bit seven of the first octet is one indicating that traffic is waiting.

2. Unnumbered Poll (UP).

   The UP command has the bit pattern 11001100.

   The UP command is used by the primary to solicit response frames from a secondary. It is set with the P/F bit set to one. If the secondary station has no traffic it will respond with a UA. If it has traffic for that primary, it will respond with a UI and include an information field in the response frame. The P/F bit in this frame will be set to zero if the station has more traffic, otherwise, it will be set to one.

   If it is a gateway, and is holding traffic that it wishes to route to another gateway that is not currently the primary, it will respond with a PI with the P/F bit set to one.

3. <u>Primary/Information</u> (<u>PI</u>).

PI is a nonreserved (in ANSI X3.66) response having a bit pattern 11011000. It is used by a gateway that is also currently a secondary and is holding traffic that is to be routed to another gateway that is not the current primary. Its use signifies a desire to be assigned the role of primary at the earliest opportunity so that the traffic can be forwarded directly to the target gateway without having to be relayed through the station that is the current primary.

4. <u>Primary</u> <u>Status</u> (<u>PS</u>).

PS is a nonreserved command/response with the bit pattern 11011001.

The PS command directs the addressed gateway secondary to assume the role of primary. The PS response is the acceptance of the role of primary.

5. <u>Unnumbered</u> <u>Information</u> (<u>UI</u>).

The UI command/response has the bit pattern 11001000.

The UI command directs the addressed secondary to accept a frame containing an information field. (It is normally preceded by a RR command/response). The information field of the frame contains header subfields that must be examined by the station in order to make appropriate disposition of the traffic. The response to a UI command is UA or RNR for positive acknowledgement, or REJ for negative acknowledgement. The UI response is used as a response to the UP command when the addressed secondary has traffic for the primary. Under these circumstances the primary will accept frames using UA, or require their repetition using REJ. It will control the flow of frames with the P/F bit.

6. <u>Unnumbered</u> <u>Acknowledgement</u> (<u>UA</u>).

The UA response has the bit pattern 1100(P/F)110.

The UA is a response used for positive acknowledgement of a frame.

7. Reject (REJ).

The REJ command/response has the bit patern 1001(P/F)000.

The REJ command/response is a negative acknowledgement and request that a frame be retransmitted.

8. Frame Reject (FRMR)

The FRMR response has the bit pattern 1110(P/F)001.

The FRMR response is used to report an error condition that is not recoverable by retransmission of the identical frame, such as receipt of a control field that is invalid, or receipt of a UI frame that had an information field exceeding the maximum length.

9. Receive Not Ready (RNR)

RNR has the bit pattern 1010(P/F)000

RNR is used as a response by a station to indicte a "busy" condition; i.e., the temporary inability to accept additional frames containing information fields. It is a positive acknowledgement of the immediately preceding frame received by the station.

10. Receive Ready (RR).

RR has the bit pattern 1000(P/F)000.

The RR command is used by the primary station to request a secondary station to reserve a buffer for a frame containing an information field. The RR response is used by a station to indicate that it is ready to receive a frame with an information field. It indicates that a buffer is available for that information field and is one way to report the end of a station "busy" or buffers full condition.

Timer Operations

Five timers are required by a GRIDNET gateway station. Of these five, only one, for the end-to-end acknowledgement response to a delivered message, is also required by the other stations on the network. Two of the five timers are only used when the gateway is also functioning as the primary station. The five timers are identified below.

1. T1 A watchdog timer that is reset by the occurance of a data stream on either the clockwise or counter-clockwise loop. When T1 expires, T2 is started in each gateway on the loop.

2. T2 A timer to avoid simultaneous attempts by more than one gateway to assume the role of primary following an outage of the currently acting primary station. The value of T2 is different at each of the gateways on a loop.

3. T3 A timer to initiate the routine transferral of primary status from gateway to gateway around the loop. T3 is started when a gateway assumes the role of primary. When T3 times out, primary status is passed to the next gateway around the loop clockwise. Note that primary status may also be transferred in response to a request prior to the expiration of T3 at the current primary.

4. T4 A timer defining the maximum permissible time for an acknowledgement to a frame to be delayed on a loop. Expiration of this timer is indication that an outage has occurred involving the addressed secondary station.

5. T5 A timer defining the maximum permissible end-to-end response time for a message acknowledgment. Expiration of this timer indicates that the message may have been lost and that the addressee is no longer accessible over the network.

Gateway Functions as a Primary Station

The functional activities of a gateway station acting as a primary are described in the functional flow diagram shown in Fig. 4. The order of activities associated with status polling, normal polling, message routing and transfer, and primary passing are shown.

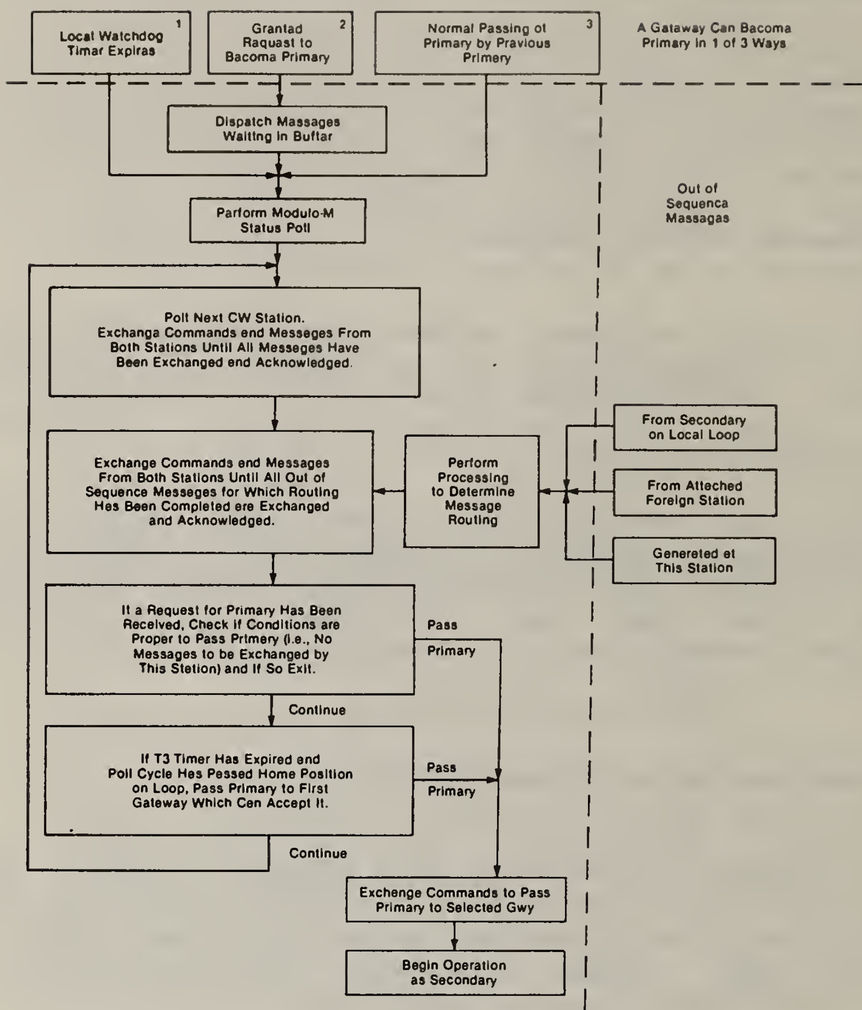

Figure 4. Gateway Functions as a Primary Station

Message Transfer

All message transfers on a loop are controlled by the acting primary station. The primary polls each active station in its status list in sequence and delivers or receives all messages to be transferred. Each message transferred to a secondary station must be acknowledged before the next message can be sent or before the poll cycle can continue. Message transfer between loops is accomplished via a buffer memory shared by the connected stations forming a gateway. A station puts a message for its connected foreign station into the buffer, and a signal is sent to inform the station that a message is ready to be picked up. After the message has been picked up from the buffer and the appropriate routing checks have been performed, the message is available for out of sequence message transfer on the new loop.

Each gateway station does an operational status check on the gateway station on the adjacent foreign loop to which it is connected. This operational status check is done to prevent messages from being routed to a gateway which is no longer a through passage. The operational status check is performed by the use of tally registers of the connected foreign station. A gateway will reset its own tally register to zero and increment the connected foreign station's tally register each time the gateway performs a buffer access or a tally register value test. Normally, the paired tally registers are continually being reset to zero by one gateway and incremented by the other. An unusually high value in one tally register indicates that the corresponding station is not operational. When this occurs, the other station assumes responsibility for any messages left in their common buffer, and performs the appropriate message re-routing and transfer functions.

GRIDNET Path Predictor Routing Algorithm

The path predictor portion of message routing generates a path to the message's destination using known available gates. The available gates are determined from the normal connectivity of the network and the gate outage information carried in the message's append list. This information is updated at each gateway. Each gateway maintains a status table that contains a list of all inoperable gates on the gateway's home loop, its adjacent loops and its second adjacent loops. In Fig. 5, each of the gateways on the loop labeled A has knowledge of the operability status of the gateways on the loops labeled B and C. When a message first enters a loop, it copies any new outage information from the gateway's status table to its append list.

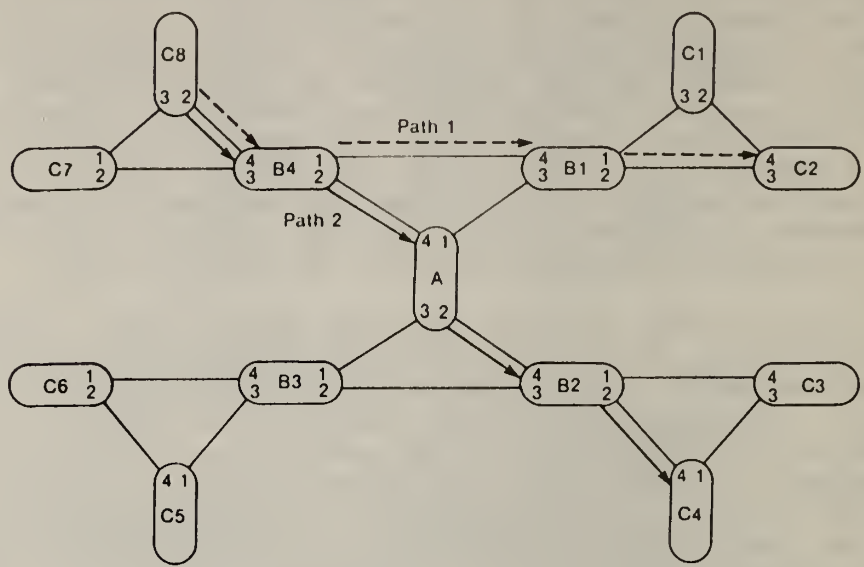

Figure 5. Loop A Has Status Information For B and C Loops

The routing portion of the append list contains a sequential list of the exit gates for each loop the message must enter to reach its destination. The loop L and N numbers are not required in the path list because of the unique GRIDNET addressing scheme. Fig. 5 illustrates this for two possible paths. To follow path 1 in the figure from loop C8 to loop C2, the only information required is that the sequence of exit gates be 2 (on loop C8), 1 (on loop B4) and 2 (on loop B1), and 0 (to designate C2 as the destination). In a like manner, the sequence of exit gates for path 2 from C8 to C4 is: 2, 2, 2, 2, 0. The message will follow the sequence of exit gates until an outage is encountered in either the current loop or the next loop in the predicted sequence. When an outage is encountered, a new path from the message's current location to its destination is calculated.

The top level routing logic for a message entering a loop is shown in Fig. 6. After outage information is copied, the next exit gate in the list is removed. If the value of the exit gate removed is zero, the message has reached its destination loop. If the value of the next exit gate is not zero, the operability status of the exit gates for the current loop and the next loop are checked. If either gateway is inoperable, a new exit gate list must be generated.

In developing the status tables, it is assumed that if loop B1 of Fig. 5 cannot communicate with its adjacent loop C1, then loop C1 cannot communicate with loop B1. Status tables are generated using a three byte status report that is received by each gateway from its next adjacent loop. One byte is for each of the three other gateways on a B level loop other than the gateway connected to loop A. Each byte is split into two groups of four bits. The left four bits contain information on the status of B level gates and the right four bits indicate which gates are available on the attached C level loop. The bit configuration for the left four bits is shown in Table 1.

Table 1
Status of Next Adjacent Loop Gates

Bit Pattern	Meaning	
0000	No Op	
0001	Gate 1	Operational and Can Communicate With Adjacent Loop Station
0010	Gate 2	
0011	Gate 3	
0100	Gate 4	
1001	Gate 1	Operational, But Cannot Communicate With Adjacent Loop Station
1010	Gate 2	
1011	Gate 3	
1100	Gate 4	

The right four bits are set for each gate that is available on the C level with the right most bit for station 1 and the left most bit for station 4.

An example of the format for these status bytes is shown in Fig. 7. In this example, the three bytes shown are for gateways 1, 2, and 4 of loop B1 of Fig. 5. Note that loop B1 is connected to loop A through its gate 3 and as a consequence, no byte in the status report is required for gate 3. In this example, it is assumed that:

- o Gateway 1 of loop B1 can communicate with loop C1 and all Gateways are operational

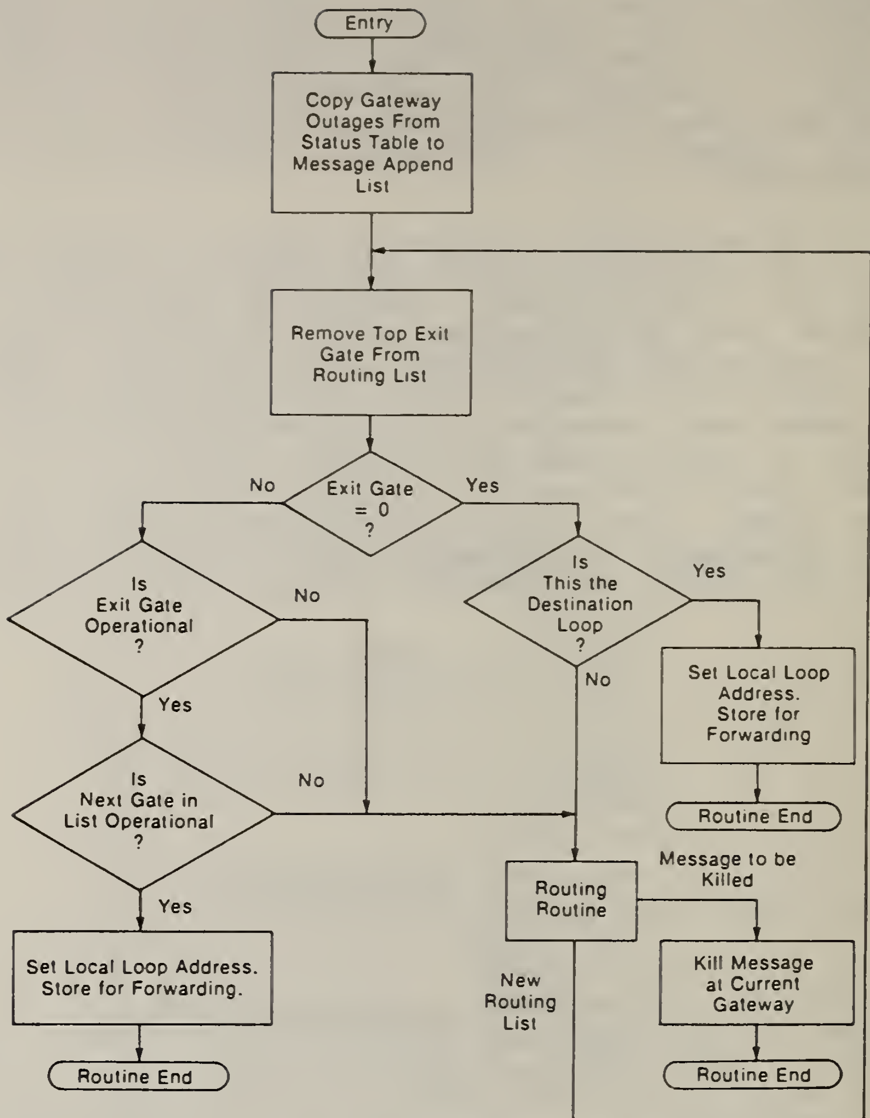

Figure 6. Logic for a message entering a Loop

- o Gateway 2 of loop B1 can communicate with loop C2 and all Gateways are operational except Gateway 2
- o Gateway 4 of loop B1 cannot communicate with loop B4

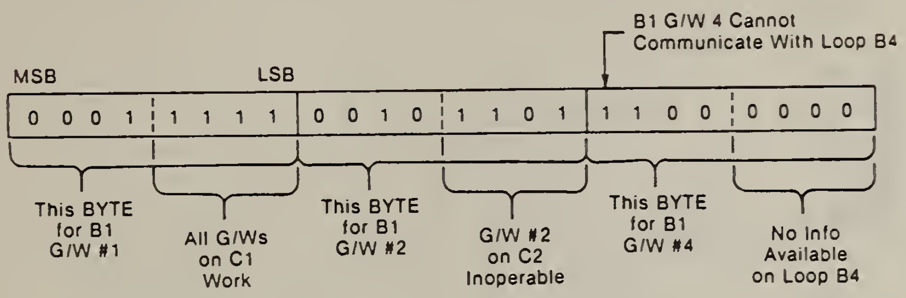

Figure 7. Status Information Format From Loop B1 to Loop A

Method of Path Prediction

The method for creating the list of exit gates works backwards from the destination determining all possible connections until the loop currently holding the message is encountered. The routine creates a temporary table from which the essential gate list is then extracted. To explain the procedure, consider the network of Fig. 8.

In the figure, the nodes at which a loop exists are designated with an "X" and a dashed line is shown between all connected nodes. In the network, all gates to and from the loop at L=4, N=8 are inoperable. When a message going from loop 8,4 to 3,9 arrives at loop 6,6, it must be rerouted because the local status table indicates that gate number 2 of loop 5,7 is out. Therefore, it is necessary to predict a new path to the destination.

To generate the sequential list of exit gates, the routing routine generates the information shown in Table 2. The table consists of a variable number of four word records. A separate record is generated for each loop as required by the algorithm. The first and second words of

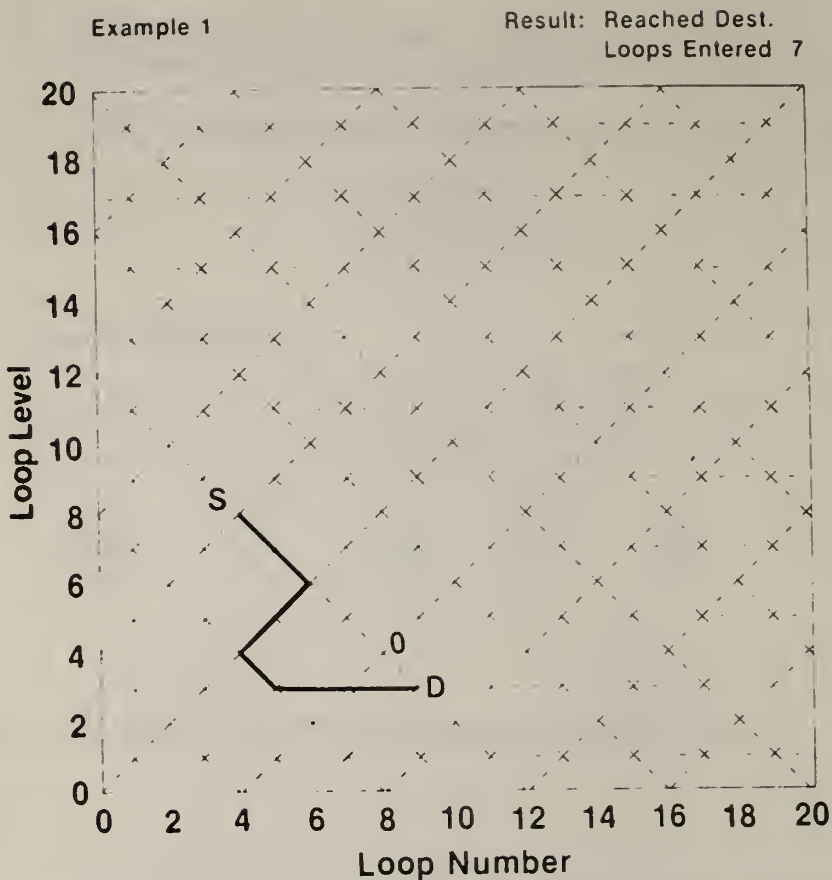

Figure 8. Routing Example

each record contain the L and N address of the loop. Word 3 is the backward entrance gate (i.e. the exit gate to be taken from the loop at level N to reach the loop at level N-1). Word 4 is the index to the record that caused the record to be created.

The first record in the table is for loop 3,9 (the destination loop). The next four records (records 2, 3, 4, and 5) are for the four loops connected to the destination loop's gateways. That is, record 2 is for the loop connected to gateway 1 of the loop in record 1, record 3 is for the loop connected to gateway 2 of the loop in record 1, etc.

After all possible adjacent loop records for record 1 (loop 3,9) have been added, records are added for the adjacent gates of record 2 (loop 3,11). This process is continued exhaustively until the loop performing the routing is reached.

Table 2. Temporary Table Generated by the Routing Algorithm

	Record No.	Word No. 1	Word No. 2	Word No. 3	Word No. 4	Level
		L	N	Exit Gate	Record Creating This Record	Hops From Destination
Destination →	1	3	9	(0)	0	0
	2	3	11	4	1	1
	3	2	10	4	1	1
	4	3	7	(2)	1	1
	5	4	8	2	1	1
	6	4	12	3	2	2
	7	3	13	3	2	2
	8	1	11	4	3	2
	9	1	9	1	3	2
	10	2	6	1	4	2
	11	3	5	(1)	4	2
	20	4	4	(2)	11	3
	32	5	5	(3)	20	4
Current Loop →	45	6	6	(3)	32	5

Exit Gate List From Current Loop to Destination: 3, 3, 2, 1, 2, 0

Records are not added under the following conditions:

1. The connecting gateway is not operational
2. A record for the adjacent loop already exists

The sequential gate list is then created by working backwards through the table. The backward entrance gates, word 3, become exit gates when going from the last record toward the first. Word 3 of the last record is the exit gate for the current loop. The next exit gate is contained in the record that created the last record. Word 4 is a link field pointing to that record. Word 3 of each linked record is copied in sequence to the path list. Table 2 shows a portion of the table used for the routing performed at loop 6,6 in Fig. 8. In the table, a fifth word has been added to each record that shows the number of hops (loops) from the destination. It can be shown that this method of routing always gives a minimum number of hops between the current message location and the destination provided there are no intervening outages.

If a path does not exist between the current location and the message's destination, the routine will run out of records to be generated before a record for the loop with the message is generated. In this case, the routing algorithm will exchange the destination and origination addresses and the message will be returned to its origin as an indication that it could not be delivered. If further outages develop that prevent the message from being successfully returned to its origin, the message will be killed. If positive confirmation of successful delivery is required, it is provided by an end-to-end confirmation message.

DESCRIPTION OF HYBRID GRIDNET

In this section, the additional channels, repeaters, switching facilities and supervisory functions required to implement a Hybrid Gridnet are described. Unless otherwise indicated, these additional facilities and capabilities are overlaid on a GRIDNET and do not replace or substitute for any of the basic GRIDNET features or capabilities that were described in the preceding sections. Operation of a GRIDNET is not dependent upon the operability, or even the existence of any of the circuit switching facilities of the hybrid system. The converse is not true; connection of the circuit switched facilities is completely dependent upon the GRIDNET, however, once a switched circuit has been established, its continued operation is no longer dependent on the packet switched portion of the network.

In the Hybrid GRIDNET, circuits are established between two stations as the result of the transmission of special connection request messages over the regular GRIDNET. These messages propagate from source to destination using the normal GRIDNET routing algorithms and communications control protocols. Circuits are established by these messages on a link by link basis using an advance reservation system to insure that two circuits that are propagating through the network simultaneously do not collide with each other. As a circuit is being established, a busy link may be encountered on the preferred route. When this happens, the GRIDNET alternate routing algorithms are invoked and the busy link is treated just as though it were inoperable. Thus, as network loading of switched circuits increases, the actual path length between an arbitrarily chosen pair of stations may be greater, but some path will be found so long as an operable and not busy path exists.

Once a path has been found and a circuit established between two stations, it should remain connected for as long as it is needed and then it should be disconnected so that its component links can be reused in forming other circuits. Three options were considered for performing the disconnect function.

The first option involved the transmission of a special bit pattern over the circuit. Special circuitry at each switch would be arranged to recognize that pattern and to disconnect the switch connections. This approach has the advantages of tolerating variable delays in the link to link circuit connection process and of providing a very fast and efficient disconnect process that could be set in motion by a brief transmission from either end of the circuit. It suffers from the disadvantage of requiring special circuitry at each switch that can recognize the unique pattern, and from the requirement that this pattern be protected during normal data transmission. There is the further hazard that an error occurrence might change a data stream into the protected pattern and cause an undesired disconnection. Finally, link failures might prevent the dissemination of the disconnect signal over the complete circuit and might leave some links still connected.

A second option involved the transmission of individual messages to each gateway station directing them to disconnect the circuit. This option requires a minimum of special purpose hardware at each switch, but otherwise it is costly in terms of overhead, not only from the standpoint of many messages, but also from the fact that a record would have to be maintained of each link that was used at each switching point in order to disconnect only the desired circuit. This scheme could also leave some links connected in the event of outages that fragmented the network.

The third option is the use of timers that will automatically cause switch disconnection if no signals are observed on the circuit for specified time intervals. This approach guarantees that circuits will be disconnected even if the network is fragmented. It requires minimal special purpose circuitry at each switch, nominally only signal detectors and timers. It has the disadvantage of requiring idle fill during circuit establishment and idle periods in order to avoid disconnection. This can actually be turned to an advantage by establishing different time periods before disconnect occurs depending upon whether signals cease on one or both channels of the circuit. Once an originating station has been instructed to connect to a circuit that is being established, it transmits continuously on its outgoing channel until the circuit is completed or the attempt to establish is abandoned. As each successive link is connected, call progress signals are sent back to the originating station. When signals are removed from one channel of the circuit, but persist on the other channel, this is interpreted as a call in progress and a long delay is invoked before disconnect occurs. When signals are removed from both channels the circuit is disconnected after only a short delay. This allows a longer period of time out to disconnect to be used during the circuit establishment phase and a shorter disconnect time out to be invoked when traffic ceases in both directions. This is the option that is selected for use in the Hybrid GRIDNET

Configuration of Hybrid Facilities

Additional facilities that are required to support a Hybrid GRIDNET include supplementary channels together with regenerative signal repeaters, busy condition detectors, timers and switching facilities. Their general arrangement is described in the following sections.

Configuration of Circuit Switched Channels

Additional physical channels are overlaid on a loop within a GRIDNET in modular groups. The basic group consists of one transmit channel and one receive channel that can connect each gateway station on the loop with those secondary stations that are located between that gateway station and the next clockwise gateway station. Each gateway station is also provided with one transmit and one receive channel to each of the other individual gateway stations on the loop. Thus, in a fully connected loop that is interior to the network, each gateway station has four transmit and four receive channels which make up four (two-way) circuits. Three of these circuits, which will be designated by the letter "T" (trunk), connect with other three gateways. The fourth, designated by the letter "S" (secondary), can be connected to one or more secondary stations that are serviced by that gateway. The arrangement of these overlaid physical channels is shown in Fig. 9.

Here, stations numbered one through four are gateway stations, and stations numbered five through 14 are secondary stations. Stations five and six have access to a transmit and a receive channel which form a circuit to gateway station one. The switching facilities at gateway station one permits this circuit to be interconnected to other circuits which may go to any other gateway station on the loop, or which may be switched through the gateway to similar circuits on an adjacent loop (not shown in this Figure). If a switched circuit was required between stations six and nine, it could be routed via six, one, two, and nine, or it could go via six, one, four, two and nine, or even six, one, three, two and nine. Note, however that if a switched circuit was established between stations six and nine this would prevent the establishment of a circuit to station five, since both five and six share a single circuit in "party line" fashion. Note also that this shared service is not mandatory. More of these basic groups of this basic

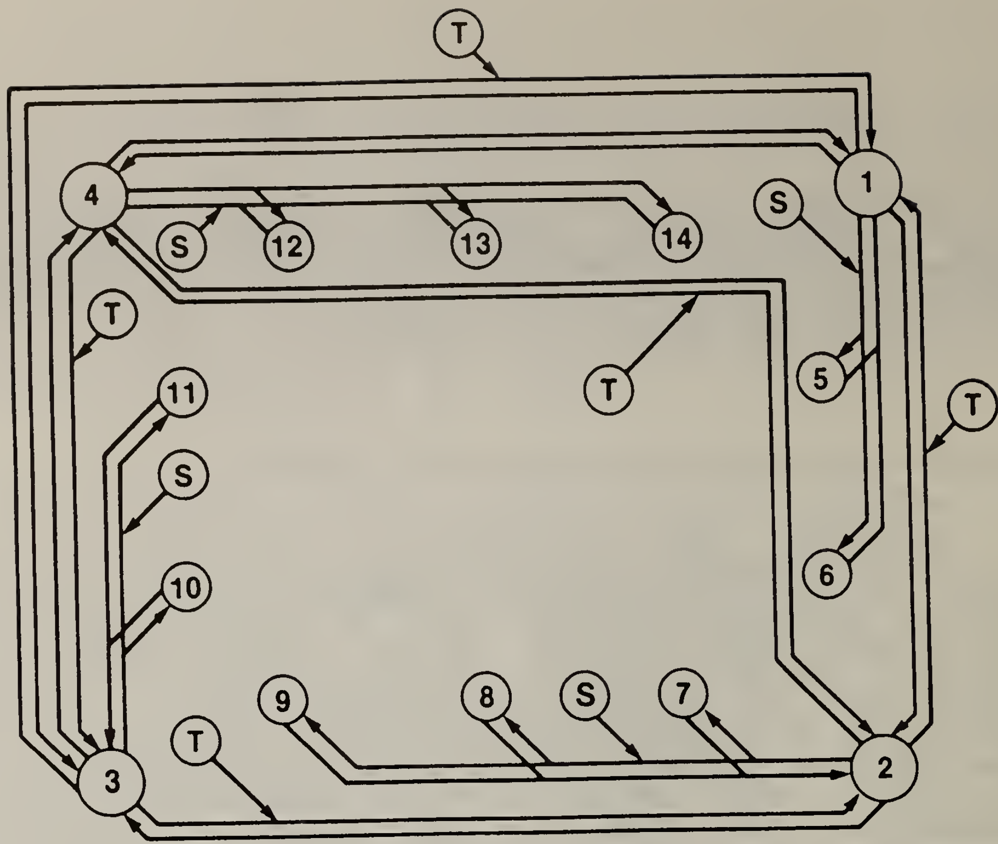

Figure 9. Arrangement of Switchable Channels

group of circuits can be overlaid on the loop to provide whatever quality of service might be desired, up to the limit of concurrent use of a switched circuit for every station on the loop. Since the circuit switching is established by GRIDNET packets, the connections between end points must be completed on a link by link basis and only one circuit at a time can be set up within a given loop, even though the loop may concurrently support other circuits that were previously established. In the balance of this report it will be assumed that only a single basic group of circuits has been overlaid on each GRIDNET loop.

Configuration at a Secondary Station

Fig. 9 does not show the dual loops and other components that support the GRIDNET packet switching capability. These have been omitted so that the routing of the switched circuits can be indicated more clearly. Even here, significant details have not been shown. For example, the circuit from gateway station one to gateway station two occupies the same physical cable as the circuit from gateway station one to stations five and six. When both of these circuits enter station five, the optical signals (fiber optic media is assumed) may be converted to electronic signals which are passed through a regenerative repeater where the signalling pulses are amplified and reconstituted after a one bit delay. Then the signals are reconverted to optical form and launched on the appropriate outbound fiber optic channel. The signals on the S circuit are always converted to electrical form for possible connection with the station before being reconstituted. The signals on the T circuit may, or may not, be converted to electrical form for regeneration depending upon the path lengths and the requirements for compensation for attenuation. The general arrangement of a secondary station is shown in Fig. 10.

In this figure, the GRIDNET communications controller is shown on the left together with the clockwise and counterclockwise GRIDNET fiber optic loop channels. A switch controller is connected to the communication controller and to a pair of busy detectors, one for the transmit channel and one for the receive channel of the S circuit. This controller is capable of operating or releasing a solid state switch that can connect the data terminal equipment of station five to the S circuit that goes to gateway station one.

The busy detectors at each station monitor both the send and receive channels of the S circuit and provide supervisory information that is employed in both the establishment and termination of connections involving this circuit. These busy detectors provide busy signals to the switch control during and for preset time intervals after the time that signalling pulses are observed on the circuit. This includes signalling in both directions. Whenever the busy signal is active, it inhibits the initiation of a request by that secondary station to establish a connection on the S circuit. If a circuit has already been established on the S circuit, the busy signal causes the switch controller to maintain the connection that has been established by the switch. When the busy signal is removed, the switch connection is automatically disconnected, freeing the S circuit for use by another secondary station.

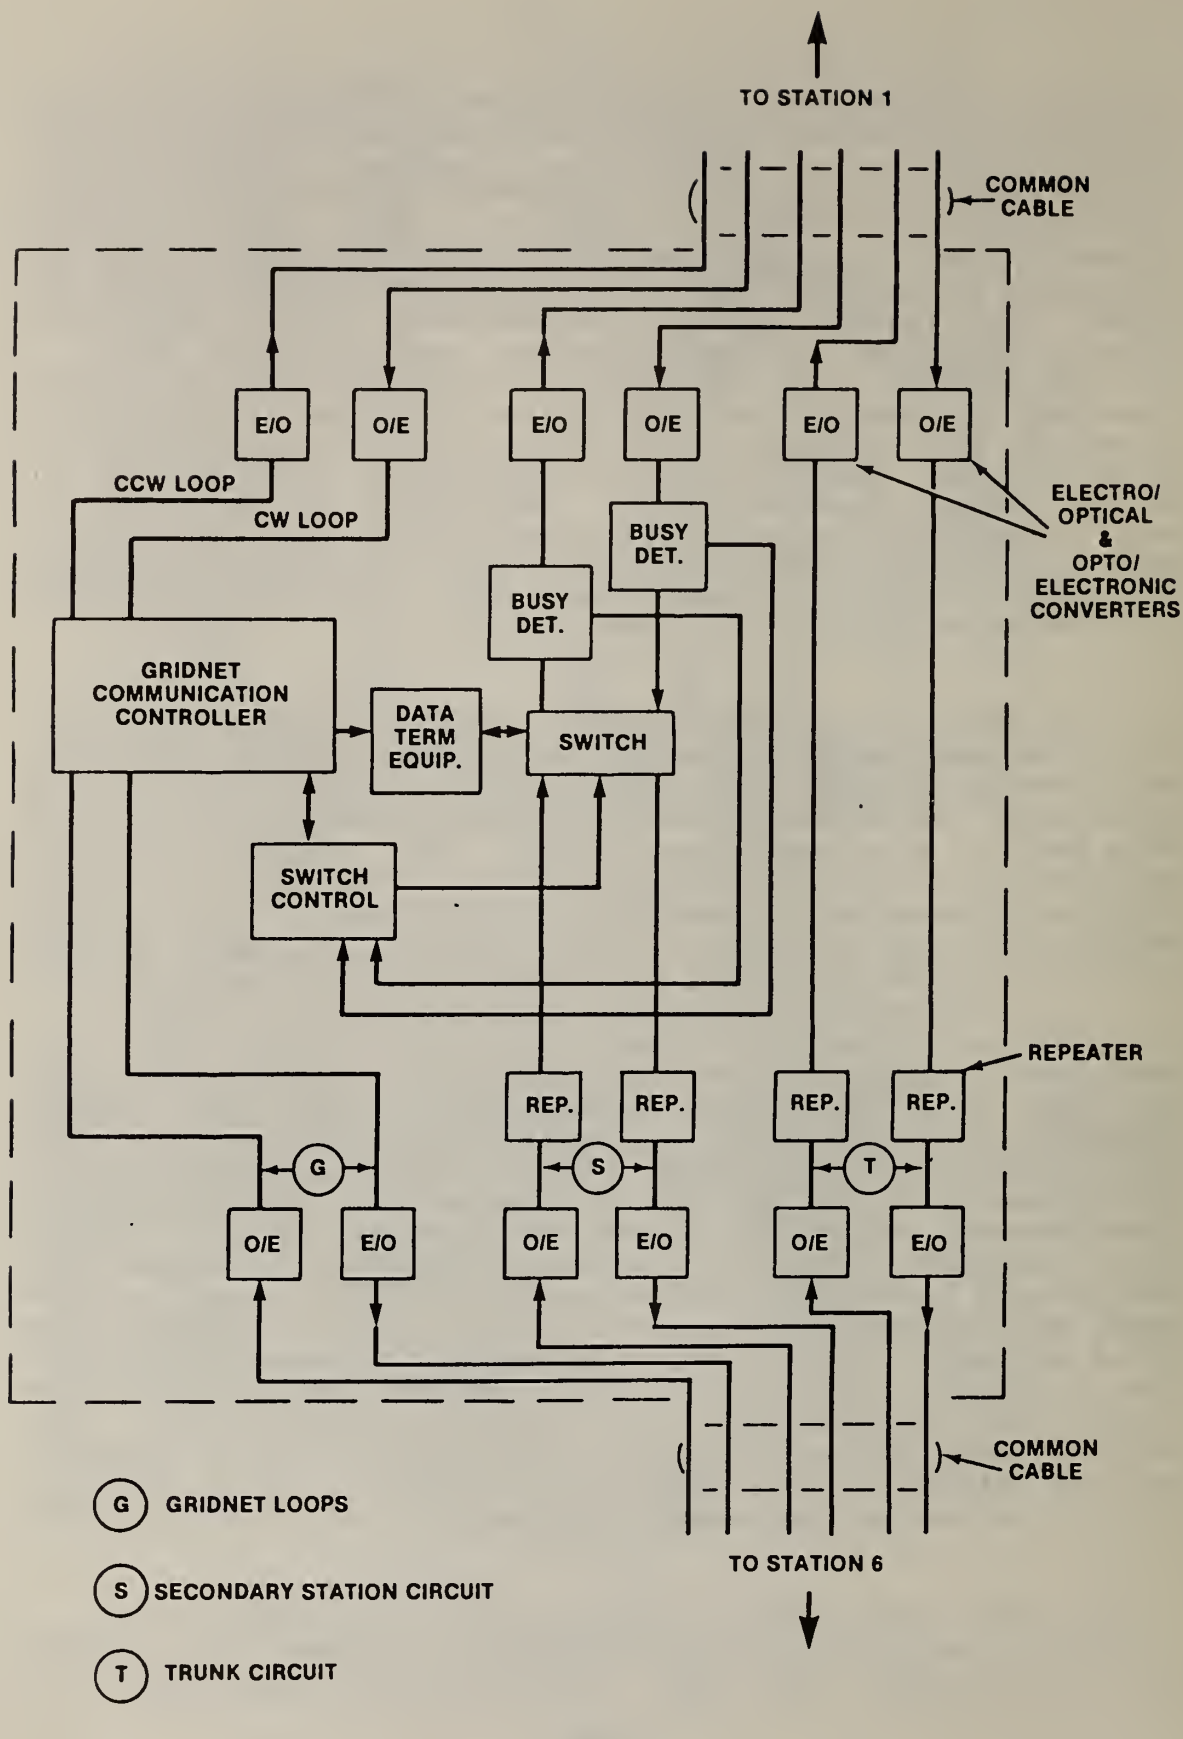

Figure 10. Secondary Station Configuration

Configuration at a Gateway Station

The configuration at a gateway station is similar in some respects to that previously described at a secondary station. It differs mainly in that it serves as a switching terminus for up to three type T circuits as well as for the S circuit of a basic circuit group. Fig. 11 is a simplified block diagram of the general arrangement.

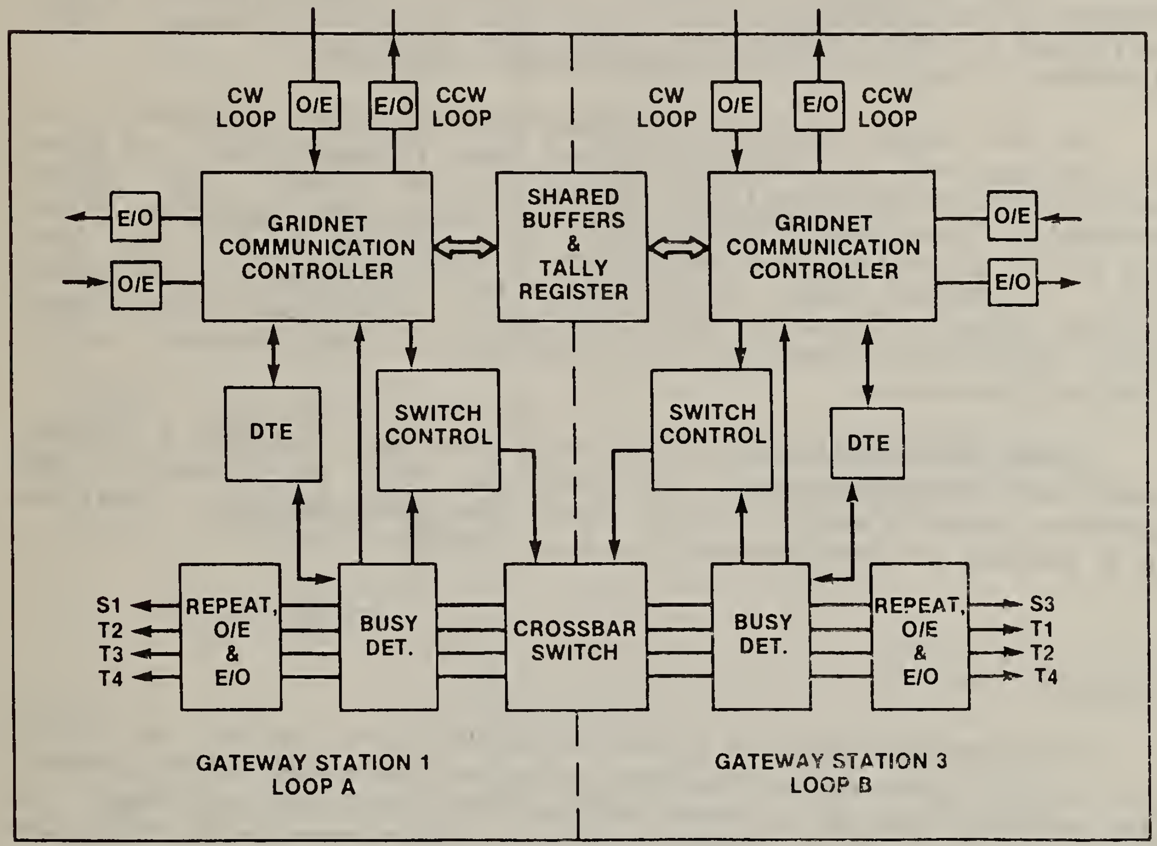

Figure 11. Gateway Station Configuration

This figure is intended to represent the gateway interconnecting station one of loop A with station three of loop B shown in Fig. 9. The two GRIDNET communication controllers are shown at the top of the figure together with their shared buffer memory and tally registers that are used in the transfer of GRIDNET messages between loop A and loop B1. Both of these stations have the capability to actuate a common solid state crossbar switch through a switch controller. Also, they each have access to a seperate switching arrangement that permits them to make a bridging connection to any of their circuits. Busy condition detectors provide signals to the crossbar switch controller that will keep connections up during the intervals while the circuit is being established, and once established, will disconnect any circuit that has been idle for longer than a shorter predetermined time interval. These detectors also provide signals to the GRIDNET communications controller so that switching is not attempted on busy circuits. Bridging connections are not inhibited by busy signals. They are used to insert call progress signals and for maintenance functions.

On the left, station one has available an S circuit to the group of secondary staions that it services. It also has T circuits to stations two, three and four, the other gateways that are on loop A. On the right, station three has its S circuit and T circuits to stations one, two and four. Although these circuits are all shown as single rather than double lines on the figure, it should be understood that each one is made of two channels, a send channel and a receive channel.

Each of the gateways has Data Terminal Equipment (DTE) that can be connected to any of the circuits via the crossbar switch and which is used when that gateway station is a source or destination on a switched circuit.

Crossbar Switch

The details of the crossbar switch are shown in Fig. 12, which also shows the two bridging switches which permit the gateway itself to send and receive on a circuit that is connected by the crossbar switch. This contrasts with the use of the DTE at either station which serves only as a source or destination for a switched circuit and is not capable of being bridged across a circuit that passes through the gateway station.

The four loop A circuits in the basic group are designated as S1 and T2 through T4. The S circuit is always given the same number as the number of the gateway station that services it. The T circuits are always given the number of the distant gateway station to which it is connected. For example, T4 of the figure is the circuit that connects gateway station one with gateway station four. At gateway station four, that same circuit would be called T1 since there, it connects to gateway station one. If more than one basic group of circuits are overlaid, subsequent groups are identified by a second digit, e. g., T4, T41 and

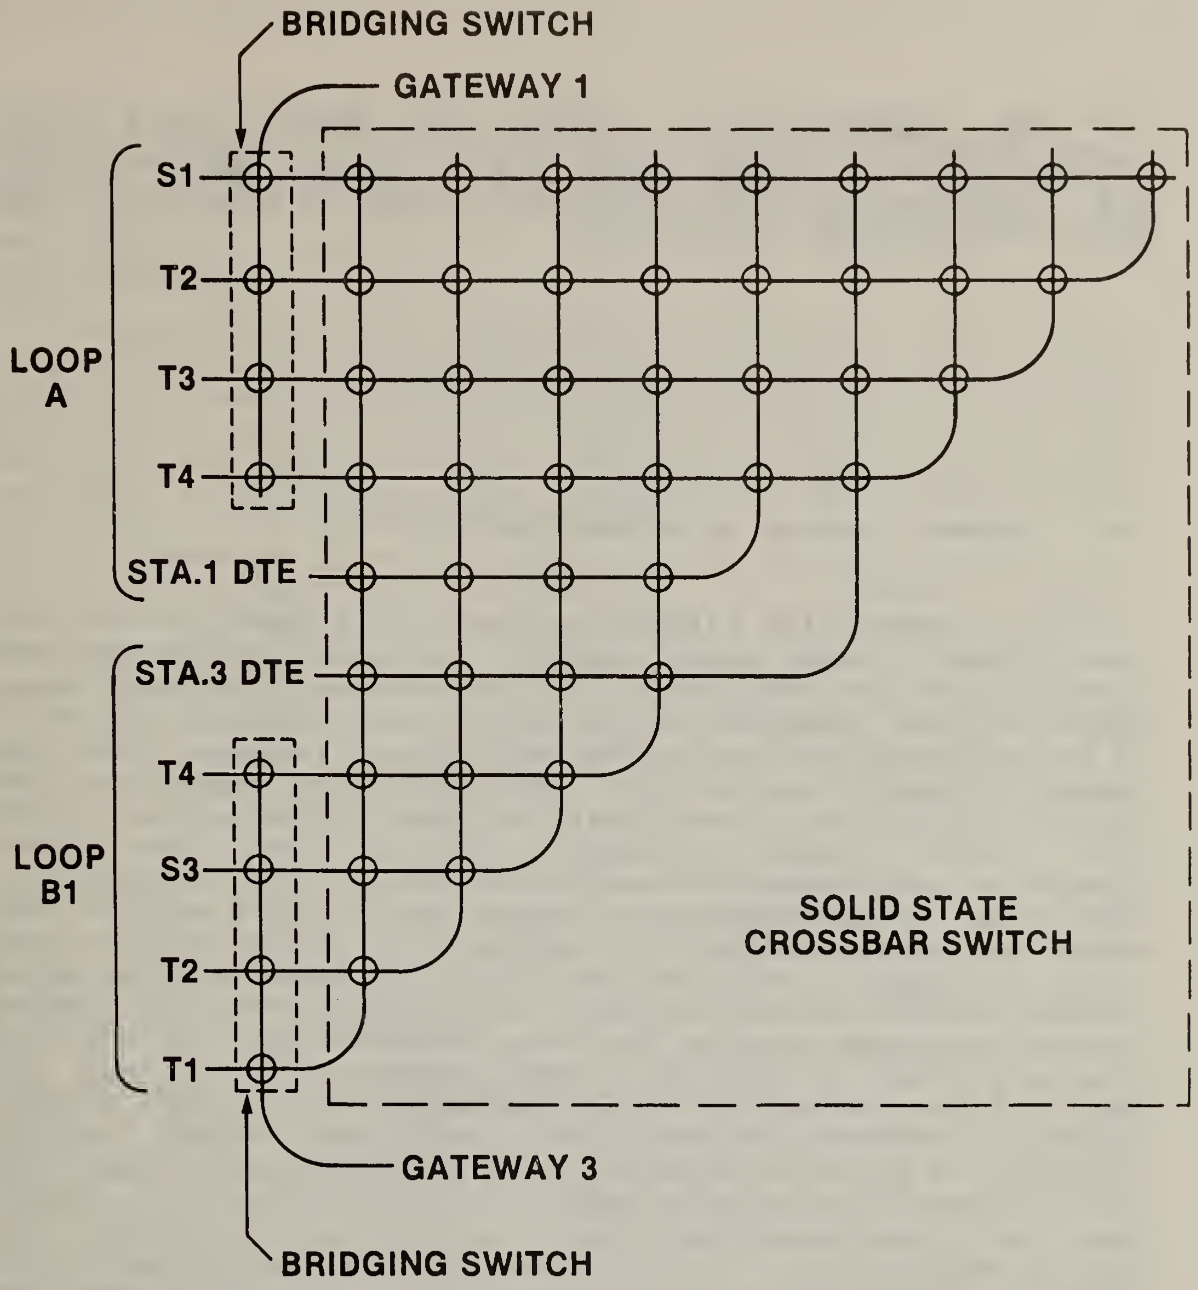

Figure 12. Crossbar and Bridging Switches

T42 would identify three trunks in three basic groups of circuits that connect from the station under consideration to station four on that loop.

The crossbar switch permits any circuit in loop A to be interconnected with any other circuit in either loop A or loop B1. It also permits the DTE in station one and station three to be connected to any circuit in either loop, but it does not permit the two DTE's to be connected together. This is unnecessary since the stations are collocated.

The crossbar switch is arranged to permit only a single connection to be made in any row or column, and the switch for a basic group of circuits will support a maximum of only five simultaneous connections.

Establishment of a Switched Circuit

If station five wishes to establish a switched circuit from itself to some other station, it firsts checks the busy condition of the S circuit. If the circuit is not busy, station five prepares a connect request message, and waits to be polled by the acting GRIDNET primary station over the normal GRIDNET. Assume, for illustrative purposes, that the acting primary is gateway station three. When gateway station three polls secondary station five, the latter transmits its connection request message. This message carries the identification of the destination station with which secondary station five wishes a physical connection. Since secondary station five is not on the S circuit which gateway station three controls, the messages must be relayed to the proper gateway on the local GRIDNET loop, in this instance gateway station one. When gateway station one has received the connection request message from gateway station three, it requests primary status using the normal GRIDNET protocol. As soon as it is granted this role, it tests the busy condition of the S circuit (to station five) to confirm that it is not busy, and then reserves that circuit for the use of station five. This reservation inhibits any other incoming request for the use of the S circuit from establishing a connection. This reservation lasts for a time interval that will permit station five to attach to the S circuit and begin transmitting signals that will actuate the busy detectors. After reserving the S circuit, station one then transmits a GRIDNET message to secondary station five instructing it to attach its Data Terminal Equipment (DTE) to the S circuit.

Upon receipt of the instructions to attach to the S circuit, secondary station five actuates its switch and, while holding the switch operated, starts transmitting continuous flag signals on the S circuit. This transmission actuates the busy detector which acts through the switch controller to keep the switch actuated for the time-out period that is used for circuit establishment. (A much shorter time-out period is used to initiate disconnection.) The flag transmissions by the call originating station are continued until the circuit to the desired destination has been completed, or until the attempt to establish it has been abandoned.

Secondary station five must not attach its DTE to the S circuit, or begin transmission on it, prior to the receipt of instructions to do so from gateway station one in order to permit circuits to be completed only after a path has been reserved for them. For example, suppose that while secondary station five was transmitting its connection request message to gateway station three, gateway one had received a message from its adjacent loop gateway requesting a connection to secondary station six. If secondary station five was already attached to circuit S, the circuit would have been busy. It would not then have been possible to complete the distant circuit and many previously established links might have had to be disconnected. By delaying the attachment of any secondary station until it is directed to do so by the primary station, collisions and other such problems can be avoided.

Gateway station one has been calculating the appropriate routing for the connection request message concurrently with the foregoing. When the routing calculations are complete, the incoming circuit (in this example, the S circuit) is tested to insure that the station initiating the connection request has attached to the circuit as requested. This test also insures that the initial portion of the switched circuit that is being connected is operable. If the test fails, a GRIDNET message is dispatched to secondary station five reporting the failure and indicating that the desired circuit cannot be completed. If the test is successful, the next action depends upon the destination address that station five wishes to be connected with. If station one is the destination, and its Data Terminal Equipment (DTE) is available for connection, it actuates its cross bar switch making the required connection to the S circuit. It then transmits an XID frame with its complete address. This indicates to station five that the circuit has been completed and two way communication is started.

If the selected route requires the use of a T circuit to another gateway on the same loop, that circuit is first tested to make sure that it is not already busy. If it is busy, an alternate route is calculated and this new circuit is tested. When a candidate circuit is found that is not busy, a message is directed to the next gateway on the selected route requesting reservation of the proposed circuit. The recipient of this message tests for busy condition on the designated incoming circuit, and if it is still not busy, reserves it for the requested use and transmits a GRIDNET message back to station one indicating that the reservation has been granted and instructing it to actuate its switch. Upon receipt of this message, station one actuates its cross bar switch interconnecting the S circuit with the reserved T circuit. It then actuates the bridging switch and transmits an XID frame containing its full address as a call progress indicator. These transmissions are sent back to station five, and reset the timers of all busy detectors enroute. The continuing transmissions of station five are sent forward by the actuation of the crossbar switch.

If the selected route involves the use of a T circuit on the adjacent loop, the responsibility for handling the switching functions is transferred to the collocated gateway on that loop by placing the connection request and calculated route in their common message exchange buffer and notifying the other station. Since both gateway stations have access to the crossbar switch and to their own set of busy detectors, this transfer of switching responsibility can be made without actuating any circuit switches and without running any risk of causing a collision between an incoming and an outgoing attempt to build up a circuit. The gateway that is responsible now tests the candidate circuit for availability and sends a message to the next gateway on the selected route and proceeds to attempt to reserve the preferred circuit as described in the preceding paragraph. If successful, it actuates the crossbar switch and transmits the supervisory XID call progress frame as described above.

The circuit continues to be established on a link by link basis. In each case where the connection is extended from one loop to another, the gateway that is on the destination side of the selected route initiates the reservation request for the next circuit. It also actuates the crossbar switch when the reservation is granted. At each step, the responsible gateway first tests the incoming circuit that was previously reserved to make sure that a successful connection has been made by the previous gateway. If that test is successful, it selects an outgoing route that does not test busy and dispatches a reservation request message to the next gateway along the chosen route. If the reservation is granted together with permission to connect, it does so and transmits an XID call progress reporting frame back to the originator via the circuit that is being established. If the reservation is refused, an alternate route is attempted, and the process is repeated. If no alternate route is available, the connection attempt is aborted and the appropriate GRIDNET message is directed back to the originator. Disconnection of that portion of the circuit that had been established is insured at the expiration of the busy signal time-out period.

When the destination of a switched circuit is a station that is located on an S circuit, the last gateway station is the controller of that S circuit. As a result no reservation request message is necessary, and the gateway can make the connection on its own initiative. It then transmits a regular GRIDNET frame to the destination secondary station instructing it to connect to the designated S circuit. When this connection has been made, the circuit establishment is completed.

There are four ways by which a gateway station can successfully act to extend a circuit that is being established, depending upon the target circuit destination. It can, 1) connect to its own DTE, 2) place the connection request in the buffer shared with its adjacent foreign loop station, 3) connect to an S circuit, or 4) connect to a T circuit. The gateway station can perform the first two of these actions while it is in secondary station status, but it will normally acquire primary status in order to connect to either an S or a T circuit.

Additional Commands and Responses

In order to establish a switched circuit as described above, the Hybrid GRIDNET requires two more commands and responses in addition to the ten that were previously defined for GRIDNET. They are defined as follows:

11. <u>Reserve</u> <u>Connection</u> (<u>RC</u>)

    RC is a nonreserved command/response with the bit pattern 1101(P/F)010. It is used to reserve a particular circuit for use in fulfilling a connection request. The RC frame includes a four byte information field that identifies the station of origin for the requested circuit and the particular circuit on the current loop for which the reservation is requested.

    The RC command/response is exchanged between gateway stations prior to actuating the crossbar switch to extend the developing circuit. Its purpose is to insure collision avoidance. The first byte in the information field is the S value of the address of the station of origin for the connected circuit. The second and third bytes are the L and N values of that station's address. The fourth byte is interpreted in accordance with the following binary representation:

```
 1 2 3 4 5 6 7 8
 |  |    |    |
 |  |    |    └--->| Overlay number (0 through 7)
 |  |    └----->| Secondary station gate number
 |  ------>| Primary Station gate number
 └-->| 0 = T circuit, 1 = S circuit
```

The RC may be used as a response to indicate that the secondary station has reserved the desired circuit and that the primary station may make the desired connection. The UC response is used if the connection cannot be made.

12. <u>Unable</u> <u>to</u> <u>Comply</u> (<u>UC</u>)

The UC response is a nonreserved response having the bit pattern 1101(P/F)001. The frame includes a one byte information field.

The UC response is issued when a secondary station is unable to comply with a request issued by the primary station. The reason is indicated by a one in any bit position in the single byte information field in accordance with the following coding:

```
 1 2 3 4 5 6 7 8
 | | | | | | | |
 | | | | | | | └->| Requested circuit is busy
 | | | | | | └->| Requested circuit is out
 | | | | | └->| Originator is not on circuit
 | | | | └->| Reserved
 | | | └->| I have other priority traffic
 | | └->| I am unable to act as primary station
 | └->| Circuit cannot be completed - task aborted
 └->| Reason not specified
```

In addition to these two new commands and responses, one of the original commands, XID, is used as a call progress signal. When a gateway has actuated a crossbar switch to extend a circuit, it transmits an XID frame back on the return channel of the switched circuit to the originator of the connection request. This frame resets the timers on the busy detectors as well as indicating the location to which the connection has progressed. In this application, the P/F bit is always set to zero since no response is expected or desired. The first octet of the four in the information field has the bit pattern 10000001. The second, third and fourth are the S, L, and N values of the address of the gateway actuating the switch to extend the connection.

Illustration of Connection Establishment

The connection establishment process described above is illustrated in more detail in the following diagrams. These reflect both processes and decision functions in flow chart form as well as the related GRIDNET command and response frames.

In these diagrams, the numbers on the left are references to the action, event or message that is horizontally aligned with them. For example, item 5 is a "Wait" function for Station 5 of Loop A. These numbers are used in the narrative explanation of the sequence of events. Ordinarily, the sequence of events advances with increasing number values.

GRIDNET command and response frames are represented by arrows. The tail of the arrow indicates the originator, and the point the recipient of the frame. The type of command or response is entered at the break in the arrow together with the setting of the Poll/Final bit which is shown enclosed in parentheses.

In these diagrams, no effort has been made to illustrate the full capabilities of the communications control protocol. For example, all frames are assumed to have been correctly received so there is no example of frame retransmission for error control.

The connection establishment process illustrated in these diagrams starts out with the origination of a connection request at Station 5 of Loop A of Fig. 13 and following the path to Station 1 of Loop A, thence to Station 3 of Loop B. From there on, a generic, iterative loop is represented that will continue the circuit establishment process to any designated destination.

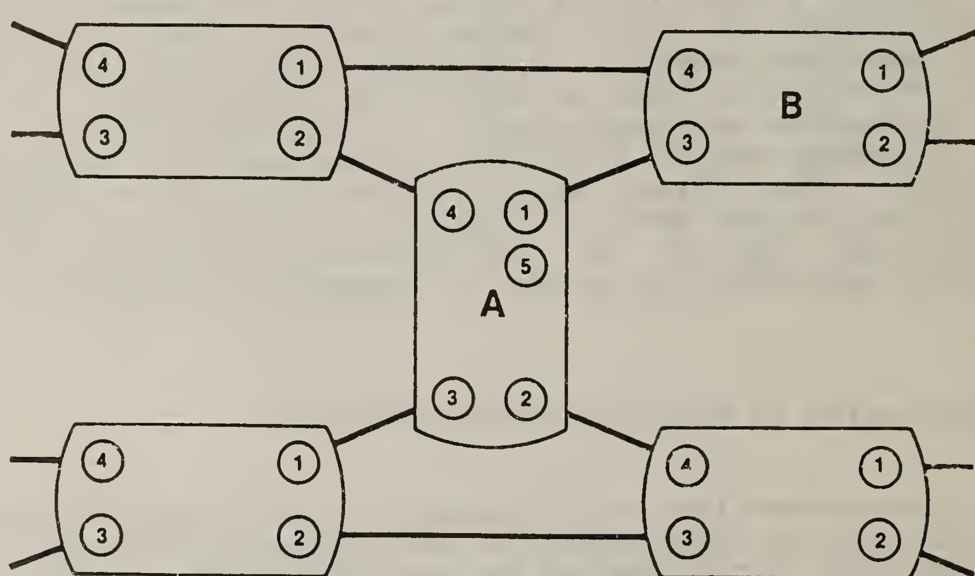

Figure 13. Example used in Flow Chart.

FLOW CHART - Sheet 1

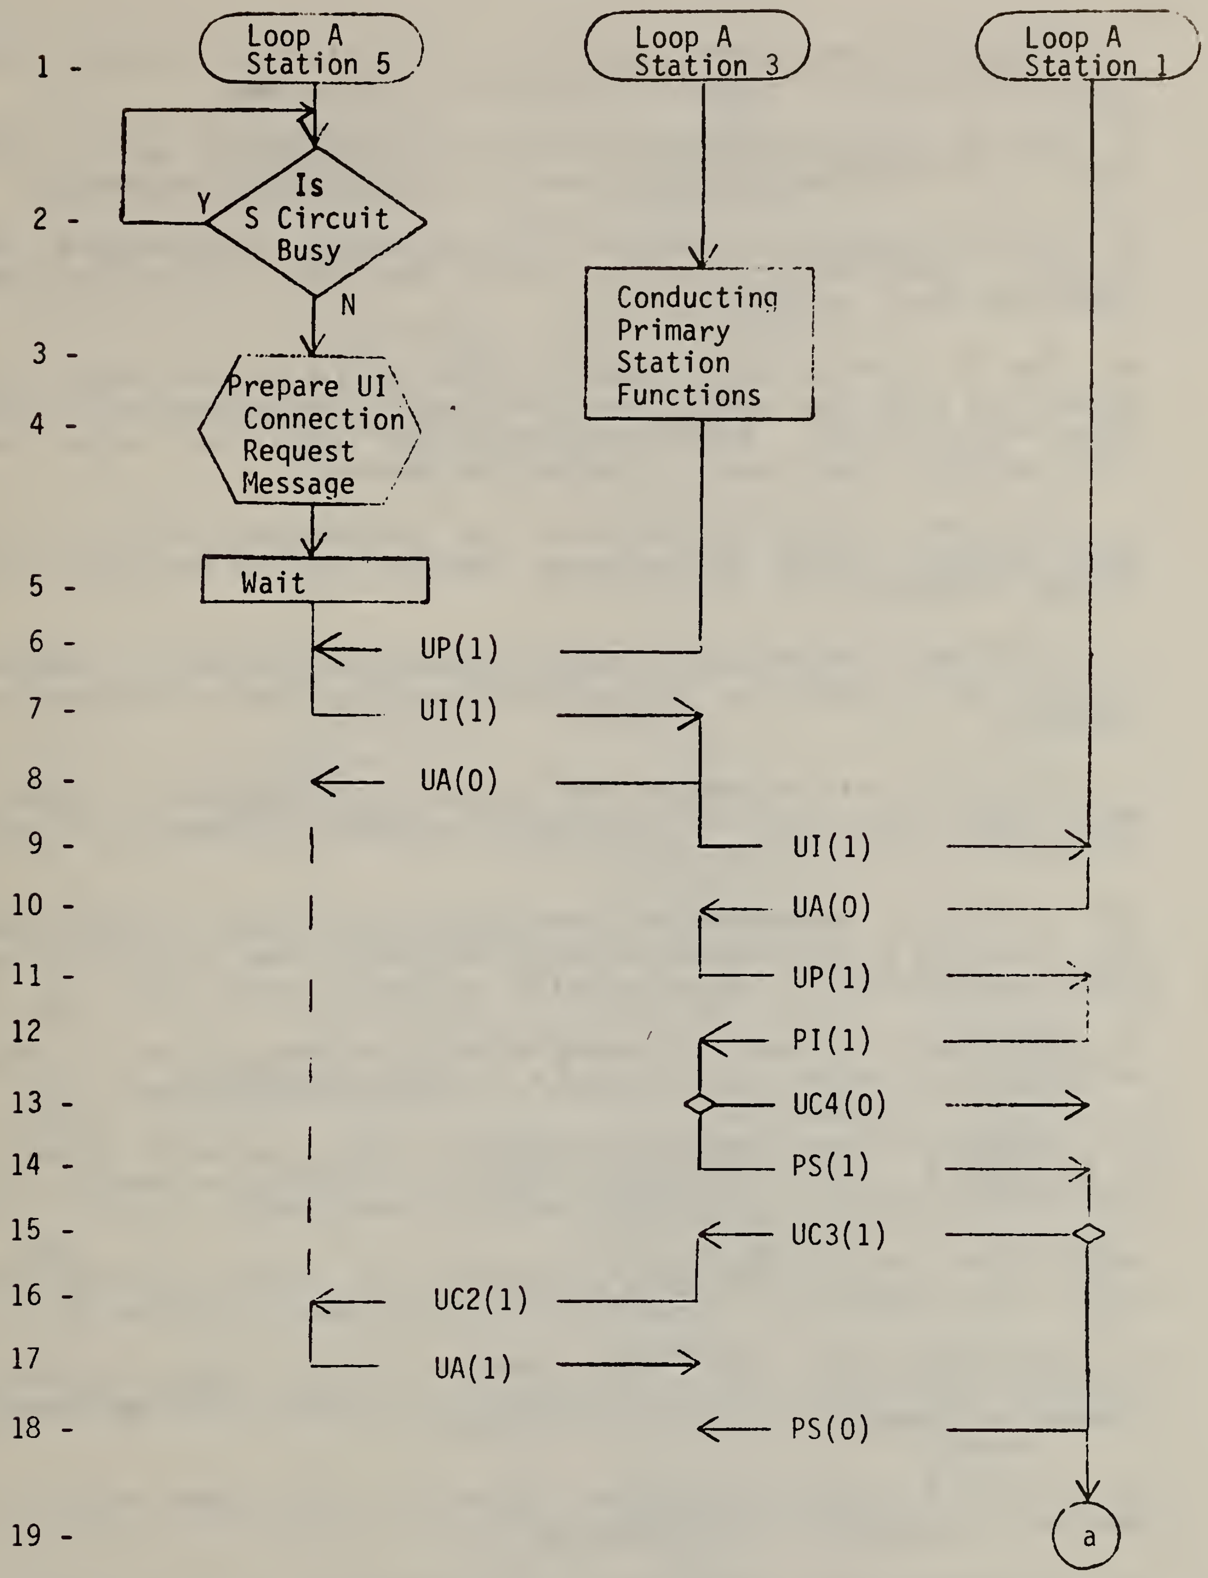

1. This line shows the identity of the stations that will be involved in the events depicted on this page.

2. The S circuit serving station 5 is tested to make sure that it is available before making a connection request.

3. In this example, station 3 is currently functioning as primary station, polling the other stations on the loop and handling their traffic.

4. The UI "connection request" frame contains the full address of both the source and the desired destination in its information field. Routing information will be added as this message progresses through the network.

5. Station 5 waits its turn to be polled by station 3. When the poll frame (6) is sent, the P/F bit is set to 1 inviting response.

7. Station 5 responds with its "connection request" frame. P/F set to 1 indicates it has nothing else that it wishes to transmit.

8. Station 3 acknowledges the "connection request" frame, and after examining its information field, transmits it (9) to station 1.

10. Station 1 acknowledges receipt of the frame using P/F bit set to 0 indicating that it has more that it wishes to send, and station 3 responds with a "poll" frame (11).

12. Station 1 requests primary status in order that it may process the "connection request" it has just received.

13. Station 3 must decide whether to grant primary status to station 1. If it has other waiting traffic to handle, it will respond with a UC4 and station 1 must wait.

14. Either immediately upon request, or after processing other work, station 3 offers primary status to station 1.

15. Station 1 must now decide whether it can accept primary status. If it cannot (for some undetermined and undefined reason) it responds with a UC3 frame.

16. If this happens, station 3 directs a UC2 frame back to station 5, which is acknowledged (17).

18. Alternatively, station 1 acknowledges acceptance of primary status. The P/F bit is set to zero indicating that no further communication is expected from station 3.

FLOW CHART - Sheet 2

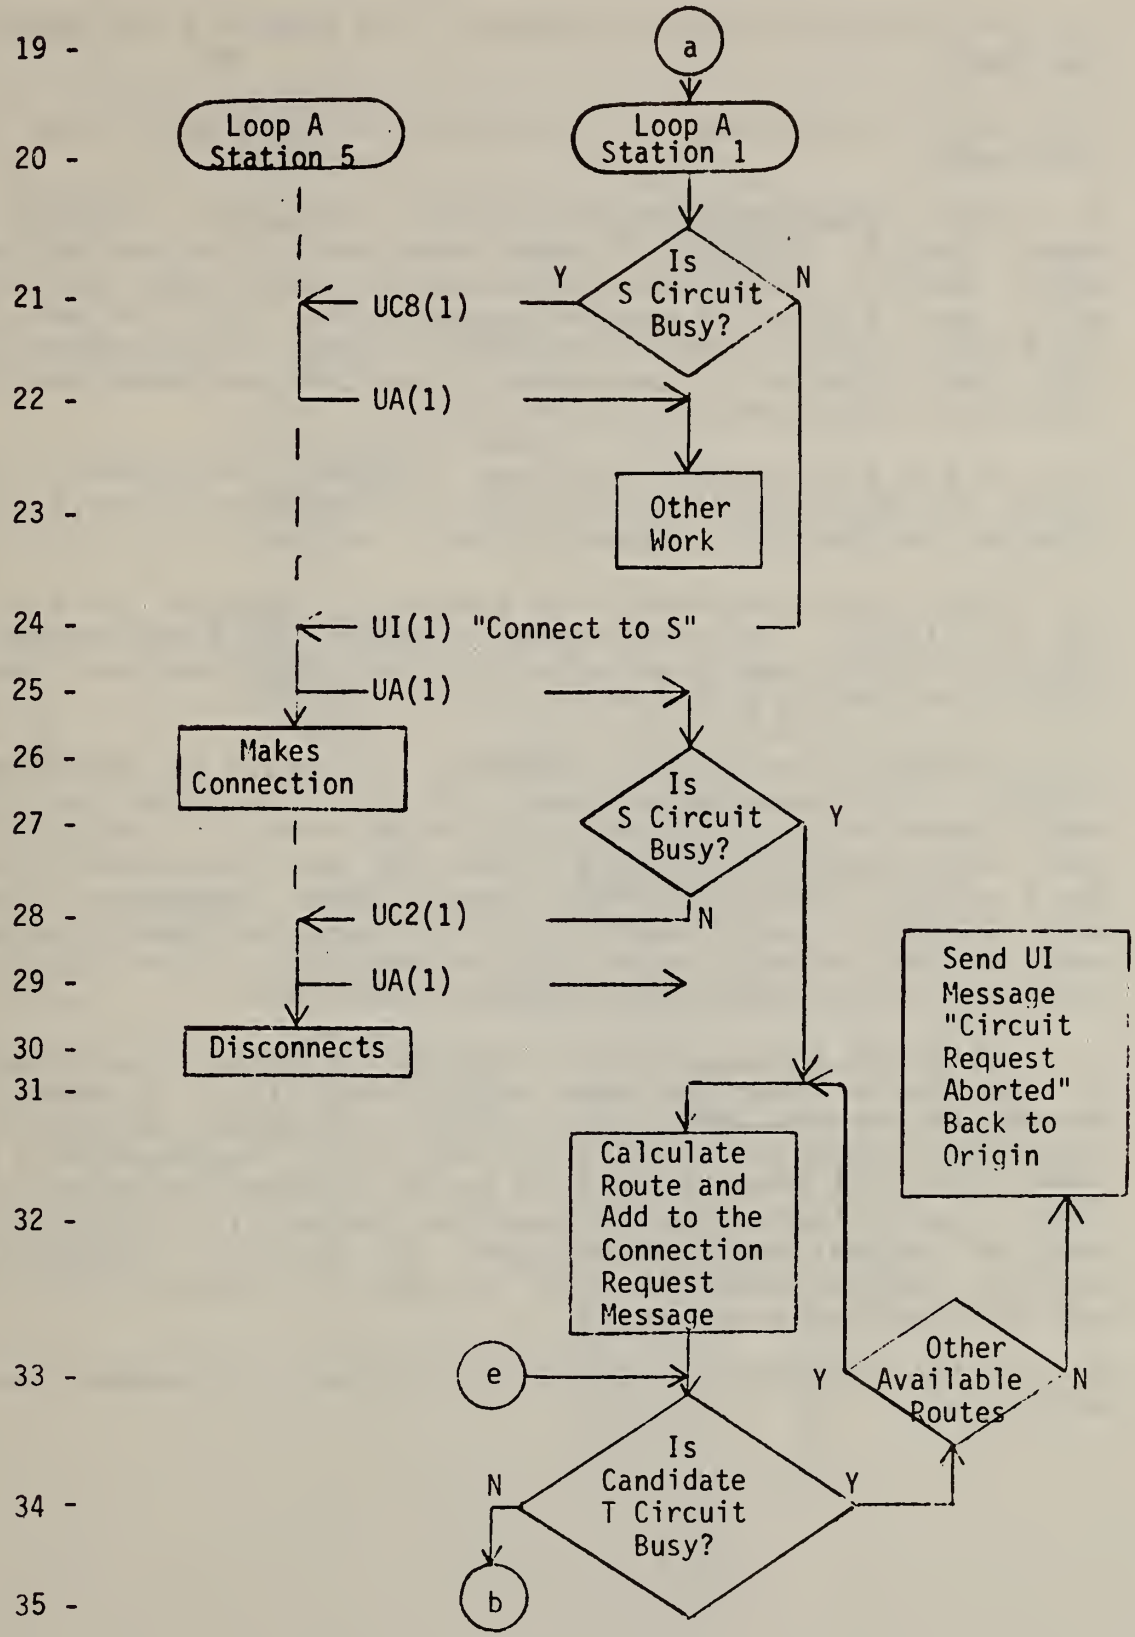

19. This is the connector between this page and the preceding page.

20. Only two stations are involved in the activities on this page.

21. After examining the "connection request", station 1 tests the S circuit to make sure that it is not already busy. Another request may have come in and been serviced while the station 5 request was enroute. If it is busy, a UC8 frame is dispatched to station 5, who acknowledges it (22) and station 1 proceeds to perform the other regular functions of a primary station (23).

24. If the S circuit is not busy, station 1 directs a UI frame to station 5 containing an information field that directs station 5 to connect to the S circuit.

25. This is acknowledged, and station 5 connects its DTE to the S circuit (26) and begins transmitting flags on the S circuit. These should actuate the busy detectors and reset their timers at all stations on the S circuit.

27. Station 1 tests the S circuit. It should be busy since station 5 is transmitting flags. If it is not busy there is some outage and a UC2 frame is directed back to station 5 (28) who responds with a UA frame (29) and disconnects from the S circuit (30). If the S circuit does (properly) test busy, station 1 calculates the route that the connection request message should follow to its destination, and appends this routing information to the message (32).

34. Station 1 then tests the candidate T circuit designated by the route that has been calculated. If it is busy, it determines whether there are other available alternative candidate routes (33). If there are none, it generates a UI frame that will be sent back to the originator (station 5 in this illustration) of the connection request (31) indicating that the request has been aborted. If there are other alternate routing possibilities, station 1 returns to step 32 and recalculates a new route.

35. If the candidate T circuit is not busy, proceed to connector "b".

FLOW CHART - Sheet 3

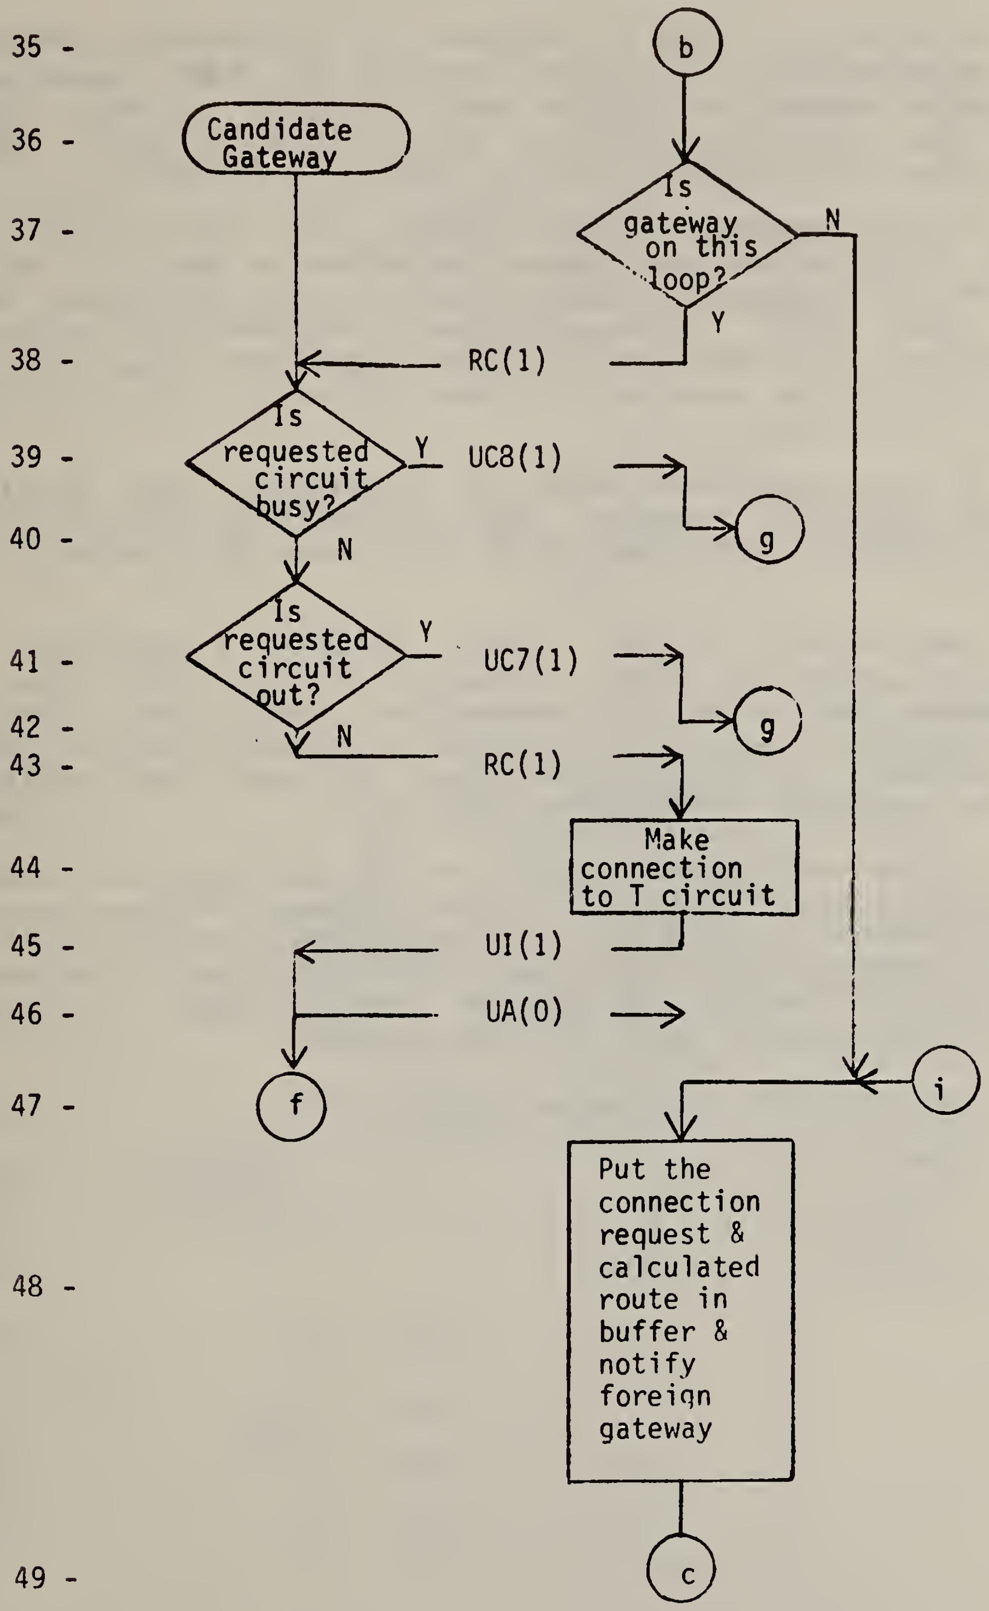

37. Station 1 of Loop A examines the calculated route to the destination for which a circuit has been requested to determine whether or not the next switching gateway is on the same loop as Station 1. If it is, an RC frame (38) is sent over the T circuit to that candidate gateway.

39. If the candidate gateway finds that the circuit for which reservation has been requested is busy, it replies with a UC8 frame, or if the circuit is out, the reply is a UC7 frame (41). Either of these responses will cause the generation of a "Connection Request Aborted" frame (connector "g") that is sent back to the originator - in this case, Station 5 of Loop A.

43. If the desired T circuit is available, the candidate gateway station reserves it to protect it from use by any other requester and sends an RC frame back to Station 1 of Loop A.

44. Upon reciept of this frame, Station 1 actuates the crossbar switch and makes the connection. It then sends the UI message (45) having an information field that contains the connection request and routing information to the candidate gateway. This is acknowledged (46) and the candidate gateway proceeds to connector "f", (47).

48. If the next gateway is not on the same loop as Station 1, then it must be that it is the adjacent foreign loop, so the connection request and routing information is put into their common shared buffer and the foreign loop station (in this example Station 3 of Loop B) is notified. Connector "i" also provides a route to this action.

49. The connector to the next page is "c".

FLOW CHART - Sheet 4

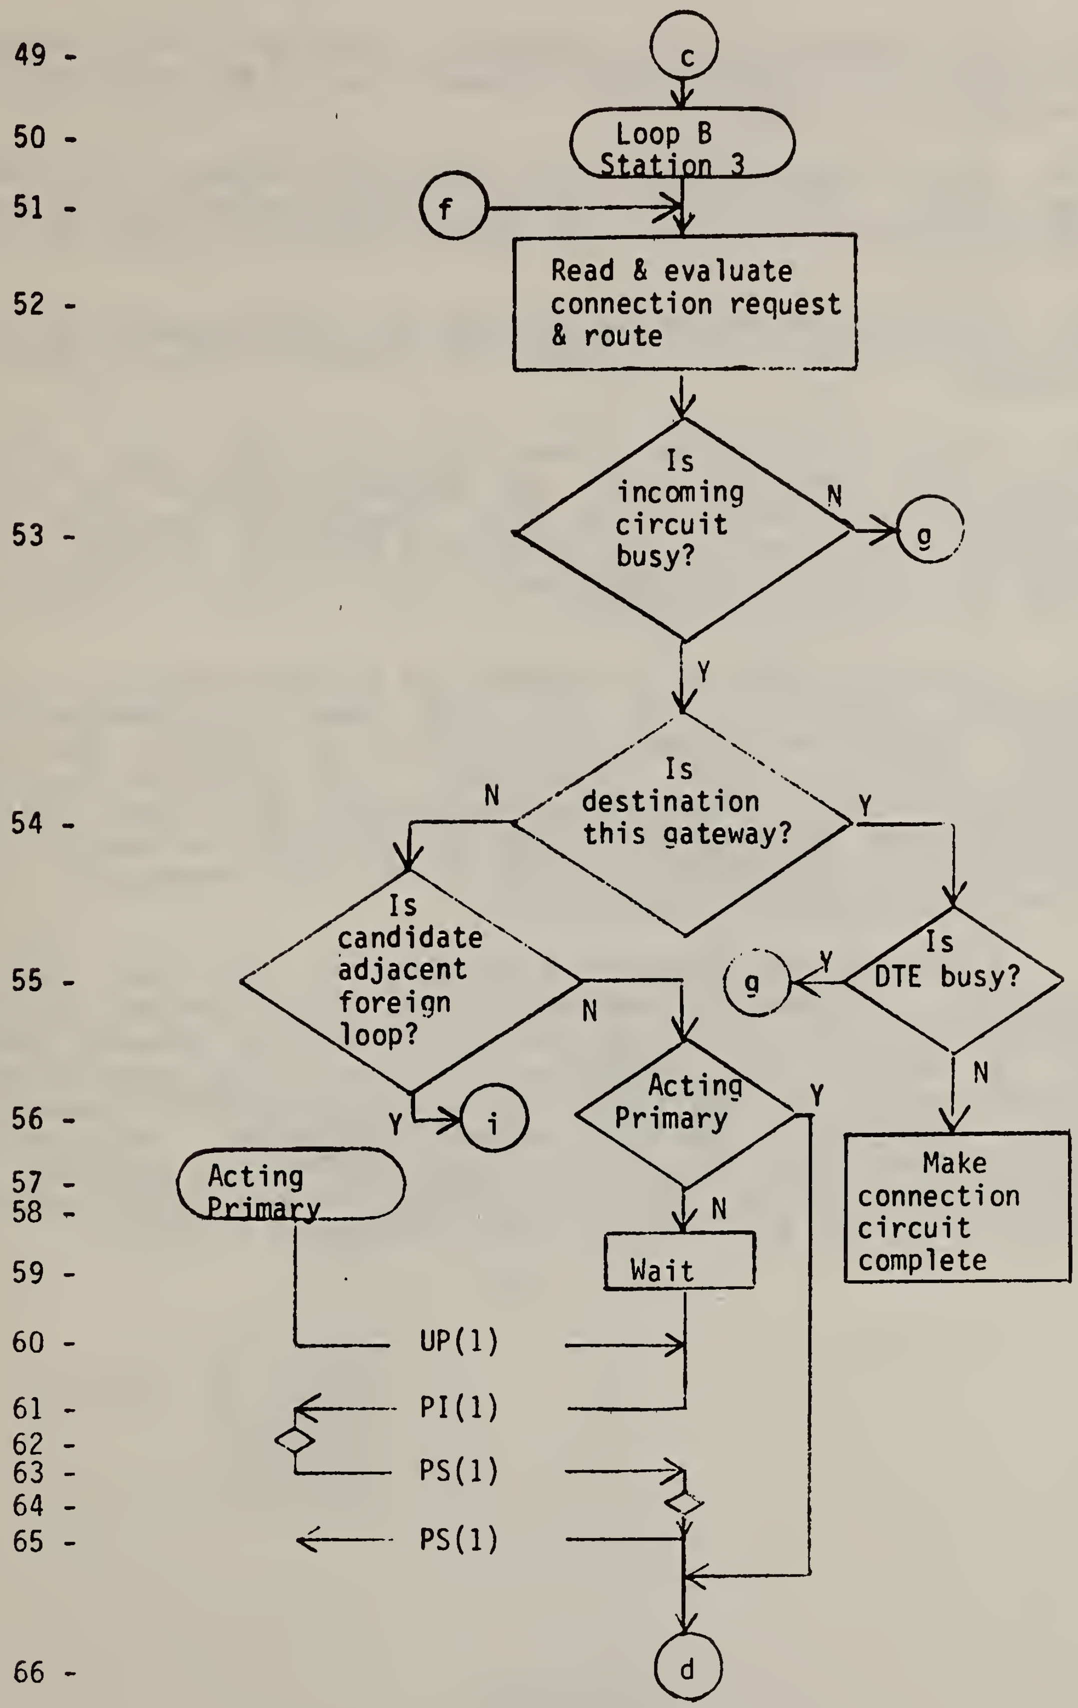

50. The identity of the performing station in this example, Station 3 of Loop B. Other gateways can enter this loop via connector "f", (51).

52. After examination of the connection request, the gateway tests the incoming circuit to make certain it is busy, and dispatches a "connection request aborted" frame if it is not, (53).

54. If the destination is the DTE at this gateway, and this DTE is not busy, the connection is made and the circuit is complete.

55. If the destination is not the local DTE, and routing directions indicate that the adjacent foreign loop is the preferred candidate for the next switching function, the process is directed to the preceding page via connector "i". This route is only followed when entry to this loop has been made via connector "f", (51).

56. If the gateway is currently the acting primary, it can proceed immediately to the next page via connector "d", (66). Otherwise it must wait (59) to be polled (60) and then request primary status (61). The decision functions involved in the requesting and acceptance of primary status are shown at (62) and (64), but the alternatives are not all illustrated, and only the granting (63) and acceptance (65) of primary status is indicated.

A gateway can complete a circuit to its own DTE or transfer a connection request to the common buffers for action on the part of its adjacent foreign loop gateway while it is acting as a secondary station, but it must acquire primary status in order to direct regular GRIDNET frames to other stations in the most expeditious manner. As a secondary it does not normally have knowledge of the identity of the acting primary, while as a primary it has full knowledge of the identity and status of all secondary stations.

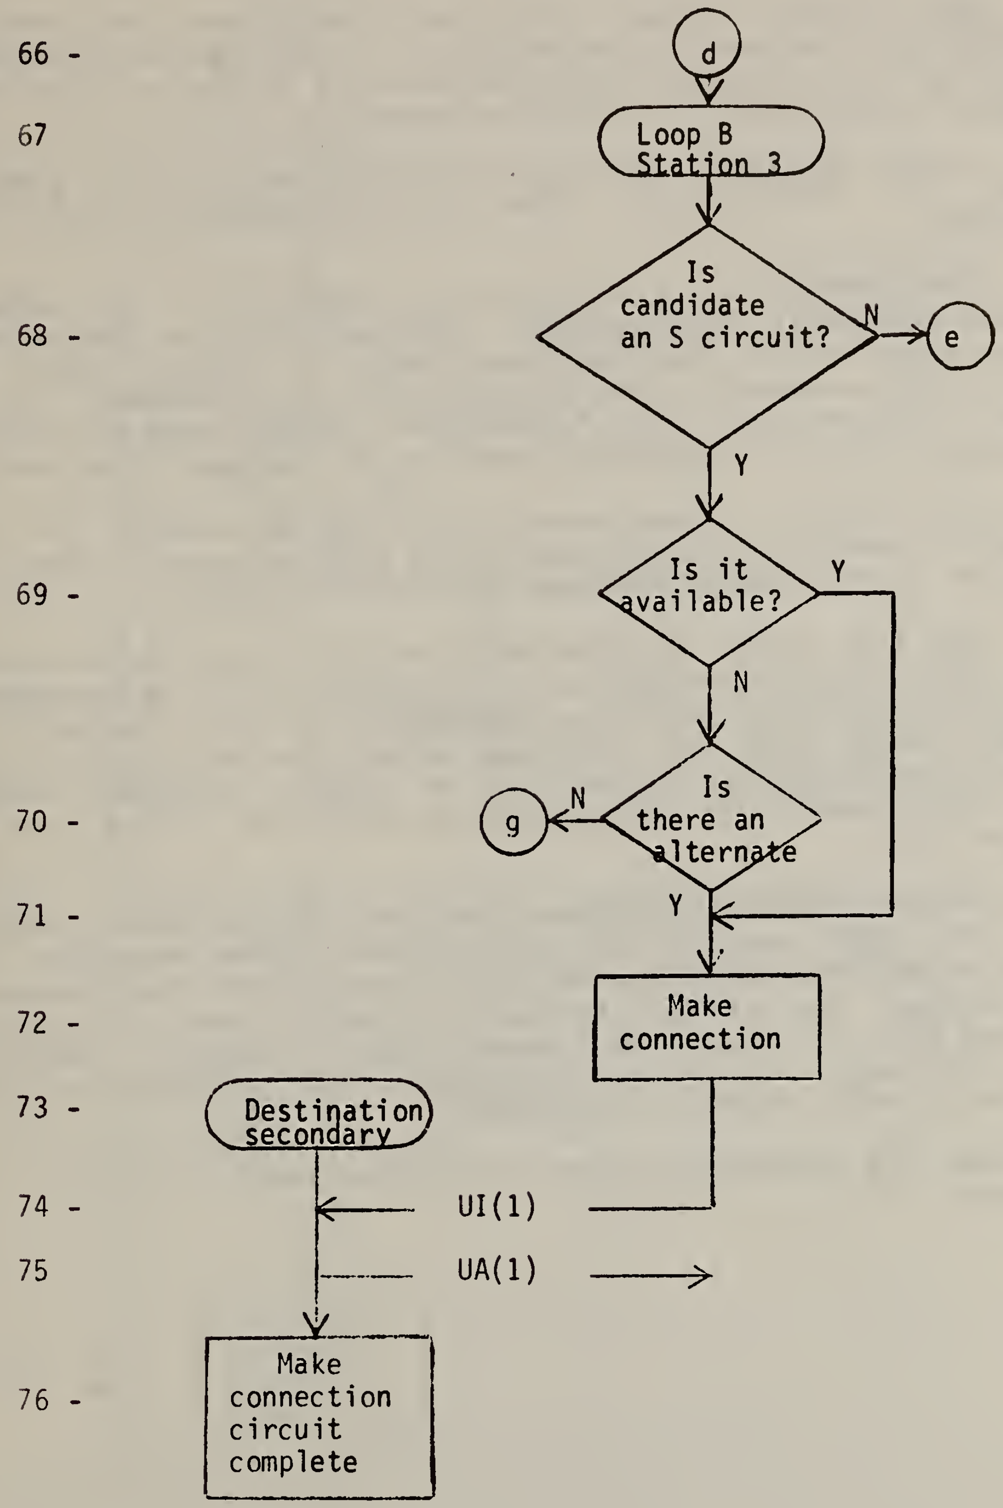
66 -
d
67
Loop B Station 3
68 -
Is candidate an S circuit?
N
e
Y
69 -
Is it available?
Y
N
70 -
g
N
Is there an alternate
71 -
Y
72 -
Make connection
73 -
Destination secondary
74 -
UI(1)
75
UA(1)
76 -
Make connection circuit complete

68. With primary status, the gateway examines the routing information to see if the candidate is an S circuit leading to one of the secondary stations that it serves on the loop. If not, it proceeds to (77) via connector "e".

69. If the candidate circuit is an S circuit, the gateway checks its availablity. That is, is it operable and is it non-busy.

70. If the preferred circuit is not available, an attempt is made to select an alternate. This can only be successful in instances where there are more than a single layer of basic supplementary channels overlaid on a GRIDNET. If no circuit is available, the "connection request aborted" frame is generated and transmitted back to the request originator.

72. If a circuit is available, the crossbar switch is actuated to make the connection, and a call progress signal is injected on the return channel to the originator.

74. Then a UI frame is dispatched to the destination secondary station instructing it to connect to the designated S circuit. Note that it is not necessary for the gateway to precede this with an RC frame since the utilization of S circuits is under the control of a single gateway rather than being shared by two gateway stations as in the case of T circuits.

75. The UA frame acknowledges that the destination secondary (73) has accepted the connection request message. If the selected secondary is unable to comply with the connection request, the P/F bit in the UA frame is set to zero indicating that it wishes to transmit more. Then in response to a following poll for that additional transmission, the secondary would respond with a UCl.

76. When the secondary makes its connection, the circuit is complete.

FLOW CHART - Sheet 6

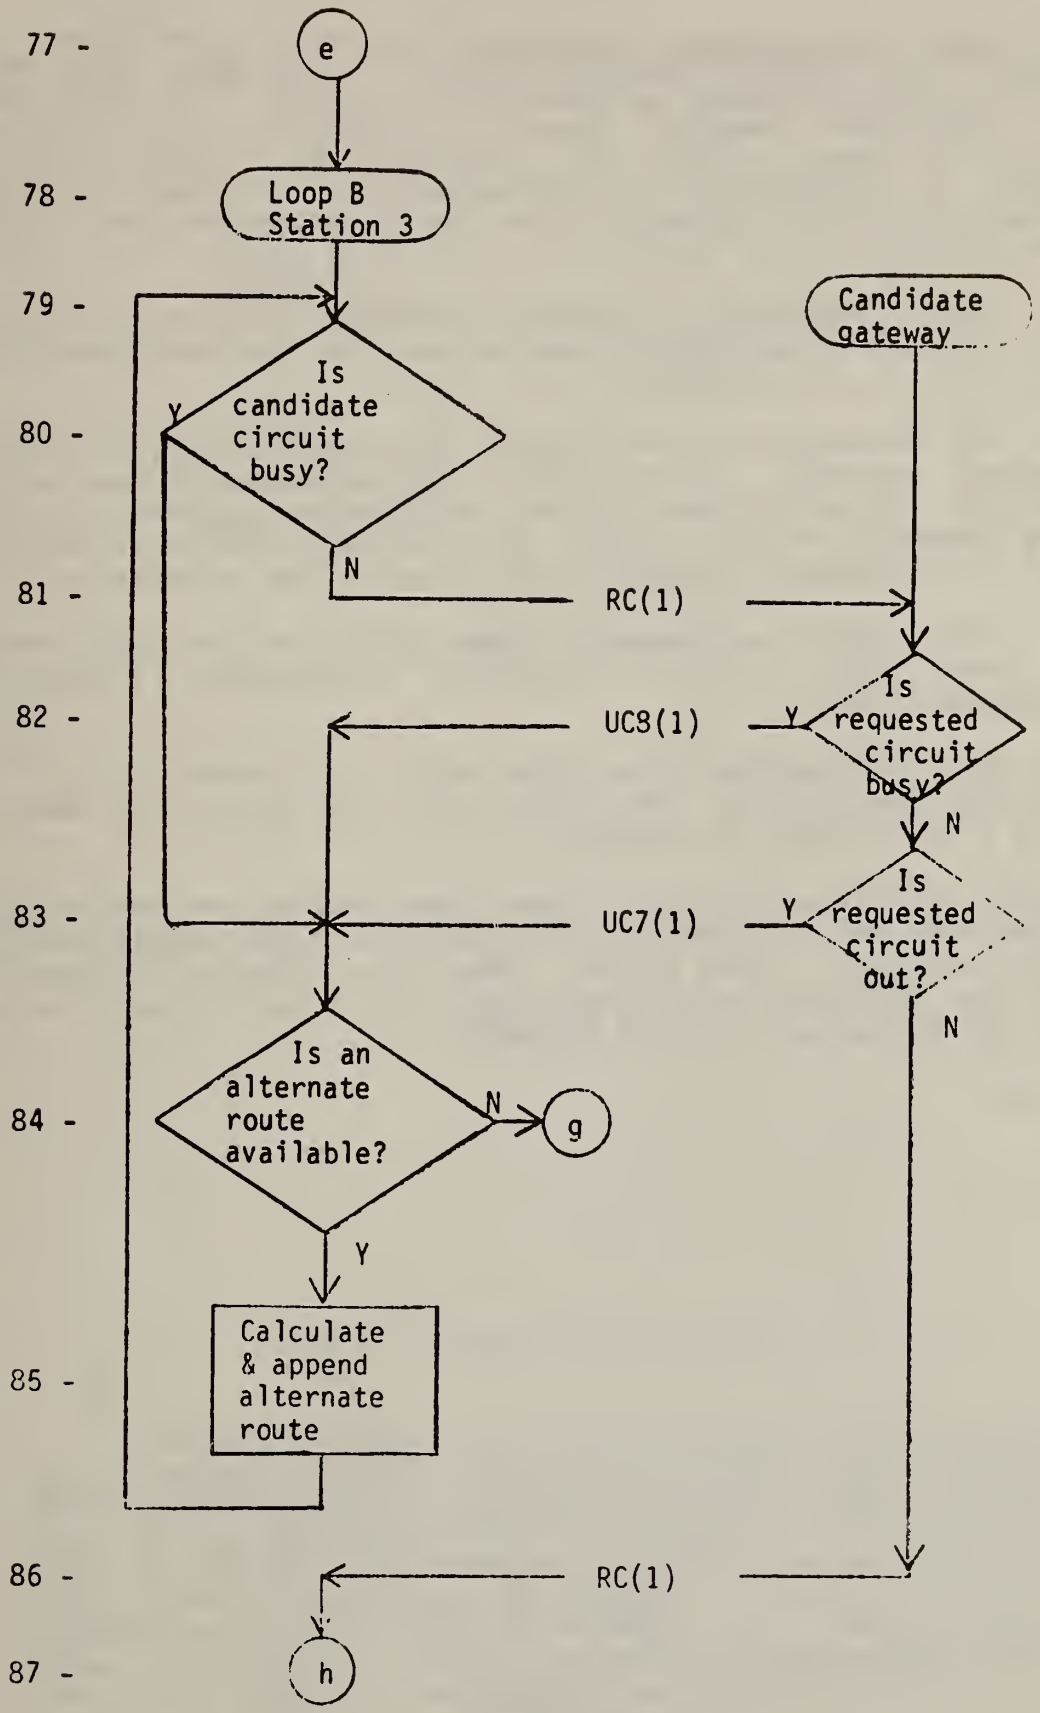

78. When the gateway (in this illustration, Station 3 of Loop B) has attained pimary status, it tests the busy condition of the candidate outgoing circuit (80). If it is busy, it determines whether there is a possible alternate route (84). If there is not, the "connection request aborted" frame is dispatched via connector "g". If there is an alternate route, it is calculated and appended to the connection request frame (85).

81. Having selected a non-busy candidate outgoing circuit, the reserve connection frame is sent to the candidate gateway, (79).

82. The candidate gateway tests the availability of the requested circuit. Unavailable responses are shown at (82) and (83) and acceptance of the reservation request at (86). Upon receipt of the latter, the crossbar switch is actuated (and call progress signals sent back to the originator).

91. The connection request frame with routing information is then sent to the candidate gateway and in turn acknowledged, (92).

94. This new gateway now loops back via connector "f" to (51) and the process continues until the circuit is completed or the attempt is aborted.

As previously noted, these flow charts are not intended to exhaustively illustrate all possible error conditions and their recovery. They merely show the typical procedures that are followed in establishing a switched circuit between two points using a Hybrid GRIDNET.

FLOW CHART - Sheet 7

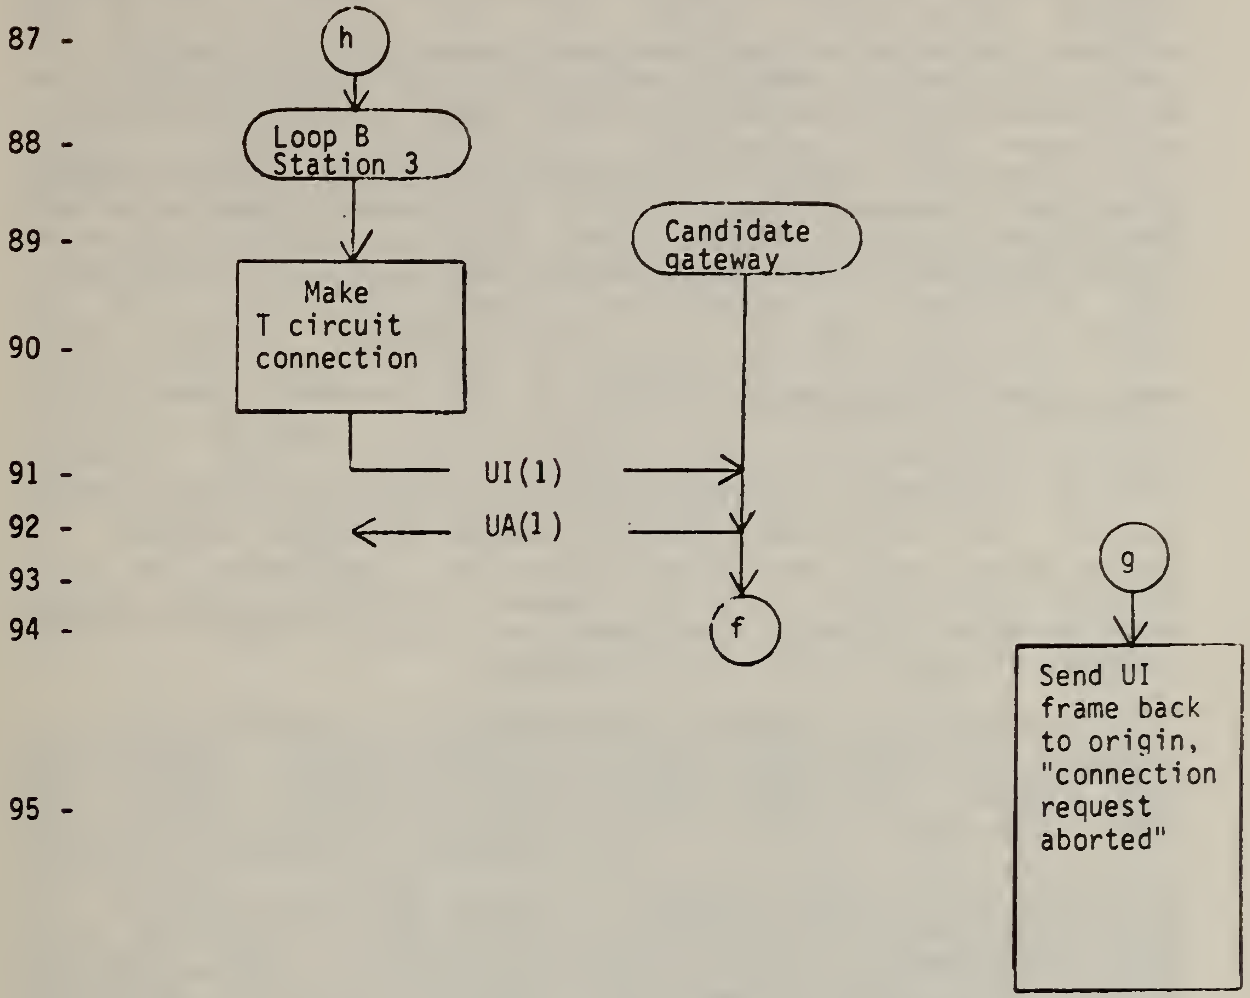

Hybrid GRIDNET Performance

Characterization of the performance of a Hybrid GRIDNET is a complex subject that is beyond the scope of this paper. Operating bit rate is a design parameter that depends upon the choice of media and the associated electro-optical components. Error performance is a function of signal to noise ratio and signalling technique. Availability may be estimated from the mean time between failures and the mean time to repair, which in turn, are dependent upon design features, component selection and operating stress as well as other factors. Connection establishment time is dependent upon the protocols and supervisory signalling used for this purpose as well as upon signalling rate, propagation delays, processing delays, and the possibility of sub-optimal route selection. Network capacity is a function of overlaid channel density and of the length and geographical distribuition of the circuits that coexist simultaneously.

Since the factors that affect connection establishment time and network capacity in the Hybrid GRIDNET are both highly complex and differ from those in other types of telecommunications networks, they will be briefly discused.

Connection Establishment Time

As a circuit is extended from its source to its destination, it may pass through a number of loops enroute. In those cases where the circuit does not originate or terminate in a given loop, but passes through it, the gateway station must exchange a minimum of either four or eight GRIDNET frames to make the connection.

Four of these frames are exchanged in order to acquire primary station status if the gateway does not already hold that status. In a loop that is internal to a network and full connected, there are four gateway stations. Thus, at any given time, there is a 0.75 probability that the gateway does not hold primary status. In the case of loops that are on the periphery of the network, there may be three or even only two gateway stations, so this probability is reduced to 0.66 or 0.50, respectively. In addition to the time required to exchange these four frames (lines 60 - 65 of the Flow Chart), the gateway may have to wait an indeterminate amount of time before its address is reached in the polling cycle. If the regular GRIDNET traffic is heavy, one or more frames containing information fields may have to be handled, and primary status may be even switched to another gateway before a poll response opportunity comes to the gateway that is handling the connection request.

When primary status has been acquired, a minimum of four frames must be exchanged in order to reserve the desired circuit (lines 81 and 86), and to forward the connection request message (lines 91 and 92) to the next gateway. If the desired circuit is not available, time is required to determine whether there is an alternate route available and to calculate the alternate path to the destination before another attempt can be made to reserve this new circuit. This is an iterative process that might be repeated again under heavy load conditions.

Estimating the amount of time that might be required for these action is not straightforward. In computer simulation of GRIDNET it was found that the per-loop delay of frames containing 8,000 bit information fields ranged from about 20 to 200 milliseconds, depending upon offered loads, in a network operating at one megabit/ second signalling rate, as measured in inner loops carrying the highest density of traffic. The knee of this delay curve was near the 50 ms. point, and this probably represented a reasonable maximum loading for the network. If the same value of delay, 50 ms. per loop, could be maintained for the connection establishment process, then a circuit spanning six loops and a distance of about 300 km might be set up in less than half a second. Such a circuit length might represent the longest directly routed circuit in a reasonably symmetric network containing a total of about thirty loops.

Network Capacity

The number of switched circuits that a Hybrid GRIDNET can support concurrently is highly dependent upon a number of factors. These include the number of loops in the network, the density of the overlaid channels, and the length and distribution of the circuits.

A full connected loop that is internal to a network is provided four S circuits and six T circuits from a single hybrid overlay. With optimum selection of sources and destinations, that loop could support ten simultaneous circuit connections. Without selection of sources and destinations, a more likely maximum utilization is five or six simultaneous circuits in the case of loops that are located near the center, and presumably most heavily loaded, portion of the network. Loops that are located near the edges of the network are expected to be more lightly loaded because they are not on the direct route between as many sources and destinations of switched circuits. In a 30 loop symmetrical network, it is relatively easy to lay out 16 paths that go from one edge of the network to the other without saturating any of the six loops in the central core. These paths range from four to seven loops each in length, and utilize from three to five T circuits in the core loops and from one to three circuits in the peripherally located loops. A more uniform load distribution could be obtained through the use of multiple hybrid basic overlays in those portions of the network where there was higher demand for direct routing.

Lacking this, load equalization will occur automatically to some extent as a result of the establishment of longer, more indirect, alternate routes when the preferred direct routes are busy.

In a symmetrical network, it would be expected that the maximum number of directly routed circuits that could be simultaneously supported would vary directly with the size of the network, and inversely with the path length when both are expressed in terms of numbers of loops. As the offered load increased beyond this point, it is expected that a few, longer, more roundabout, circuits would be established but that these would be accompanied by a higher probability of connection attempt abortions as a result of blockage.

The foregoing comments can provide only a general impression of the capacity of a Hybrid GRIDNET. More meaningful data could best be developed through computer simulation of a model of the proposed network that reflected the traffic patterns that it is expected to support.

CONCLUSIONS

A Hybrid GRIDNET concept has been presented in which additional data communications channels, switching facilities, and supervisory and control features have been overlaid on a GRIDNET. The result is a network that can provide both packet switched and circuit switched services. The packet switching service that is provided by the network provides a highly survivable communications capability that offers an opportunity for a high level of utilization of network facilities. This type of service results in end-to-end packet delay times that may vary with routing and network loading, however, and services such as digital voice and video cannot easily accomnodate variations in interpacket arrival times. They are most conveniently handled by end-to-end circuits in which the transmission delays are either fixed or varying at a predictable rate. A HYBRID GRIDNET provides this additional class of service. The GRIDNET advantages of distributed processing and of alternate routing based on near neighbor operability status information have been retained.

Using simple additions to the GRIDNET communications control protocol, switched connections are established as the result of the exchange of regular GRIDNET frames between the stations involved in the switching processes. The alternate routing algorithms are the same as those used in GRIDNET, so a circuit from source to destination is established if there is any operable path between these locations.

As a switched circuit is being established, candidate links are reserved in advance in such a way as to avoid any posibility of collision with another circuit that might be in the process of establishment concurrently. Call progress signals are issued to the calling station as the circuit is extended past each enroute switching point in the network. If the requested connection cannot be completed for any reason, the calling station is notified through the use of a regular GRIDNET packet.

The performance characteristics of a Hybrid GRIDNET are highly complex functions of many variables. Estimates of representative values for connection establishment time and for network circuit capacity are best developed as a result of computer simulation of a model that accurately reflects the proposed network configuration and offered loads.

REFERENCES AND BIBLIOGRAPHY

1. R. T. Moore, R. J. Carpenter, A. W. Holt, A. L. Koenig and R. B. J. Warnar, "Phase II Final Report Computerized Site Security Monitor and Response System", National Bureau of Standards, NBSIR 79-1725, PB 294 343, NTIS, Springfield, VA 22161.

2. R. T. Moore, "GRIDNET", National Bureau of Standards, NBSIR 80-2149, PB 81-144370, NTIS Springfield, VA 22161.

3. L. H. Brandenberg, B. Gopinath, and R. P. Kurshan, "On the Addressing Problem of Loop Switching", Bell System Technical Journal, 50:7, September 1972, pp. 1445-1469.

4. L. H. Brandenberg and B Gopinath, "A Table Look-up Approach to Loop Switching", Bell System Technical Journal, 51:9, November 1972, pp. 2095-2099.

5. D. G. Cantor and M. Gerla, "Optimal Routing in a Packet Switched Computer Network", IEEE Transactions on Computers, C-23, October 1974, pp. 1062-1069.

6. V. Cerf, "Internet Addressing and Naming in a Tactical Environment" Internet Group Documentation, IEN No. 110, August 1979.

7. D. Cohen and J. Postel, "Source Routing", IEN No. 95, USC- Information Sciences Institute, May 1979.

8. R. J. Cysper, Communication Architecture for Distributed Systems, Systems Programming Series, Addison Wesley, Inc., 1978.

9. D. J. Farber and J. J. Vittal, "Extendability Considerations in the Design of the Distributed Computer System (DCS)", National Telecommunications Conference (NTC 1973), Atlanta, November 1973, pp. 15E.1-15E.6.

10. D. J. Farber, "A Ring Network", Datamation, February 1975, pp. 44-46.

11. G. L. Fultz and L. Kleinrock, "Adaptive Routing Techniques for Store-and-forward Computer-communication Networks", Proc. 1971 Int. Conference on Communications, June 1971, Montreal, pp. 39.1-39.8.

12. R. G. Gallager, "Local Routing Algorithms and Protocols", Proc. IEEE 1976 Communications Conf., Philadelphia, June 1976, pp. 20-17-20-20.

13. R. L. Graham and H. O. Pollak, "On the Addressing Problem for Loop Switching", Bell System Technical Journal, 50:8, October 1971, pp. 2495-2519.

14. R. J. Jueneman and G. S. Kerr, Explicit Path Routing in Communications Networks", Proc. Third International Conf. on Computer Communications, Toronto, August 1976.

15. W. J. Kropfl, "An Experimental Data Block Switching System", Bell System Technical Journal, 51:6, July-Auguat 1972, pp. 1147-1165.

16. J. M. McQuillan, "Adaptive Routing Algorithms for Distributed Computer Networks", Report No. 2831, Bolt Beraneck and Newman, Cambridge, MA., May, 1974. Also available from NTIS, AD 781437.

17. J. M. McQuillan, I. Richer, E. C. Rosen, "An Overview of the New Routing Algorithm for the ARPANET", Proc. of the Sixth Data Communications Symposium, Pacific Grove, CA, November, 1979, pp. 63-68.

18. R. Perlman, "Internet Routing and the Network Partition Problem", Internet Group Documentation, IEN No. 120, October 1979.

19. J. R. Pierce, "Network for Block Switching of Data", Bell System Technical Journal, 51:6, July August, 1972, pp. 1133-1145.

20. J. Postel. "DoD Standard Internet Protocol" Internet Group Documentation, IEN No. 128, January 1980.

21. H. Rudin, "On Routing and Delta Routing: Techniques for Packet Switching Networks", Proc. IEEE 1975 Communications Conf., San Francisco, June 1975, pp. 41-20-41-24.

22. J. F. Shoch, "Internetwork Naming, Addressing, and Routing", IEEE Computer Society International Conf. (COMPCON Fall 1978).

23. C. A. Sunshine, "Source Routing in Computer Networks", Computer Communication Review of the ACM 7, No. 1, January 1977.

24. A. C. Yao, "On the Loop Switching Addressing Problem", SIAM Journal of Computing, 7:4, November 1978, pp. 512-513.

25. H. A. Graf, et al., "GRIDNET Simulation Final Report, Vol. 1, System Description and Results", General Electric Report No. 81HV008, 7 August 1981, PB82-116013, NTIS, Springfield, VA 22161.

26. R. T. Moore, "A Discussion of Gridnet Simulation Results", National Bureau of Standards, NBSIR 81-2414, PB82-142894, NTIS, Springfield, VA 22161.

27. Tse-yun Feng, "A Survey of Interconnection Networks", Computer, December 1981, pp. 12-27.

NBS-114A (REV. 2-80)

U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET (See instructions)	1. PUBLICATION OR REPORT NO. NBSIR 82-2588	2. Performing Organ. Report No.	3. Publication Date October 1982

4. TITLE AND SUBTITLE

HYBRID GRIDNET
Packet and Circuit Switching in a Single Network

5. AUTHOR(S)

R. T. Moore

6. PERFORMING ORGANIZATION (If joint or other than NBS, see instructions)

NATIONAL BUREAU OF STANDARDS
DEPARTMENT OF COMMERCE
WASHINGTON, D.C. 20234

7. Contract/Grant No.

8. Type of Report & Period Covered

9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (Street, City, State, ZIP)

Defense Nuclear Agency
Washington, D.C. 20205

10. SUPPLEMENTARY NOTES

☐ Document describes a computer program; SF-185, FIPS Software Summary, is attached.

11. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here)

GRIDNET is a packet switching network, composed of multiply connected dual loops, being developed for the Defense Nuclear Agency in order to provide highly survivable data communications among a large number of sites. This report describes a concept for overlaying such a network with additional channels and switching facilities that may be used to establish point-to-point circuits on a demand basis. Switched connections are established following the exchange of appropriate frames between the stations that are involved. These exchanges use the regular GRIDNET packet switching facilities to provide essential supervisory and control functions including collision avoidance by means of circuit reservation in advance of connection. Switched circuits are automatically disconnected whenever no traffic is observed on them for a designated interval of time.

Alternate routing of a switched circuit is provided by the alternate routing algorithms that guide GRIDNET packets. These algorithms do not distinguish between candidate links that are unavailable because they are busy or because they are inoperable. In either case, an alternate route to the destination is selected as long as one exists.

12. KEY WORDS (Six to twelve entries; alphabetical order; capitalize only proper names; and separate key words by semicolons)

Alternate routing; circuit switching; communications networks; distributed control; integrated switching; packet switching survivability.

13. AVAILABILITY

[X] Unlimited
☐ For Official Distribution. Do Not Release to NTIS
☐ Order From Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402.
[X] Order From National Technical Information Service (NTIS), Springfield, VA. 22161

14. NO. OF PRINTED PAGES

62

15. Price

$9.00