The components providing the means to transfer data from one end user to another within a system incorporating SNA comprise the transmission subsystem. This paper discusses the organization of the subsystem, its logical and physical aspects, and the components involved in its operation.

The transmission subsystem in Systems Network Architecture

by P. G. Cullum

Systems Network Architecture (SNA) is the set of structural concepts and operational protocols underlying IBM's Advanced Function for Communications and the products supporting it.[1-3] One of its goals is to insulate the application program and the human user of the communication system from the moment-to-moment operation of the communications facilities in the system. The end users can then concentrate on the area of their chief concern, the functions related to application processing. All functions required to transfer the data from one end user to another are supplied by the communication system components known collectively as the *transmission subsystem.* These components include data links, cluster controllers, communications controllers, components in individual terminals, and most of the functions in the host access method. The end user specifies what is to be sent and what are the parameters of transmission and is freed from further concern with the details of transmission. This paper examines the structure and the operation of the SNA transmission subsystem, particularly the routing of data through the network.

Logical view of an SNA network

An SNA network can be described in two ways: the physical facilities—the CPUs, terminals, controllers, teleprocessing lines, etc.—or the logical functions—data transmission, integrity checking, data mapping, etc. For a complete understanding of the network, both views are important.

Figure 1 presents the logical view of an SNA network; the physical view is discussed later in the paper. The network consists of two parts: the transmission subsystem and a number of network addressable units (NAUs) surrounding it. The NAUs are logical entities, representing the various ports through which end users may access the communications facilities. Some (referred to as primary NAUs) appear in the host computer; others (secondary NAUs) are logical functions in cluster controllers or terminals.

Figure 1 Logical view of the communication system

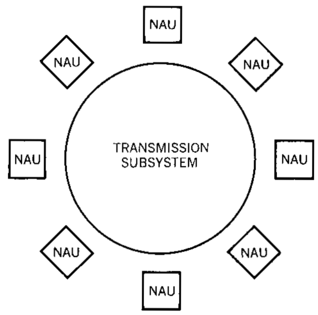

The functions and internal structure of NAUs are important factors in SNA, but a full explanation of them is outside the scope of this paper. Briefly, an NAU provides both data management to present information to the end user and protocol management to govern the rules of conversation between end users. For further information consult Reference 2. The NAUs and the transmission subsystem are collectively known as the communication system.

sessions

When one end user wants to communicate with another, the conversation, called a session, is started by using a control request to the communication system. A session is analogous to a human conversation. It typically consists of several back and forth exchanges covering, perhaps, a number of topics. Eventually, one of the end users ends the session by using another request to the communication system. As an example, a session might consist of an accounting clerk at a terminal communicating with the accounts receivable program in the host CPU. The clerk would initiate a session with the accounts receivable program, a number of account transactions would be processed, and the session would be terminated. The clerk might then establish a session with another application, perhaps accounts payable, which would proceed in a similar way.

The key features of sessions are:

- A session is a temporary logical connection between two, and only two, NAUs.

- Certain NAUs (those for application programs in the host CPU) may have simultaneous sessions with several other NAUs, supporting a number of end-user conversations at once. For example, the accounts receivable program may be simultaneously conversing with a number of accounting clerks. These conversations are independent, and the data transmitted in them is kept logically separate by the communication system. Of course, the application program may make any appropriate correlation or mixture of data it receives in its various sessions.

- In addition to its "user" session(s), each NAU has a session with the communication system itself (specifically with an NAU called the system services control point) for control operations such as initiating and terminating sessions.

- A session typically contains several user transactions.

- During a session, a logical path is established through the transmission subsystem between the two NAUs; no other NAU can intrude on the session by inserting data in this logical path (although a number of concurrent sessions may be multiplexed onto the same physical path).

types of NAUs

A description of the types of NAUs follows:

System services control point—The NAU responsible for the management of the communication system: network operator communication, configuration control, network startup, session initiation, system error recovery, etc.

Logical unit—The type of NAU serving as the port for application-oriented end users.

Physical unit—A physical unit exists in each node in the network. The physical unit's session with the system services control point is used for controlling the physical resources in the node and those directly attached to it. The physical unit of a node participates in the activation and deactivation of the data links attached to the node and the nodes attached to the links. It also may provide control functions for an operator located at the node. For example, the physical unit in a cluster controller at a branch office may allow the branch office manager to enable and disable various pieces of I/O equipment attached to the cluster controller. SNA does not specify the nature of the operator's functions. They are determined by the nature of the cluster controller providing them.

Layers of the transmission subsystem

As noted previously, one of the characteristics of SNA is the separation of the application-oriented functions from the transmission functions. The advantage of this separation increases with the introduction of a series of data links—rather than a single data link—between the application program and the terminal or cluster controller with which it wants to communicate, since management of these multiple links is a complex task. Further, the need for simultaneously sharing data links among a number of application programs makes it impossible for one application program to manage a link. Thus, link management and routing of data from one link to the next have been isolated from the application program and located in the transmission subsystem.

Two layers of functions are within the transmission subsystem—an outer layer (transmission control) with components uniquely in support of an individual session and an inner layer (path control and data link control) that is used in common by all sessions and that is, hence, referred to as the "common network" (see Figure 2).

outer layer

A number of functions related to transmission require unique support for each session, for example, session initiation and termination and data integrity. This support is supplied by a transmission control element (TCE in Figure 2). Each time an NAU enters into a session, it is connected to a transmission control element in its node. Thus, if an NAU has multiple sessions simultaneously, it has a transmission control element for each session. Located in a communications controller between the two end points of the session, a third transmission control element serves as a conversion mechanism (called the boundary function) enabling the two NAUs to work together. There are three principal facets of this mechanism:

Figure 2 The transmission subsystem

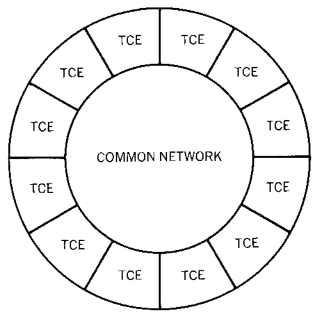

1. The secondary NAU is typically not able to process data as quickly as the primary NAU can send it. For this reason, a process called "pacing" is used in which data is temporarily held at the transmission control element in the communications controller until the secondary NAU can accept it. In conjunction with this, there is also pacing between the host transmission control element and the communications controller transmission control element, so that the latter is not overrun with information to be held until it can be forwarded to the destination.
2. The formats of control information used by terminals and cluster controllers are different from those used by host nodes and communications controller nodes. The transmission control element in the communications controller translates between these different formats.
3. In order to carry out these responsibilities, this transmission control must process the system messages controlling the session, e.g., the requests to start and end the sessions.

transmission control element

Figure 3 shows the detailed structure of the transmission control element in relation to the common network and (in the case of the two end point transmission control elements) to the NAU. The control element consists of three parts: connection point manager, session control, and network control. The "main-line" section is the connection point manager, whose function is the management of data transmission to and from the NAU. In each

Figure 3 Structure of the transmission control element and its relation to the NAU and the common network

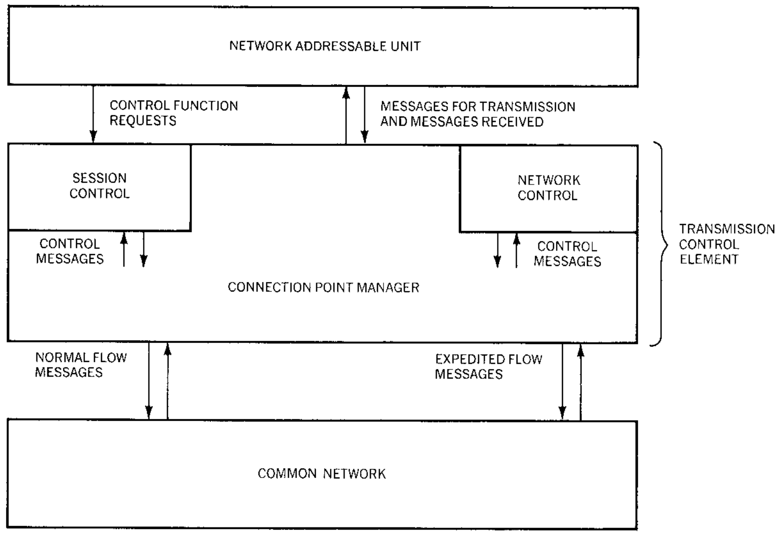

session are four data flows, two from the primary NAU and two from the secondary NAU. The end user's data is transmitted in the "normal flow" from its NAU. Control functions originating in the transmission subsystem on behalf of the session, as well as some of the control functions originating in the NAUs, flow on the "expedited flow" to avoid being detained by NAU protocols or by a queue of data from the end user awaiting transmission. The connection point managers for the session keep the two flows in each direction separate and give the expedited flow priority in their processing.

Each message that an NAU sends on the normal flow—all of the end-user messages and some of the control functions—is assigned a 16-bit sequence number by its connection point manager. (This number is unrelated to the synchronous data link control (SDLC)[4] send and receive sequence counts attached to information frames as they are placed on an SDLC data link.) When a message is sent, the connection point manager notifies its NAU of the sequence number assigned. When a message is received from the network, the connection point manager ensures that it is in sequence and that no prior messages were lost. Since some of the messages received from the network are for the NAU, whereas others are generated and processed within the transmission subsystem, the connection point manager has the additional function of routing incoming messages either to the NAU or to other components of the transmission control element.

A second component of the transmission control element is session control, which provides the session management functions for the NAU. These functions include session initiation and termination and session error recovery. Some of these functions are available to all NAUs; others are limited to primary NAUs. For an example of the latter, consider the session-initiation process.

When an NAU wants a session with another NAU, it sends a system message to the system services control point, requesting clearance for the session. Unless the system services control point finds impediments to the session (e.g., the other NAU is unknown or is already in session and cannot support another session), it notifies the primary NAU to start the session. The primary NAU uses session control functions to build a request (called the bind session request) that is the initial transmission sent in the session and that establishes the rules to be used to control the session. Once the bind session request is created, it is passed to the connection point manager and from it to the common network for transmittal to the session control element at the secondary NAU. The secondary session control uses the information in the request to establish the control facilities that will be used in the session. It also passes the information to the secondary NAU so that the required NAU protocols can be established.

The third component of a transmission control element is network control. It is used for the internal administrative communication of the communication system that is not directly related to the end user.

Physical view of an SNA network

the common network

The inner core in Figure 2, the common network, is itself composed of two strata: path control and data link control. To understand their function, consider the physical view of a network shown in Figure 4.

Physically, the network consists of a number of nodes interconnected by data links. These nodes have been divided into classes according to their function. In each network, there is a single host, which oversees the network and in which the majority of application programs reside. Attached to the host are communications controllers that perform data routing, multiplexing, and concentrating functions, and remote communications controllers may be attached to them. Also attached are various cluster controllers and terminals, both those for specific industry systems and those for use in applications in many situations. Physically, the collection of nodes has a tree structure.

The logical view of the network can be related to the physical view by noting that, in general, the NAUs are scattered around

Figure 4 Physical view of the communication system

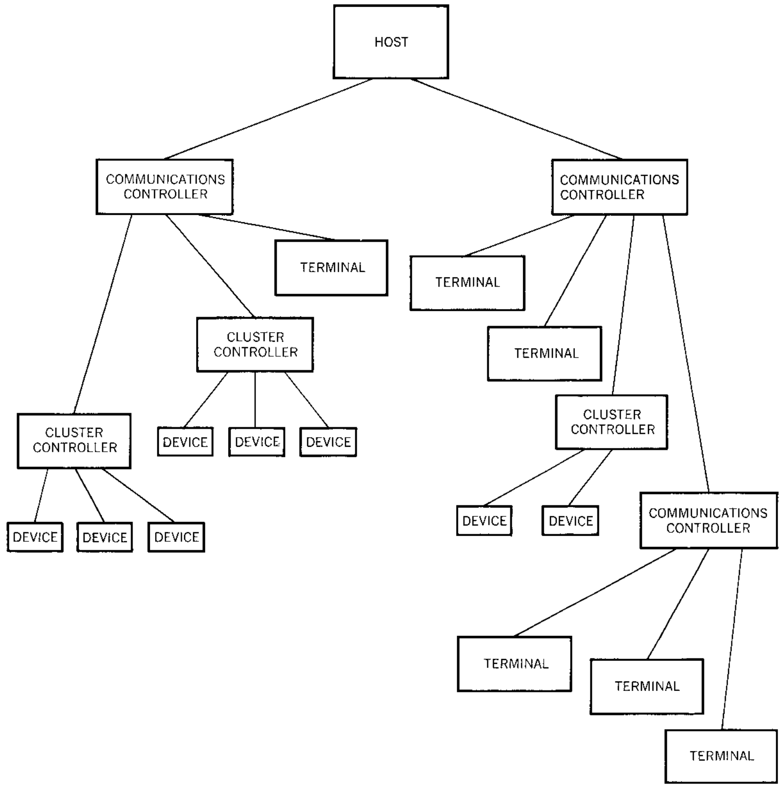

the periphery of the network in the host, cluster controllers, and terminals—the exceptions are the physical units, which appear in each node of the network.

To transmit information from one NAU to another, the proper sequence of data links and nodes must be selected, and the actual transmission of the message must be managed. These two functions are assigned to path control and data link control, respectively.

The responsibility of path control is the routing of data received from the various connection point managers through the network. This routing is accomplished in a stepwise fashion. Path control function is in each node in the network. For each message received, either from a connection point manager in the node or from another node, path control determines which node is the appropriate "next step" in the logical path and which data link should be used for transmission. The message is then transferred to data link control, whose function is to manage the movement of messages across the individual links from this node to the adjacent nodes. Data link control, like path control, is present in each node of the network. It manages the protocols

Figure 5 Progression of a request through the network showing the information present at each point

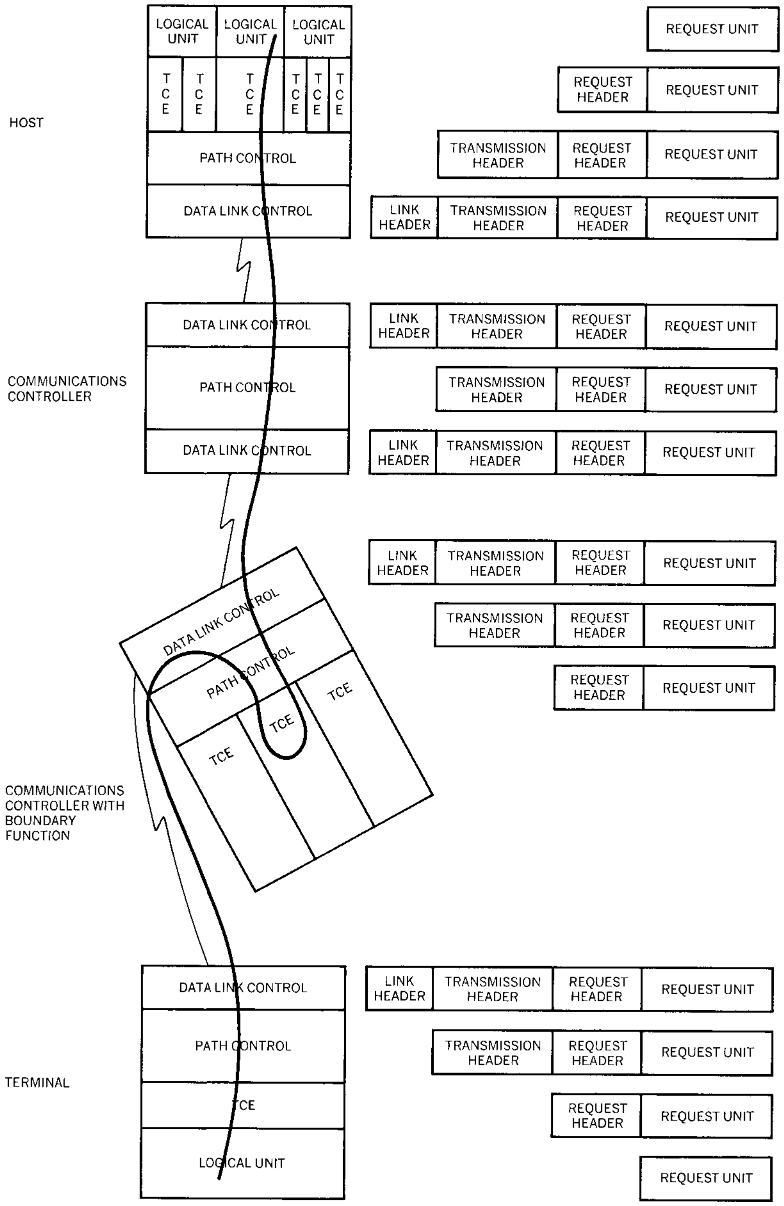

required to transmit information across the data link. For example, data link control for SDLC data links supplies the send and receive sequence counts and the frame check sequence for outgoing information frames and verifies this control information for the frames it receives. The data link control elements at the nodes attached to a data link work cooperatively to manage the data traffic on that data link.

Figure 5 shows the route a request unit travels from the time it is created by a logical unit until the destination logical unit receives it. At each step along the path, the applicable control information is shown.

Assignment and use of network addresses

When a specific network installation is created, each NAU in the network is assigned a unique network address by which it is known to the other NAUs and the transmission subsystem. NAUs connected via dial links are an exception. Their network addresses are assigned by the system services control point when the connection is made.

Figure 6 Network address format

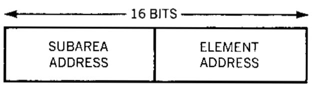

The entire network is partitioned into regions called subareas. Each subarea is assigned a subarea address. Each NAU is in one and only one subarea. Within a subarea, the NAUs are each assigned an element address. The network address, shown in Figure 6, reflects this two-part structure. The first field in the address specifies the subarea; the second specifies the element address within the subarea. When the installation first defines its network configuration, it determines how many of the 16 bits in the network address will be allocated to define the subarea and how many will be reserved for elements.

Figure 7 shows the same configuration that is shown in Figure 4, giving more detail, including the division of the network into subareas. Also shown are the logical units, physical units, and the system services control point, along with the network address that might be assigned to each.

Each subarea, except the ones directly controlled by the host (i.e., the subarea containing host application programs and the subareas containing channel-attached cluster controllers), contains a communications controller through which all the data flows into and out of the subarea. It disperses incoming data to the various cluster controllers and terminals in the subarea and gathers the data from them to be sent elsewhere in the network.

information units and formats

The communication between NAUs is of two types. First, one of the NAUs sends a request unit to the other. A request might be a line of data to be printed, a card image to be punched, an inquiry against a data base, etc. Second, when the receiver completes processing the request, it sends a response unit to report the outcome of the request. The response can be inhibited by the requesting NAU, which can ask that no response be returned or that a response be returned only if an error occurs during processing.

Figure 7 Sample configuration showing assignment of network addresses

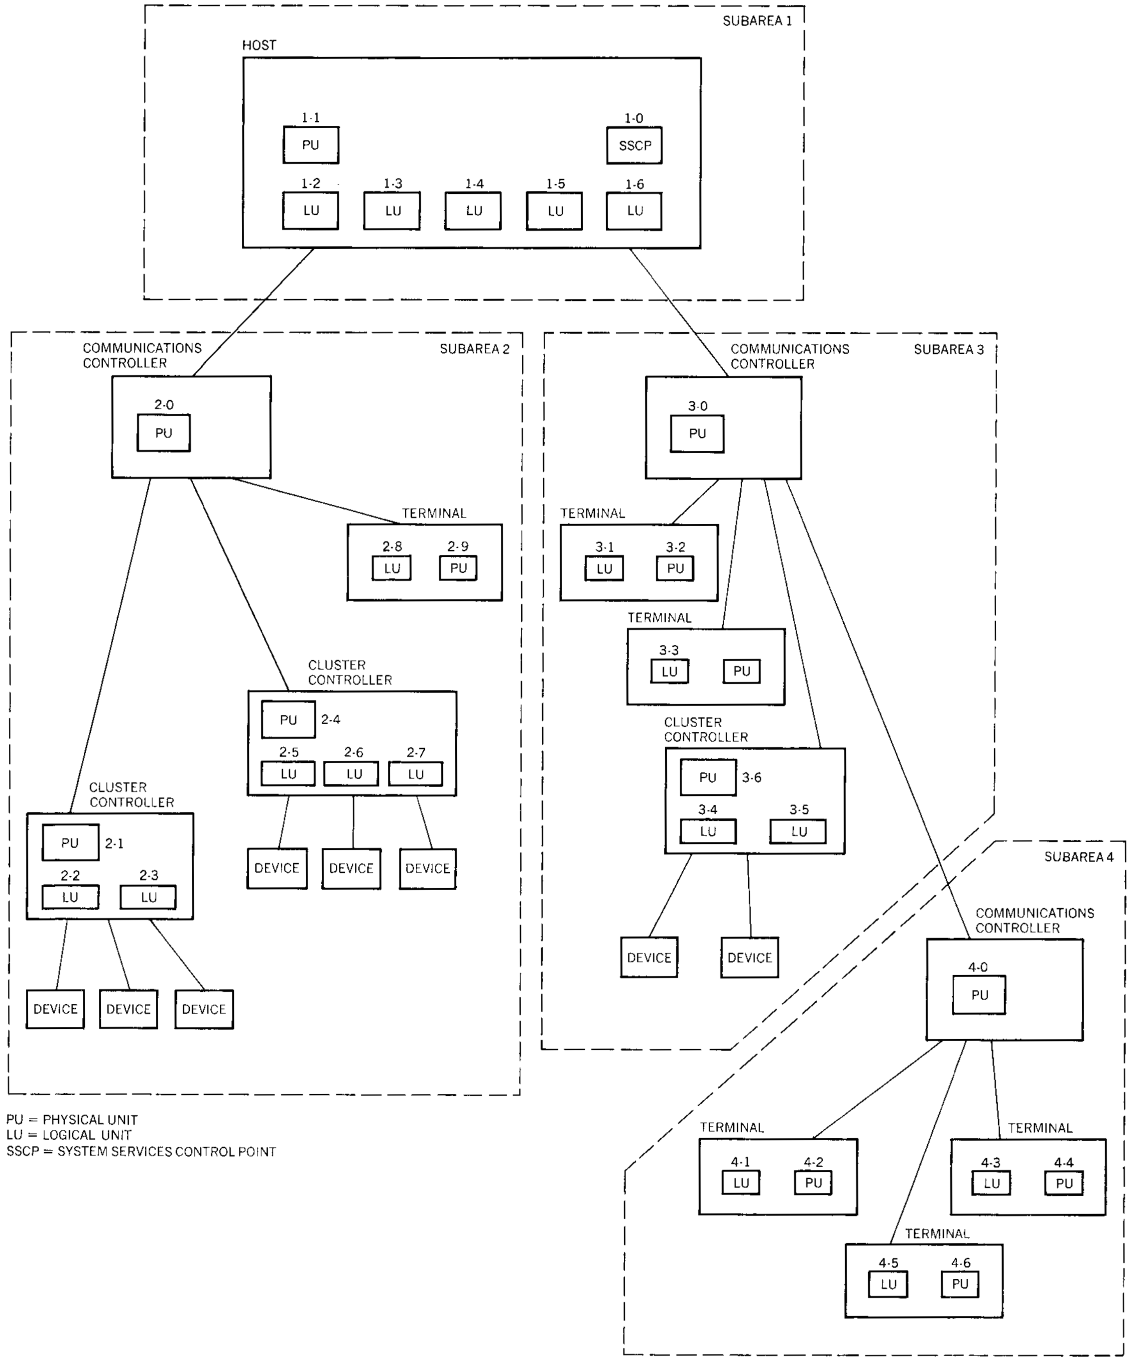

In front of each request or response unit that is sent, the connection point manager attaches a request or response header, containing various control information. For example, the header indicates whether the unit sent is a request or a response and, consequently, what format the header has.

There are four categories of requests: two (called function management data and data flow control) processed by the NAU and two (called session control and network control) processed within the transmission subsystem. The category of the unit sent is indicated by the header.

The form of response to be used (no response, respond only on error, definitely respond) is indicated in the request header. The response header indicates whether an error occurred. If an error occurs, the response will contain four bytes of sense information. A pacing indicator in both the request header and the response header is used to regulate the rate of information flow into the secondary NAU so that its processing capacity is not exceeded.

Table 1 Routing table

Destination subarea	*Outgoing link*	*Outgoing station address*
1	1	1
2	2	1
3	2	2
4	3	1
5	3	2
6	*	
7	4	2

In addition to the indicators used by the connection point manager, a number of indicators are carried in the header for the convenience of the NAUs. These indicators regulate some of the protocols the NAUs use in their communication; they are not inspected by the transmission subsystem.

When the connection point manager has attached the header to the unit sent, it passes them to path control for routing into the network. Path control attaches a header, called the transmission header, to each message it receives. The transmission header contains the information needed to route the message through the network, including the origin and destination addresses, the sequence number assigned, and the length of the message. In some cases, for transmission efficiency or to match buffer size limitations of the next node, path control may choose to divide the message into segments that it transmits separately. In this case, each segment is preceded by a transmission header. Indicators in the header allow a path control element receiving a message in segments to reconstruct the entire message.

In the host and in each communications controller, a routing table, unique to each, specifies how path control in this node should route the messages it receives from a connection point manager or from data link control. Table 1 shows an example of a routing table. It is organized by destination and indicates the "next step along the path" to which this path control should forward messages it receives for a given destination. The table entry for a given destination indicates the link on which the message should be transmitted and the station on the link which is to receive it. For the first nodes along the path that the message travels, only the subarea portion of the address is of interest, since all messages for nodes in the subarea take the same path until they arrive at the communications controller controlling that subarea. Thus, only the subarea portion of the destination address is used as a search argument for the routing table. At the communications controller in the destination subarea, the routing table has an indication (represented in Table 1 by an asterisk) that it is not to forward the message to another communications controller, but it is to look into a supplementary routing table and process the message further, based on the element address. Path control collects the message (if it had been segmented) and passes it to the session's transmission control element in the communications controller, which processes it as described above. One facet of this processing, the conversion of the control information format, requires further discussion.

local addressing and format conversion

Cluster controller nodes and terminal nodes have limited addressing capability. They use a simplified version of address called the "local address", so called because it is used and is meaningful only locally to the node and does not have a network-wide interpretation as does the network address. A format identifier present in each transmission header signifies which form of address is used (as well as distinguishing other differences) in this transmission header. The transformation between the network address and the local address is made by the communications controller to which the cluster controller or the terminal is attached.

The local address for NAUs in cluster controllers and terminals is one byte in size. There is no internal structure to the local address, and addresses may be simply assigned sequentially when the cluster controller or terminal is manufactured, or in the case of a programmable cluster controller, when the installation defines the logical units it contains. Use of the logical address allows the cluster controllers and terminals to be insensitive to the configuration changes most installations will undergo as communications requirements evolve. The same local address may be used to refer to many different NAUs. For example, if a particular application program exists in cluster controllers in each of several branch offices of a business concern, the application could have the same local address in each controller. Another example is the physical unit in the cluster controller or terminal that always has local address 0. Since the communications controller boundary function is aware of the physical location of each NAU, it uses this information to provide the uniqueness lacked by the local address.

When a session is established between an NAU in a cluster controller or terminal and one in the host, the communications controller dynamically assigns a local address representing the host NAU, which will be used throughout the session. In messages from the host NAU, the origin and destination network addresses in the transmission header are replaced with their local equivalents before being sent to the destination NAU. When a transmis-

Figure 8 Sample configuration showing assignment of network addresses and local addresses to NAUs in terminals and cluster controllers

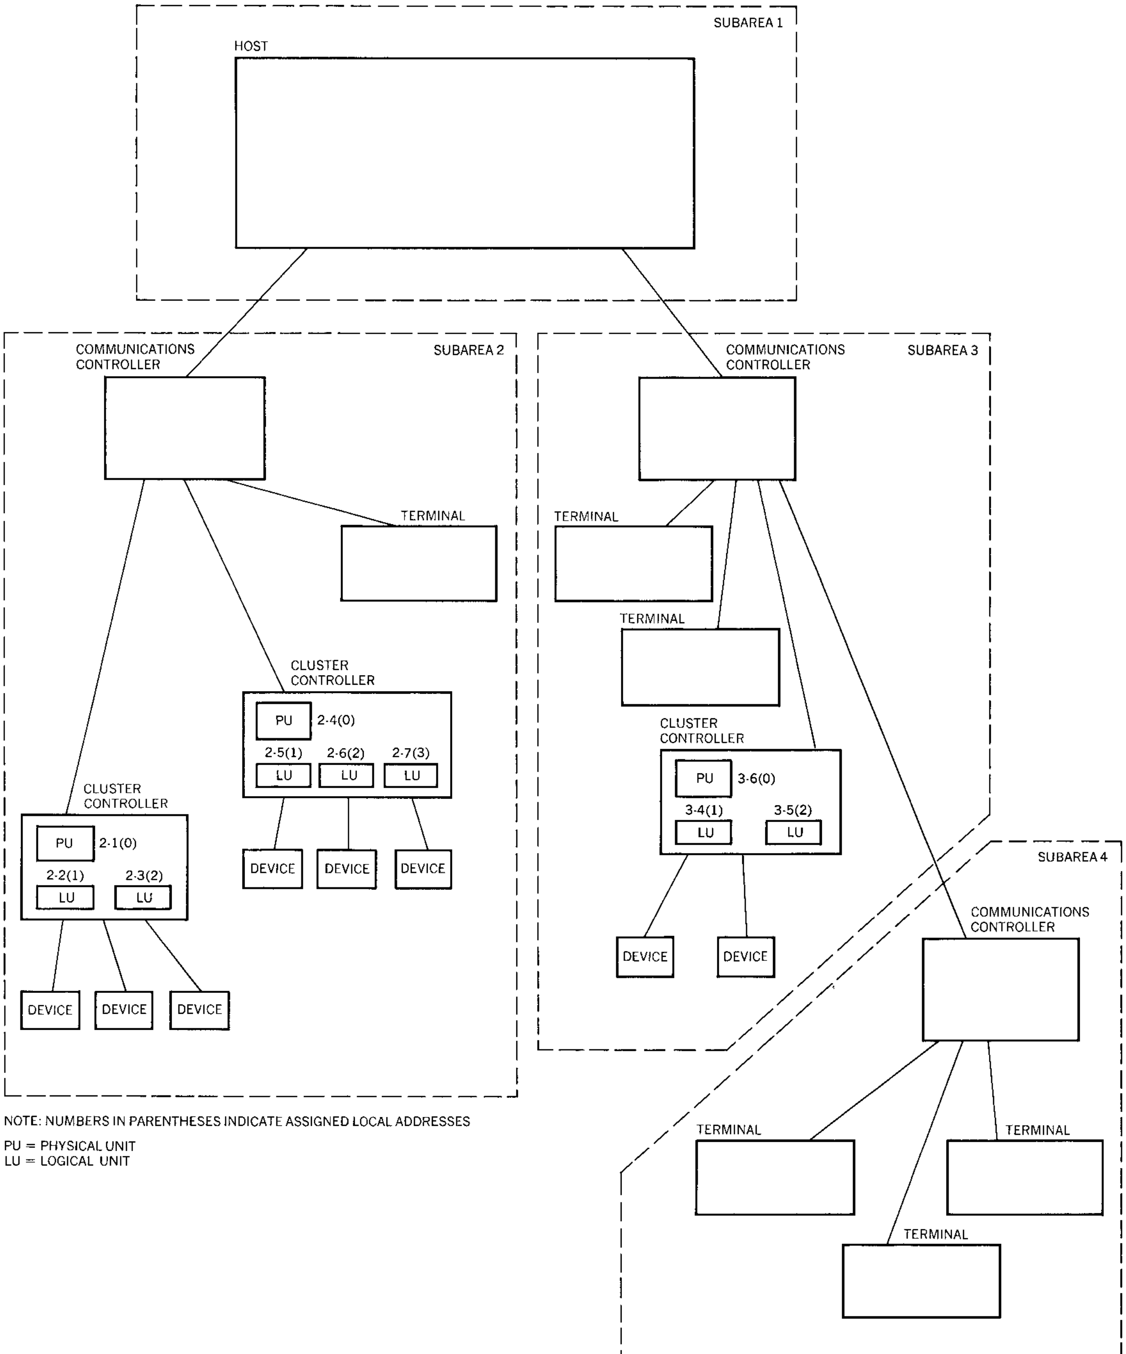

sion is received from a cluster controller or terminal, the communications controller uses the identity of the originating node (which it determines from the link and station address of the cluster controller) and the local addresses in the transmission header to determine the network addresses to be reinserted in the header. Figure 8 shows a sample configuration indicating the local addresses that might be assigned to NAUs in the cluster controller and terminal nodes.

very low-function terminals

Some very low-function terminals require even greater addressing simplicity. Local addressing for these terminals is somewhat different, although the translation required in the communications controller is analogous.

For these terminals, the two-byte origin network address and the two-byte destination network address are converted into a single one-byte value called the local session identifier. Figure 9 illustrates the identifier. Bit 1 of the identifier indicates whether the message is flowing in a session with the terminal's physical unit or in a session with a logical unit in the terminal. If the message is going to or from a logical unit in the terminal, bits 2–7 contain the local address of this logical unit. Since each logical unit in the terminal can have two concurrent sessions, one with the system services control point and one with another logical unit, bit 0 is used to indicate the session to which the message belongs. Thus, the identifier contains both the origin and the destination addresses for the message. For messages originating at the host NAU, bit 0 is the origin address and bits 1–7 form the destination address. For messages originating at the terminal, bit 0 is the destination address, and bits 1–7 form the origin address.

Figure 9 Local session identifier

0	1	2
SSCP/LU	PU/LU	LOCAL ADDRESS

SSCP = SYSTEM SERVICES CONTROL POINT
LU = LOGICAL UNIT
PU = PHYSICAL UNIT

Concluding remarks

In the data processing world of the future, communications requirements will become more complex and will involve many new application areas. These new applications will need to jointly use the system's communication resources (CPUs, teleprocessing facilities, and terminals) while the management of these resources will grow more complex.

SNA provides the management of the communication system so that the end user of the system can develop and use applications, ignoring the system's complexity and sharing its resources with the other end users. SNA also provides the system structure that allows both conceptual and physical growth of the network in the future.

CITED REFERENCES

1. J. H. McFadyen, "Systems Network Architecture: An overview," in this issue.
2. *Systems Network Architecture General Information*, No. GA27-3102, IBM Corporation, Data Processing Division, White Plains, New York.

3. *Systems Network Architecture Format and Protocol Reference Manual*, No. GA27-3112, IBM Corporation, Data Processing Division, White Plains, New York.
4. R. A. Donnan and J. R. Kersey, "Synchronous data link control: A perspective," *IBM Systems Journal* **13,** No. 2, 140–162 (1974).

Reprint Order No. G321-5024