GA27-3102-0
File No. S370-09

Systems

Systems Network Architecture General Information

IBM

Preface

This manual is designed for those persons having a basic understanding of data communications who need general information about Systems Network Architecture (SNA). It includes basic descriptions of the terminology, concepts, and scope of SNA. The information contained in this manual is current with the September, 1974, announcement of Advanced Function for Communications.

This manual does not provide instructions for implementing SNA, nor does it describe any specific equipment or programs that may implement SNA. For details of specific SNA implementations or information on SNA implementation subsets, the reader should refer to the appropriate publications, as available, that describe equipment or programs to be incorporated in an IBM SNA communications system.

Specific IBM equipment and programs which are consistent with SNA are identified in *Advanced Function for Communications System Summary*, GA27-3099. See notice (below) for the availability and current level of the applicable manuals.

First Edition (January 1975)
Changes are periodically made to the information herein; before using this publication in connection with the operation of IBM systems or equipment, refer to the *IBM System/360—System/370 Bibliography* (Order No. GA22-6822) for the editions that are applicable and current.

Requests for copies of IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

This manual has been prepared by the IBM Systems Development Division, Advanced Systems Architecture, Dept. E97, P.O. Box 12195, Research Triangle Park, North Carolina 27709 and published by department E01. A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be sent to the above address.

Contents

Figures

Chapter 1. Introduction

Purpose

This manual presents an overview of Systems Network Architecture (SNA). This architecture (SNA) brings together multiple products in a unified communication system design upon which new teleprocessing applications can be planned and implemented; SNA defines both the functions and the functional structure for IBM communication products. This manual discusses the general requirements and objectives of SNA to outline the scope of the architecture. The basic concepts and terminology introduced here are needed to relate the architecture to specific IBM communication products. However, details of physical packaging or structure of individual product implementations are not discussed here; where applicable, consult the specific product publication of interest for such information.

Communication System Requirements and SNA

A total communication system solution is necessary to provide the foundation for advanced applications, as well as to provide a base for IBM communication system product and programming development. This solution must provide a growth-oriented environment that minimizes the effort required to install new applications and to maintain or extend network configurations.

Systems Network Architecture (SNA) is an overall system solution. SNA formally defines the functional responsibilities of communication system components. In an SNA structure, all nodes (linked elements) adhere to these definitions. The scope of SNA definitions ranges from bit-level message header formats to the protocol of message sequences and to the classification of network nodes according to functional capability.

SNA relieves the communication system user of many network control and network resource management requirements, allowing him to concentrate on application functions. The IBM support allows resources to be shared across a wide range of user applications.

Advances in technology have made it possible to design communication products that perform control functions beyond those previously done. SNA communication products can perform functions that were formerly done by the main processor. The functions distributed to these products may include the management of communication lines, device control, data formatting, and in some cases, execution of application programs.

SNA defines system protocol for the support of distributed functions, so that the SNA products can be used to greater advantage. Both cost and complexity are reduced by avoiding redundant procedures; this encourages the orderly development and growth of communication-based systems.

The distribution of functions to SNA communication products increases the capacity and scope of the communication system; this can provide advantages such as:

- *Improved response time.* Local transaction processing, ranging from simple editing to local data base access, allows transactions to be handled at the point of origin. Only transactions requiring access to the central data base or central programming resources need be forwarded to the main processor.

- *Decreased communication line costs.* Handling transactions at the point of origin can decrease communication line expense by decreasing traffic. The capability to support several applications within SNA products can also reduce cost by decreasing the number of communication lines required.

- *Decreased main processor load.* The distribution of communication line management, device control processing, and transaction processing all contribute to reducing the load on the main processor.

- *Improved availability.* Critical functions can continue to be handled locally by SNA products following the failure of the main processor or of a component in the communication path.

Features of SNA

- *Distributed Function*

SNA supports distribution of functions among the physical components of a communication-based system. Basically, a given function is allotted to paired functional elements at the two ends of a communication path. The function is distributed in order to improve performance; to enhance reliability, availability, and serviceability; or to add capabilities.

- *Attachment Independence*

SNA defines paths between *end users* of the communication system. These paths are managed by the IBM-provided communication products. The end users (programs, devices, or operators) are presented with access to the paths that does not depend on the physical network configuration. Thus, modifications or extensions to the network configuration may be made without affecting the end user.

The architecture defines the structure and the data routing protocol between an origin and a destination. *Network addresses* identify the origin and destination end points of the path. With the aid of routing tables, they indirectly specify the *physical linkage* between these end points. The physical links are shared resources, under the control of IBM-provided support. This includes transmission error recovery procedures and path routing, as well as the blocking or segmenting of data for transmission over the data links.

- *Device Independence*

In a typical communication system, an application program communicates with the operators via different types of devices. SNA permits an application program to communicate at a level independent of unique device requirements.

- *Configuration Flexibility*

SNA unifies data link control disciplines and network protocols. It also defines the functions required in each unit so that different units can operate together in a network, sharing data links and using uniform procedures. This simplifies support requirements; it allows application programs and communication products to be added or changed without affecting other elements of the communication system.

Chapter 2. Basic Concepts

The basic concepts of SNA were developed to provide a structure that would satisfy the requirements of customer communication systems. The following discusses these concepts, as background to the more detailed later discussion of communication system organization.

Separation of Function

A key concept of SNA is the division of the communication-system functions into a set of well-defined logical layers. These functions exist in earlier product support programs, as shown in Figure 2-1. However, their separation into logical entities is generally not formalized. When these functions are formally allocated between the user application program and the IBM-provided communication products, the structure itself must be formalized. The major functional layers defined by SNA are:

- Application layer
- Function management layer
- Transmission subsystem layer

SNA is structured in layers (see Figure 2-2) for two basic reasons:

1. To permit changes to be made in one layer without affecting other layers.
2. To allow interactions between functionally paired layers in different units. This pairing is required to support distribution of function.

Transmission Subsystem Layer

The *transmission subsystem* is concerned with the routing and movement of data units between origins and destinations. The transmission subsystem does not examine, use, or change the contents of these data units. This separation, where the routing of a data unit is independent of the contents of the data unit, means that a change in the method of transmission between nodes requires no change in the data unit itself. Therefore, the support provided by the function management layer can be used across a variety of physical connections.

Paths through the network may be shared by many applications. The paths may consist of several physical components with interconnecting data links. The transmission subsystem provides the control necessary to manage these shared resources.

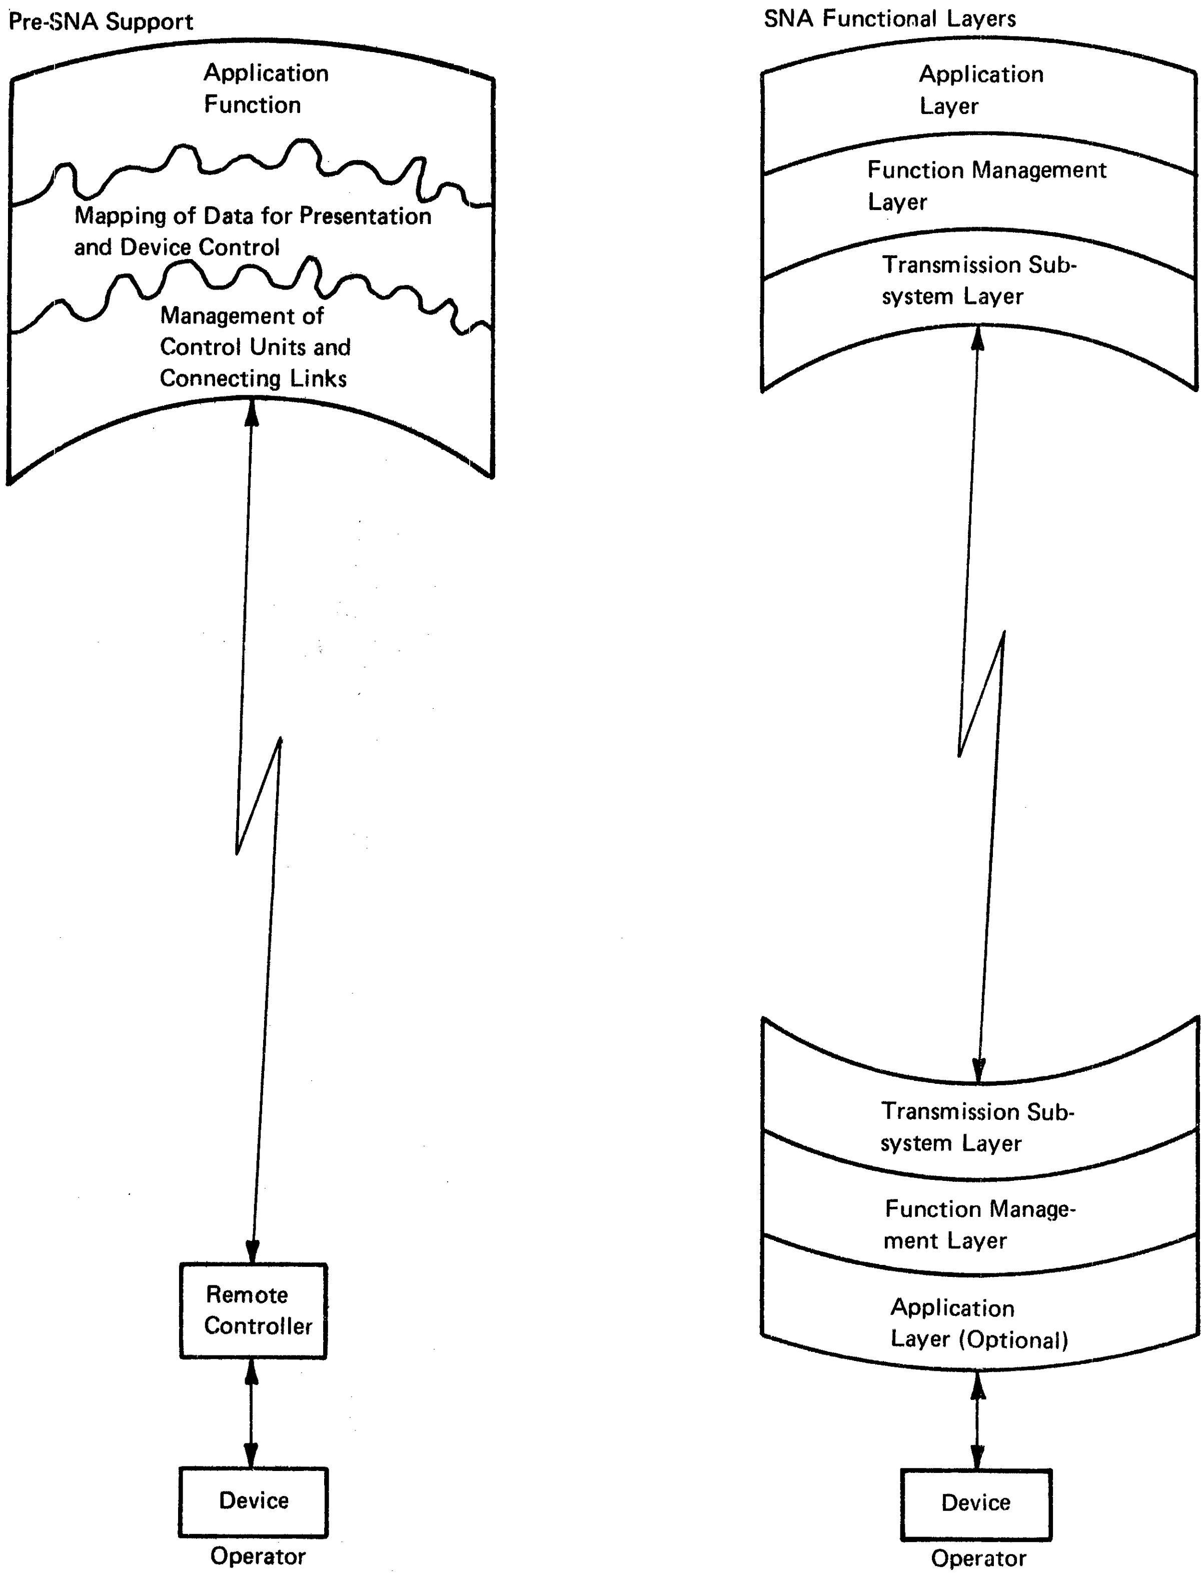

Figure 2-1. Functional Layering: Old vs New

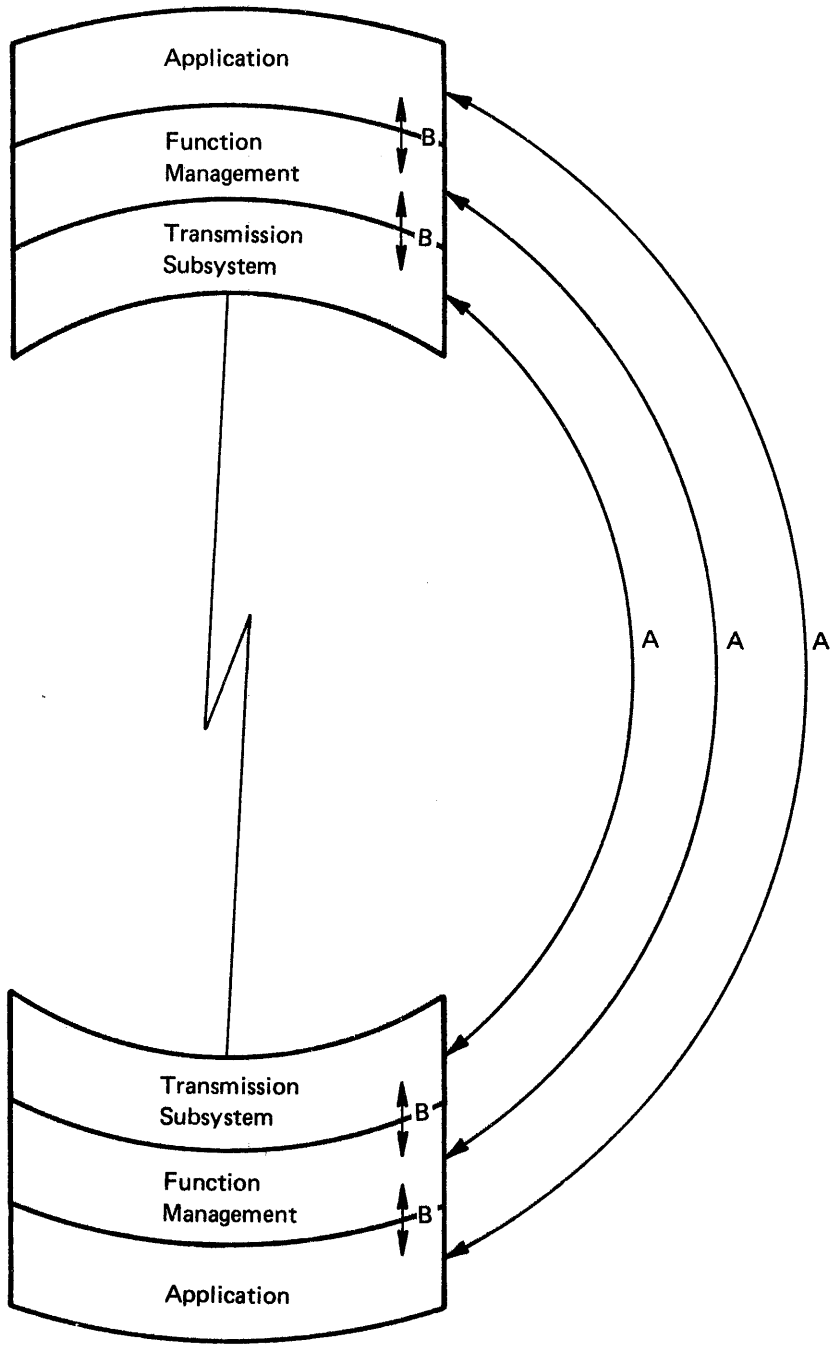

A: Equivalent Layer Communication
B: Adjacent Layer Communication

Figure 2-2. Communication Between Layers

Function Management Layer

The application layer employs a set of requests to invoke the services of the *function management* (FM) layer.

The function management layer is concerned with the presentation of information from one application layer to another application layer. Separation of the function management layer from the application layer and from the transmission subsystem layer allows device-specific transformations to be distributed out of the main processor and into the new products as shown in Figure 2-1.

Application Layer

The *application* layer is concerned only with application functions. This layer performs the user's application processing and need not be involved in the protocol or procedures for controlling a communication line or routing data units through the network.

SNA Structure

The layers introduced in the previous section have an inner structure that is discussed in the following.

End Users

As shown in Figure 2-3, *end users* are the ultimate sources and destinations of information. End users include programs and operators (such as terminal users and network administrators) as well as certain physical device media (such as cards, tapes, etc.). The structure of SNA allows end users to be independent of, and unaffected by, the specific services and facilities used for information exchange. Communication between end users requires two kinds of activity: adapting data and control conventions peculiar to each of the end users into the general form used for transmission; and actually transmitting the data across a data link, or series of data links.

Communication System

The *communication system* includes the elements in all nodes supporting end user to end user communication. SNA defines the logical structure, formats, protocol, and operational sequences among elements of the communication system.

Network Addressable Unit

The *network addressable unit* (NAU) is a resource managed by the communication system. It provides a port for end user access to the communication system. NAUs are the origins and destinations of information units flowing in the communication system.

Each NAU has a unique *network name* by which end users identify the NAU. Each NAU is also assigned a *network address* by the communication system. The network address is used within the transmission subsystem and uniquely identifies the location of the NAU within the communication system.

A formally bound pairing, called a *session*, must be established between two NAUs before their end users can communicate. The communication system does this when an end user invokes a connection protocol for initiating a session.

This protocol requires that the communication system be provided with the network name of the NAU to be linked to. The session is established between the named NAU and the NAU used by the requesting end user.

One aspect of establishing a session is selecting the appropriate *presentation service*—a function management element described in a later section. Presentation services (PS) provided by the communication system are allocated or bound to the NAUs as required to support the session.

Another aspect of a session is the path connecting the two NAUs. The endpoints of this path are uniquely identified by the network addresses of the two NAUs. The structure and use of the network addresses permit efficient routing of information flowing within sessions.

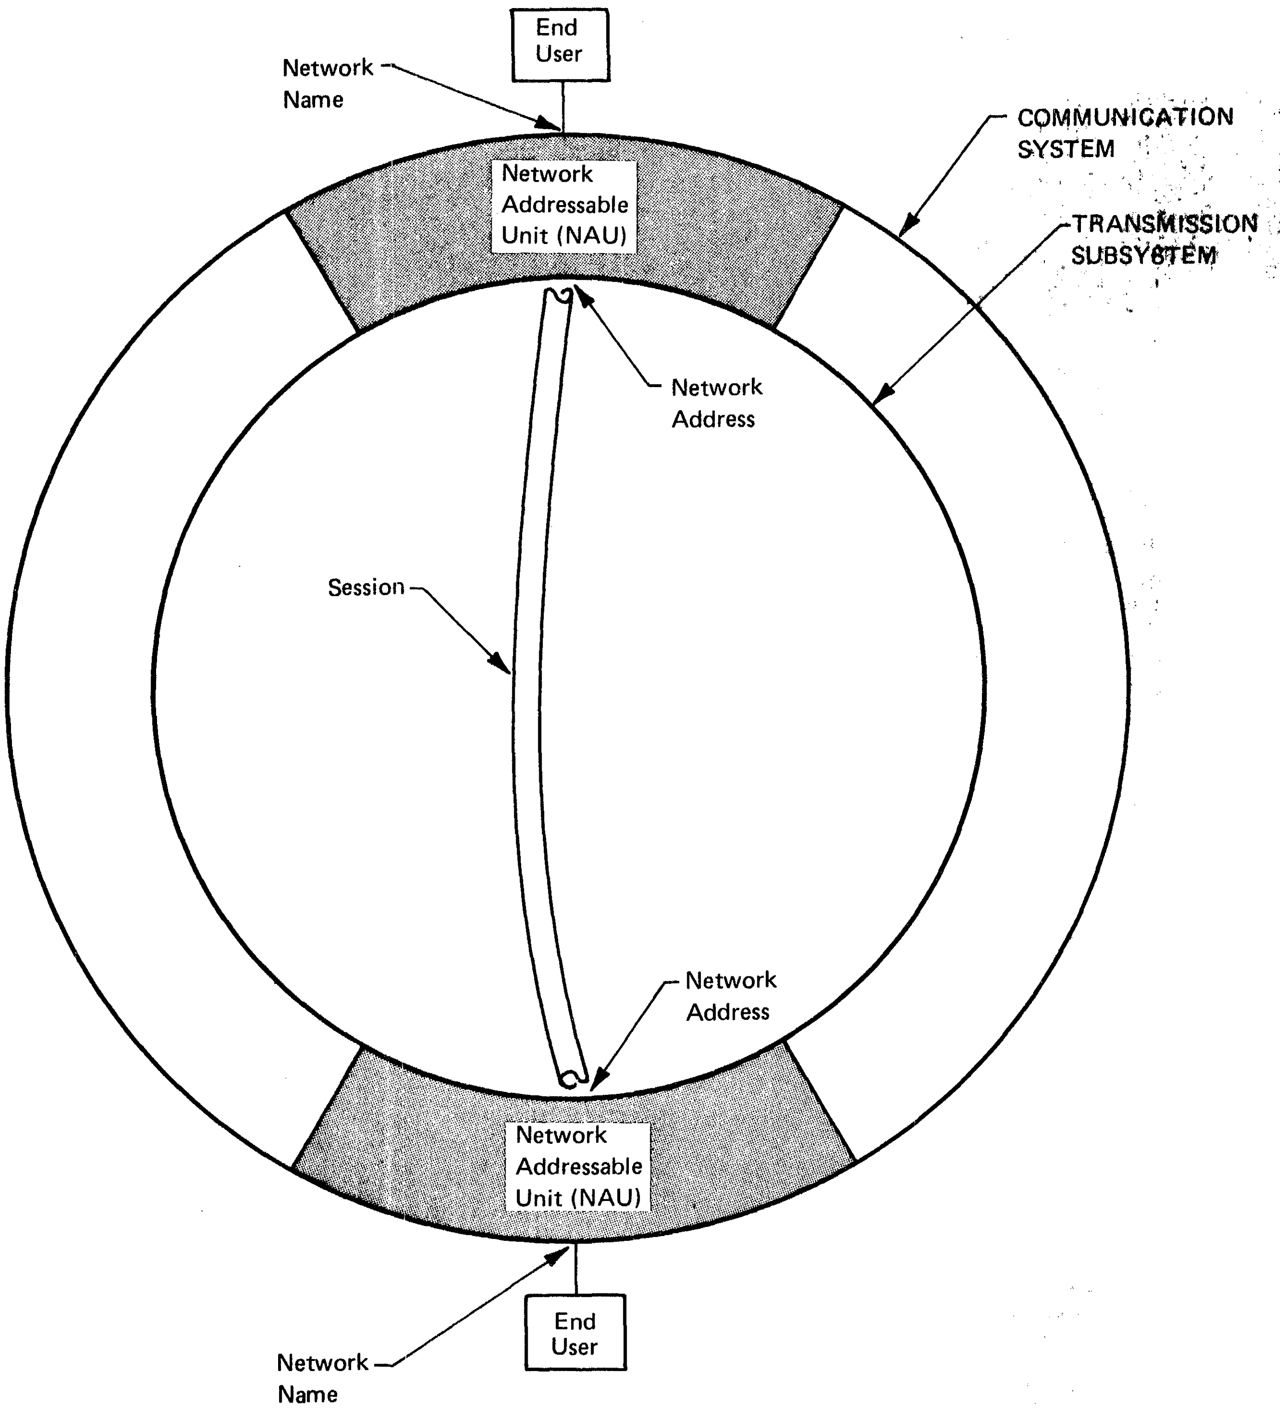

Figure 2-3. Basic SNA structure

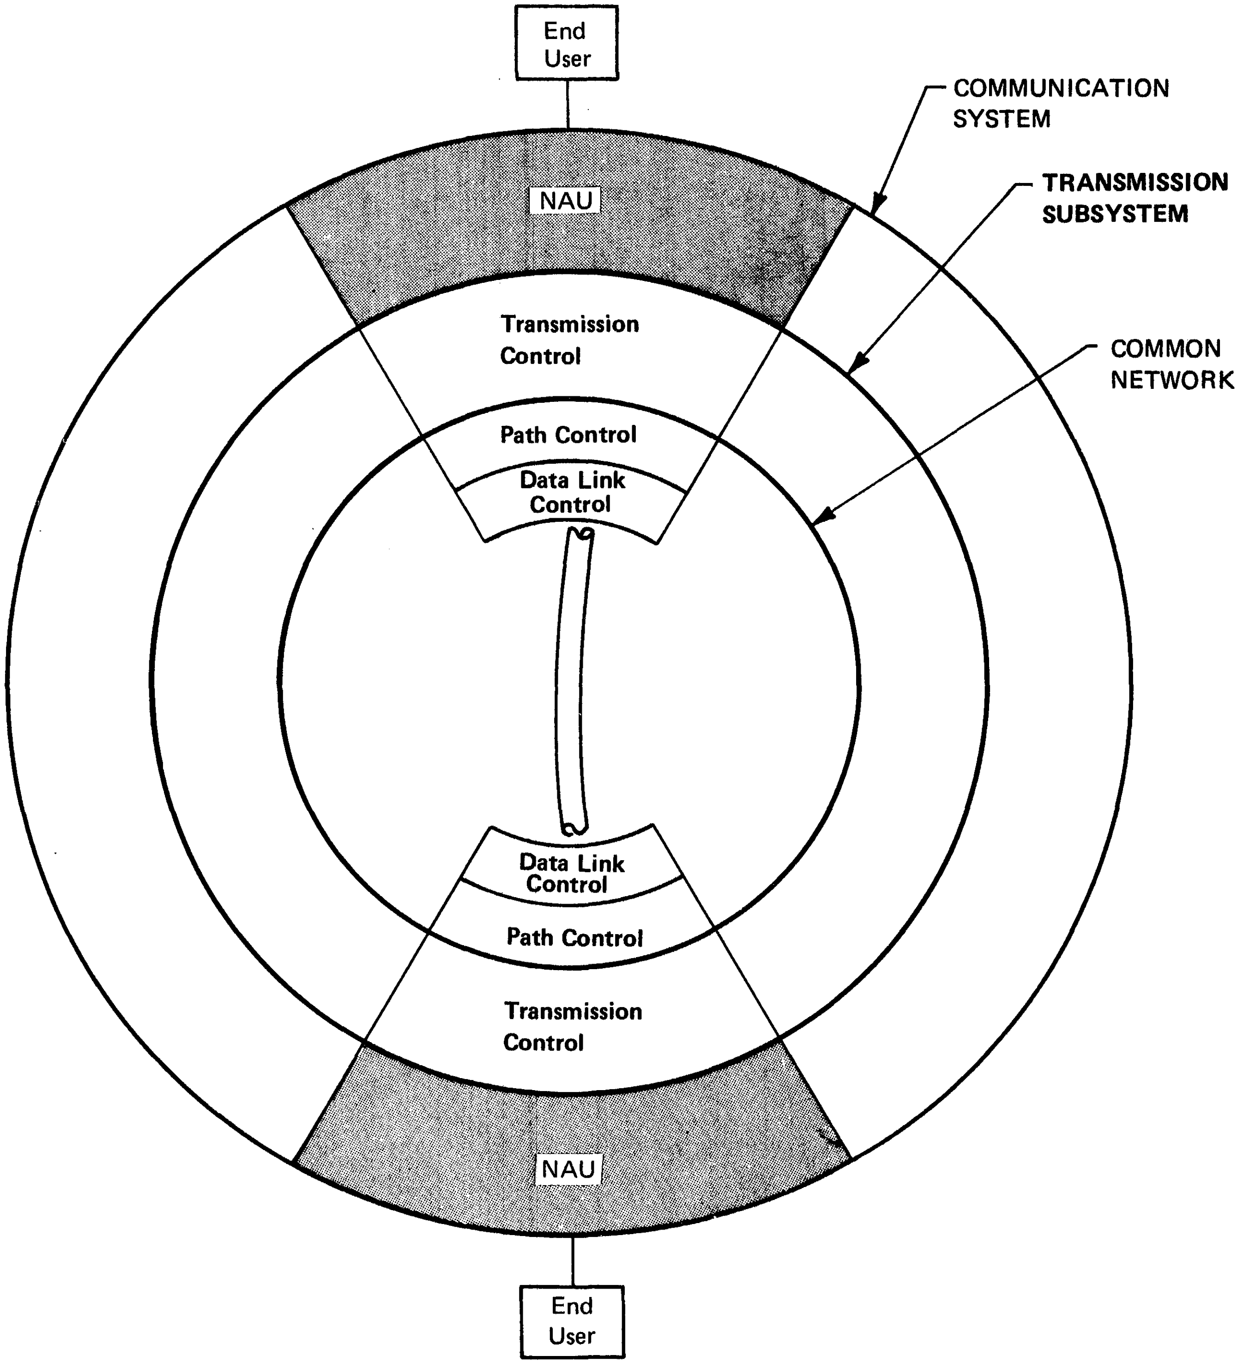

Figure 2-4. Transmission Subsystem

Transmission Subsystem

The transmission subsystem provides data exchange between NAUs. Three types of elements (see Figure 2-4) make up the transmission subsystem. *Data link control* elements manage the links between nodes. *Path control* elements provide routing of data units over the paths between network addresses. Collectively, the path control and data link control elements make up the shared *common network*. The common network routes data units between transmission control elements. *Transmission control* is the third element of the transmission subsystem. Transmission control elements control sessions and manage the flow of data into and out of the common network.

A NAU may be able to support multiple sessions with other NAUs in the communication system. The transmission control elements, using pairs of network addresses (as shown in Figure 2-5), uniquely identify and maintain the integrity of each session. One network address identifies the transmission control element itself and the other identifies a corresponding transmission control element.

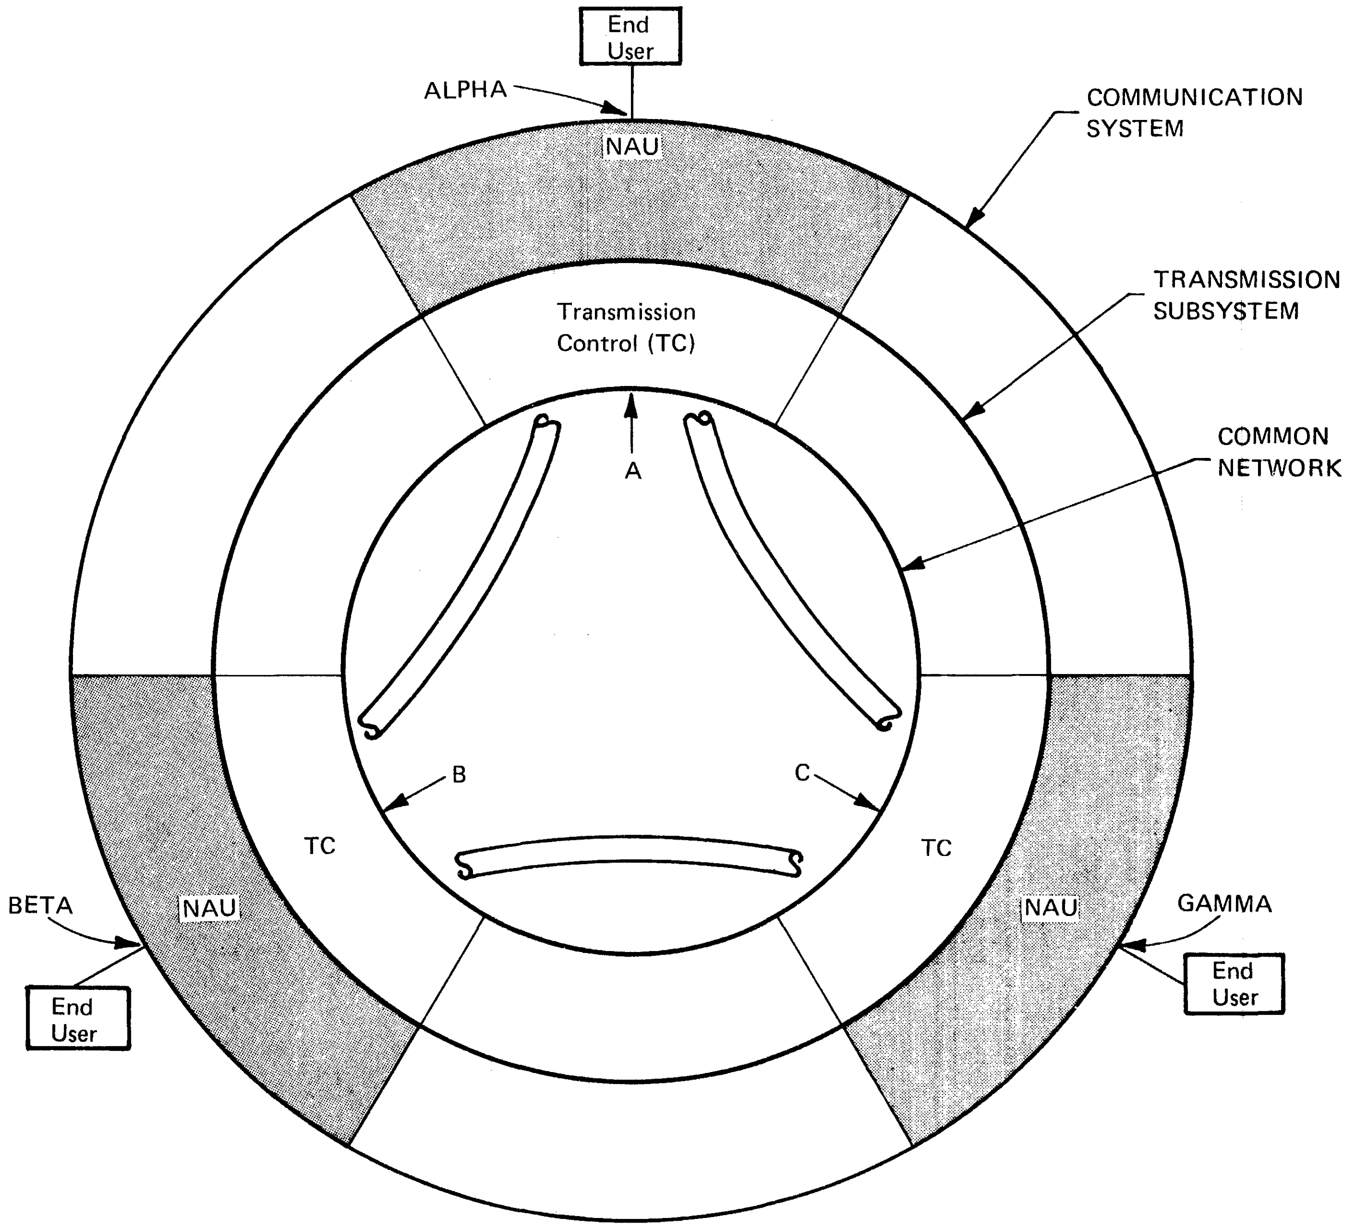

Legend:

Network names: ALPHA, BETA, GAMMA

Network addresses: A, B, C

Sessions: A to B, A to C, B to C

Figure 2-5. Multiple Session Support

Types of Network Addressable Units and Their Sessions

SNA defines three types of network addressable units (NAUs). (See Figure 2-6.) The first type is the *system services control point* (SSCP). A set of command processors perform the services provided by the SSCP. One of the command processors, called *network services*, is responsible for the general management of the network (such as bringing up the network, establishing sessions, or recovering when a network component has failed to maintain contact). Some of the processes managed by the SSCP are initiated by commands from network operators or administrators who are responsible for operation of the network. Other SSCP processes serve requests (for example, for sessions) from terminal operators.

The second type of NAU is the *physical unit* (PU). Each node in the network whose existence has been defined to the SSCP has a PU. (Communication controllers provide PU services for certain low-function, attached terminals as a boundary function.) The SSCP establishes a session with each PU in the network as part of the bring-up process. This session is used to control the physical configuration and the communication system resources associated with the node and also to collect maintenance and operational statistics.

The third type of NAU is the *logical unit* (LU). The SSCP also establishes a session with each LU in the network as part of the bring-up process. The LU is the port through which an end user accesses the SSCP-provided services such as session establishment. The LU is the port through which the end user also accesses the presentation services provided by the communication system to support end user to end user (or LU-LU) communication. A logical unit can support at least two concurrent sessions, one with the SSCP and one with another logical unit. Some LUs can support multiple sessions with other LUs.

As shown in Figure 2-6, three kinds of sessions are defined among NAUs:

1. LU to LU
2. LU to SSCP
3. PU to SSCP

Following is a summary of function management (FM) services provided within the different NAU types. A later section has more details.

1. *Presentation services:* the FM service within an LU for the LU to LU session. Since a logical unit may have sessions with many other logical units, multiple presentation services components may comprise the FM layer within an LU.
2. *Logical unit services:* the FM service within an LU for the LU to SSCP session.
3. *Physical unit services:* the FM service within a PU for the PU to SSCP session.
4. *Network services:* the FM services components within an SSCP providing configuration, maintenance, and session services for sessions between the SSCP and PUs or LUs. The SSCP also provides operator services for interaction with a network operator/administrator.

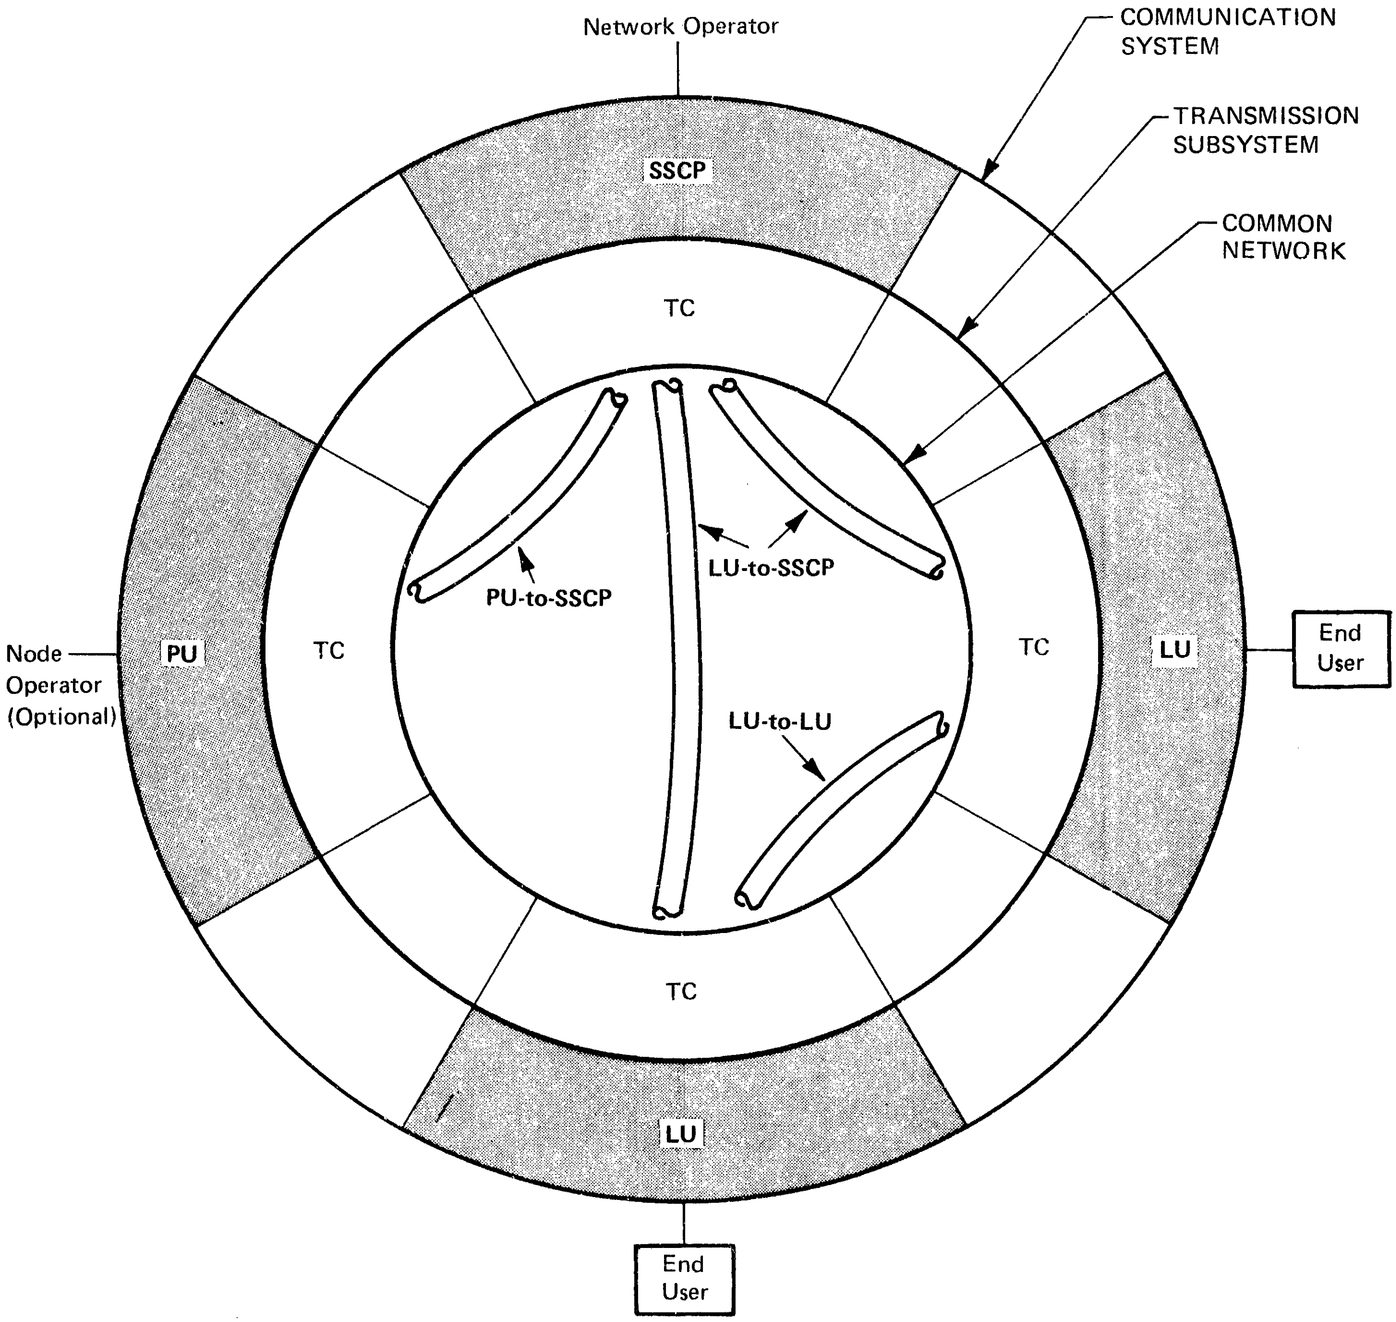

Figure 2-6. Types of Network Addressable Units and Their Sessions

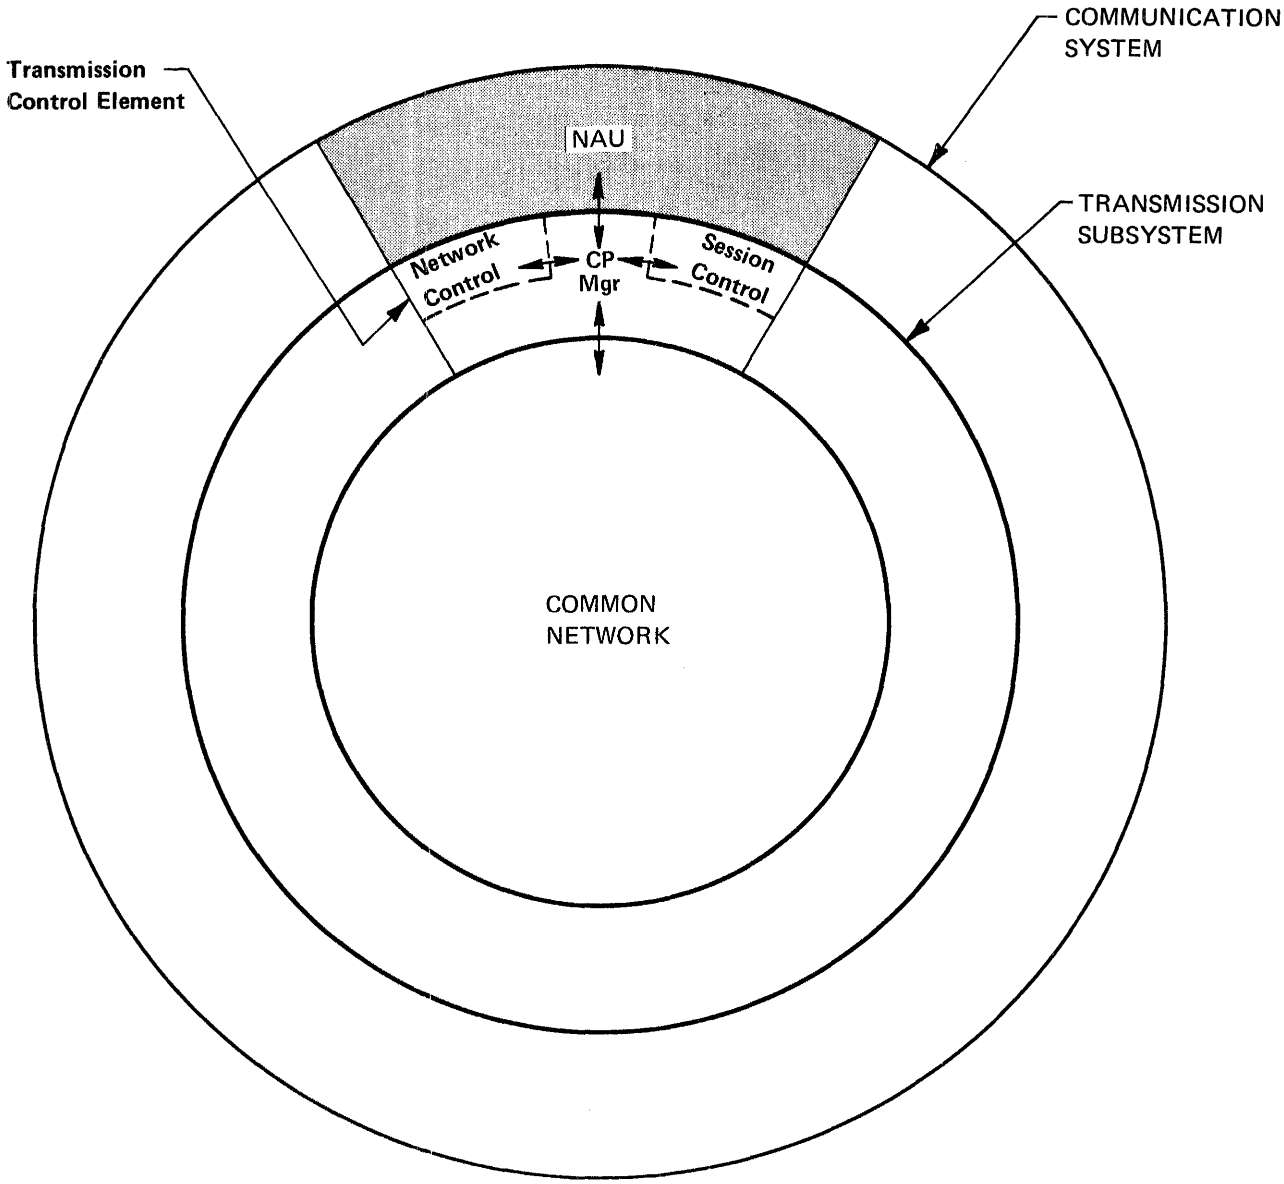

Figure 2-7. Transmission Control Element

Transmission Subsystem

Transmission Control Element

The *transmission control* element provides its users with direct access to the transmission subsystem. This direct access is used by function management services within a NAU when *requests* or *responses* are sent across the common network to counterpart function management services. Transmission control consists of three components: the *connection point manager, session control*, and *network control.* (See Figure 2-7.)

The *connection point (CP) manager* component provides a common mechanism by which session control, network control, and NAUs communicate with their corresponding elements through the common network. The unit of information the CP Manager receives from the NAUs, session control, or network control is a request-response unit (RU). The CP Manager also receives additional control information from them.

The CP Manager uses the control information to build a *request-response header* (RH), which is prefixed to each RU sent by the CP Manager and which is interpreted by the CP Manager receiving the RH-RU combination. (See Figure 2-8.) This RH-RU combination is called a *basic information unit* (BIU). The functions performed by the CP Manager include routing a request or response to its destination, assigning and checking sequence numbers for each BIU transmitted or received, coordinating responses with requests and keeping them in proper order, pacing data traffic, and creating or reacting to exception status and sense information.

The *session control* (SC) component is used to establish a session and to obtain resources required for a session. Session control also provides facilities to clear data flowing within a session if a catastrophic error occurs, and then to resynchronize the data flow after such an error.

The *network control* (NC) component provides a means for CP Managers and path control to communicate through the common network, using sessions that have been established for NAU-to-NAU communication. Such communication provided by the network control component is thus accomplished without creating an additional session.

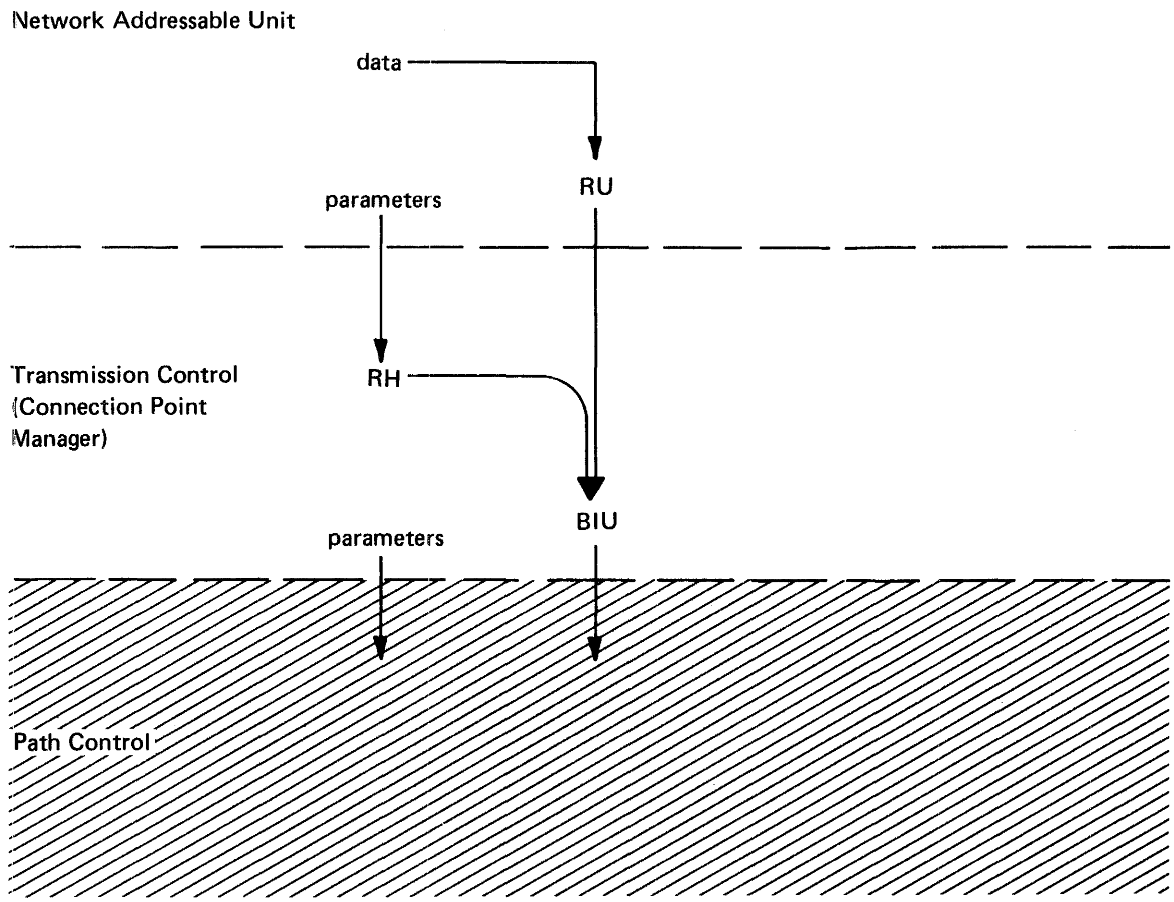

Figure 2-8. Basic Information Unit (BIU)

Path Control Element

Path control (PC) elements manage the shared data link resources of the common network and route basic information units (BIUs) through the common network. Path control, which is aware of the locations of NAUs and of the paths connecting them, maps the BIUs handled by transmission control into the data units which are transmitted.

Figure 2-9 illustrates the sequence of steps performed by path control in readying a BIU for transmission,. The first step is the division of a BIU into segments for piecemeal transmission segment by segment. This is an optional step that depends on factors such as the bigness of the BIU to be transmitted as well as accommodation of transmitted unit sizes to buffer sizes at the target node. Figure 2-10 (a) illustrates segmenting.

Path control places the control information needed by succeeding path control elements in a transmission header (TH). The TH contains addressing, mapping (segmenting indicators), sequencing, and other information needed for the transmission of data. The TH is attached to each transmitted BIU or BIU segment (if segmenting was performed). The combination of TH and BIU (or BIU segment) is called a *path information unit* (PIU).

Some of the information in the TH is created within the path control element. Other information is passed by the CP Manager to the path control element in parameter form. This information is used by path control as well as by the elements that generated the information. An example of this is the sequence number that is generated by the connection point manager, and is used within both the CP Manager and path control elements.

Once the TH has been affixed, path control performs a final step of optionally blocking one or more PIUs into a *basic transmission unit.* Blocking is useful when several PIUs destined for different (or even the same) locations are ready at the same time for transmission over a link common to and shared by several paths. Blocking is characteristic of the transmissions between host and communication controller. Figure 2-10(b) illustrates blocking.

Path control also masks the primary-secondary relationships that exist between data link control elements, thereby allowing transmission relationships to be symmetric.

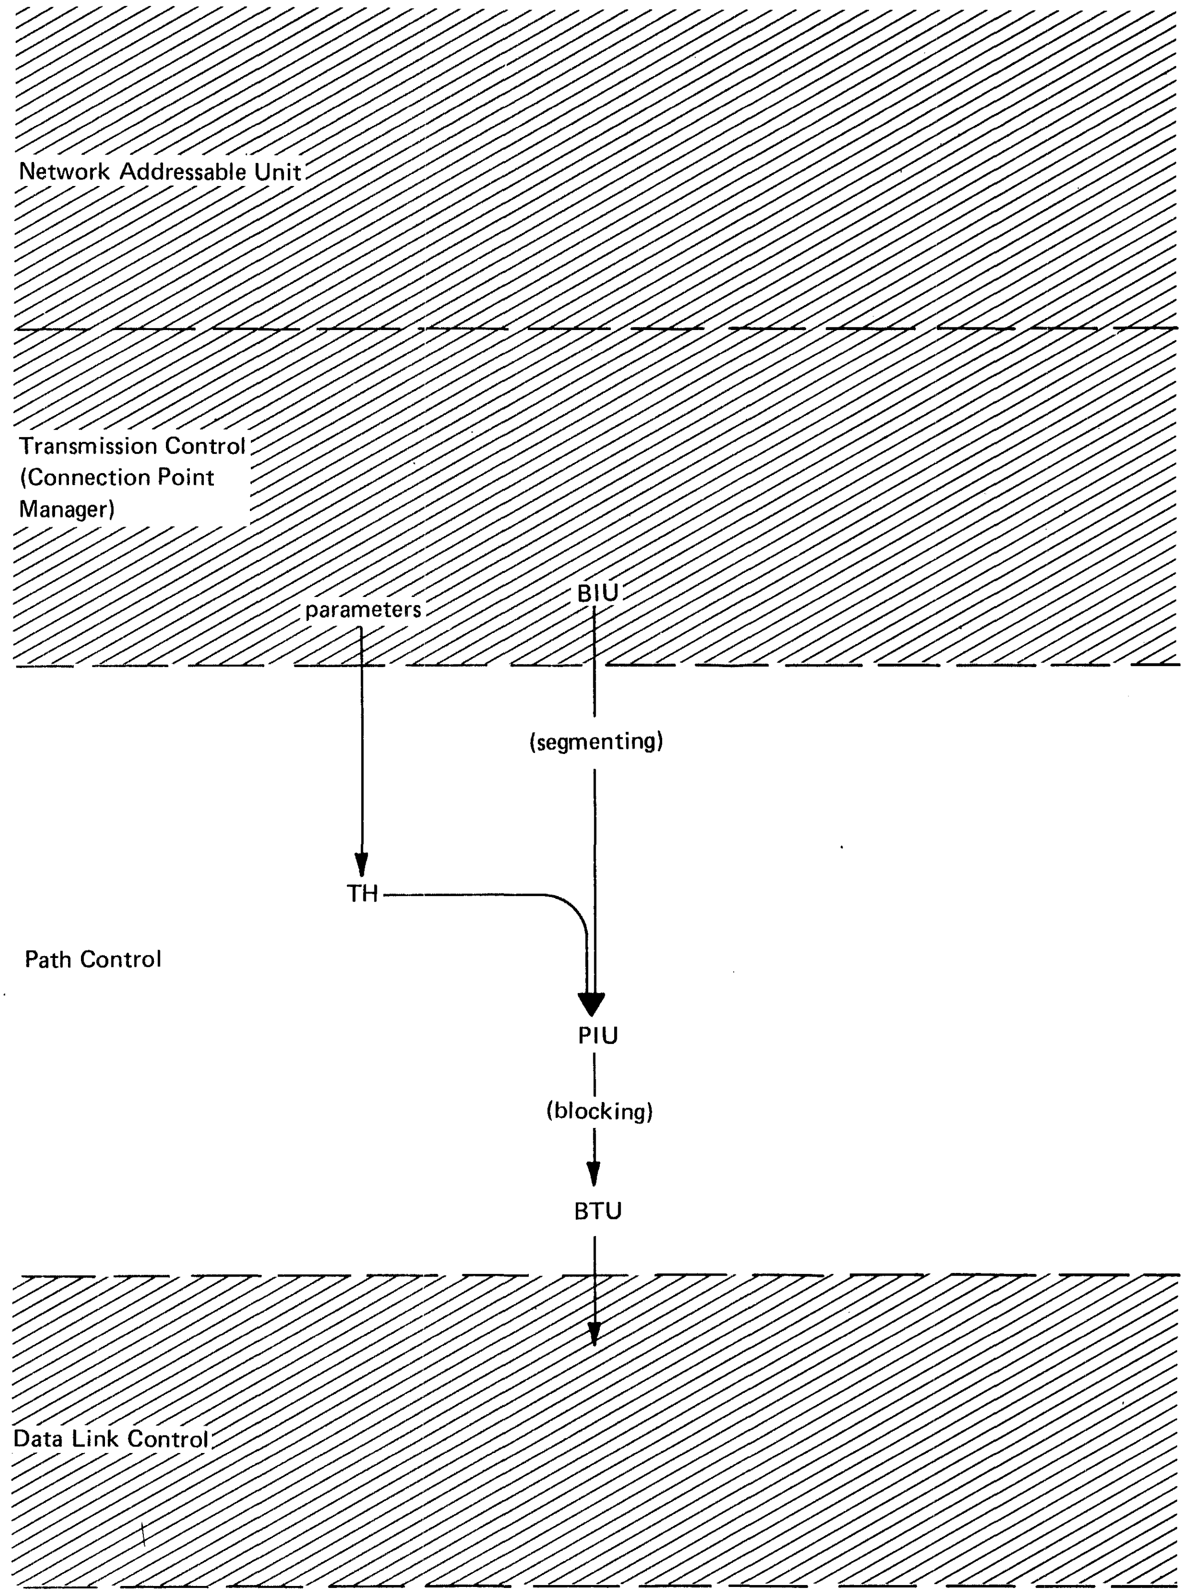

Figure 2-9. Path Information Unit (PIU) and Basic Transmission Unit (BTU)

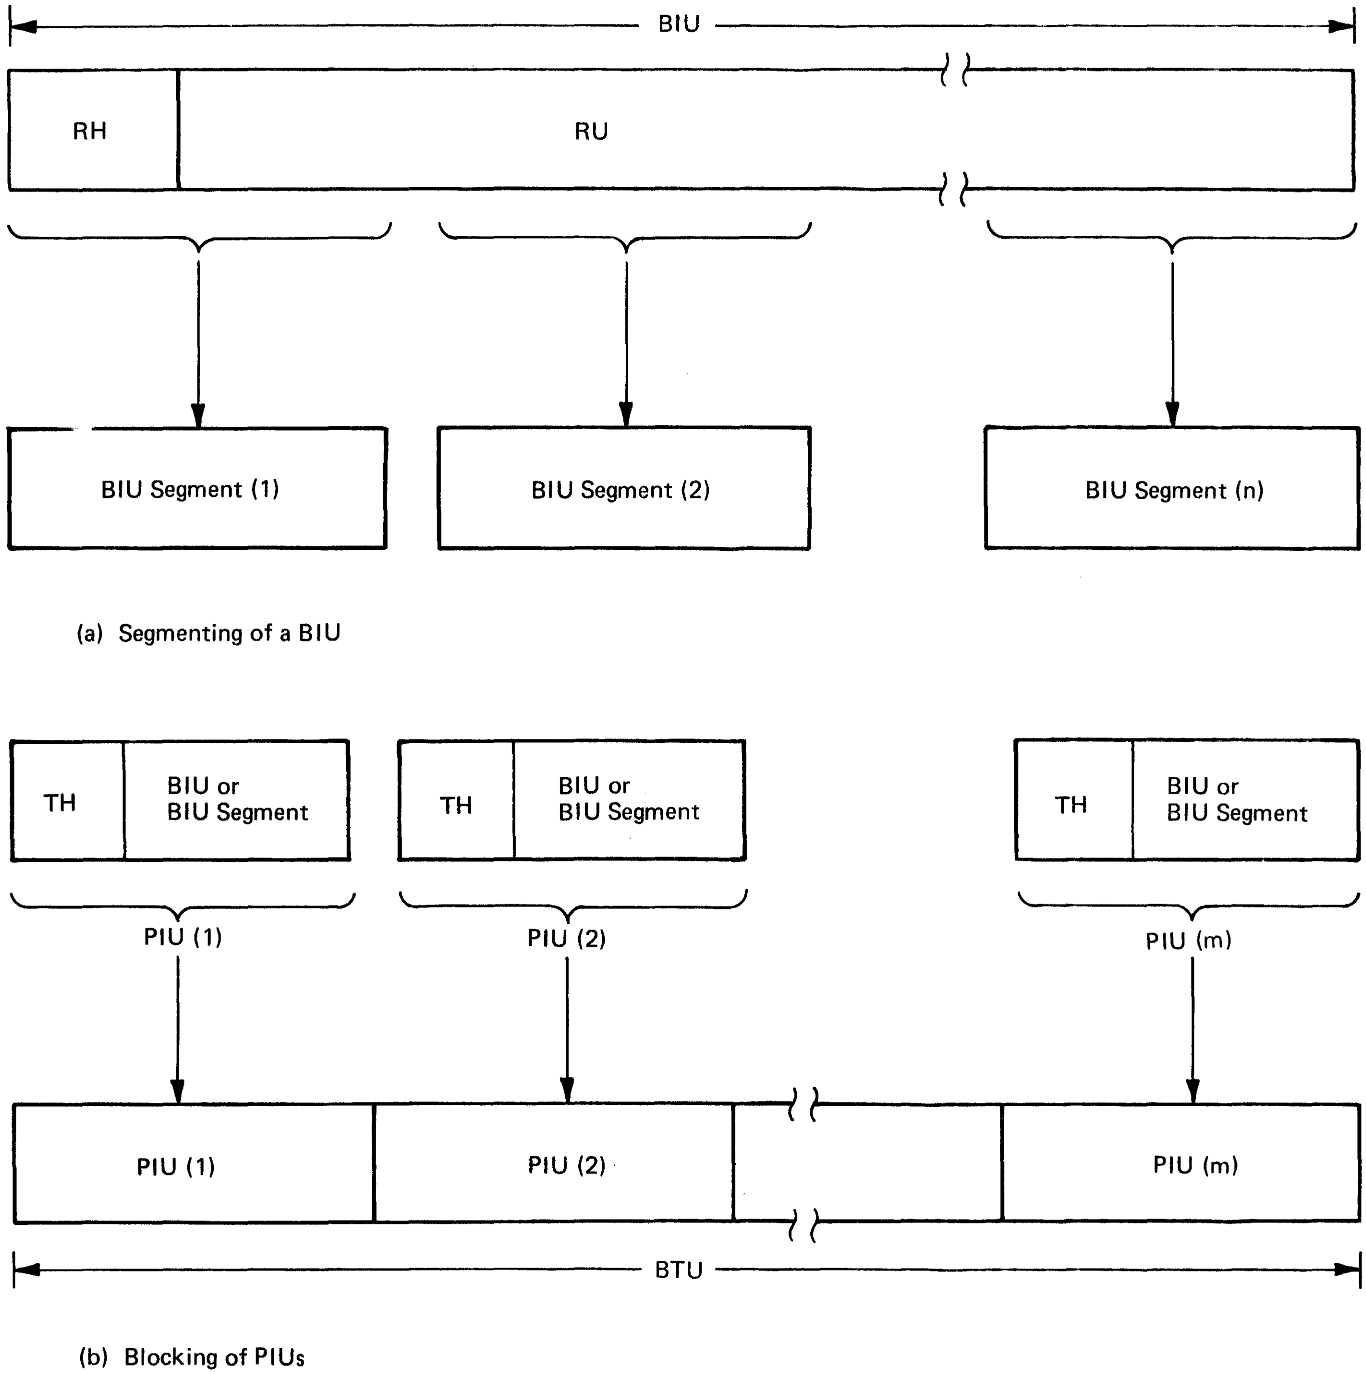

Figure 2-10. Segmenting and Blocking

Data Link Control Element

The *data link control* (DLC) element manages an individual data link. The DLC elements in each node manage the links attached to that node. DLC elements may function as either primary stations or secondary stations, depending on the physical configuration.

The procedures and protocol used to transmit data depend on the type of link being controlled. The System/370 data channel with an associated protocol is one form of data link control, and Synchronous Data Link Control (SDLC, used for transmission over common-carrier communication facilities) is another form of data link control. The DLC elements of the various stations on a given link cooperate in managing the link.

The data unit transmitted between nodes by the DLC element is called a *basic link unit* (BLU). Figure 2-11 shows the relationship of BLUs to BTUs.

Transmission Subsystem Flow

Figure 2-12 shows the elements of the transmission subsystem and presents the flow of a request-response unit (RU) through three nodes. The numbers on the right-hand side of Figure 2-12 are used in the following description.

The transmission subsystem is used to deliver RUs, unchanged, from one NAU (1) to another NAU (11).

The NAU (1) passes a request-response unit and associated parameters to the transmission control element used for the session with the other NAU (11). The transmission control element adds an RH to the RU (2), forming a *basic information unit* (BIU). The RH includes parameters that specify the type of RU being transmitted, and options and parameters associated with that RU; information in the RH is used by both the receiving transmission control element (10) and the data flow control function manager in the receiving NAU.

Transmission control passes the BIU and additional information to path control. The additional information includes the origin and destination network addresses, a unique identifier (or sequence number), and other parameters. Path control adds a *transmission header* (TH) to the BIU (no segmenting shown), forming a *path information unit* (PIU). If blocking is performed, several PIUs can be put together to form a *basic transmission unit* (BTU). The BTU shown (3) is a single PIU.

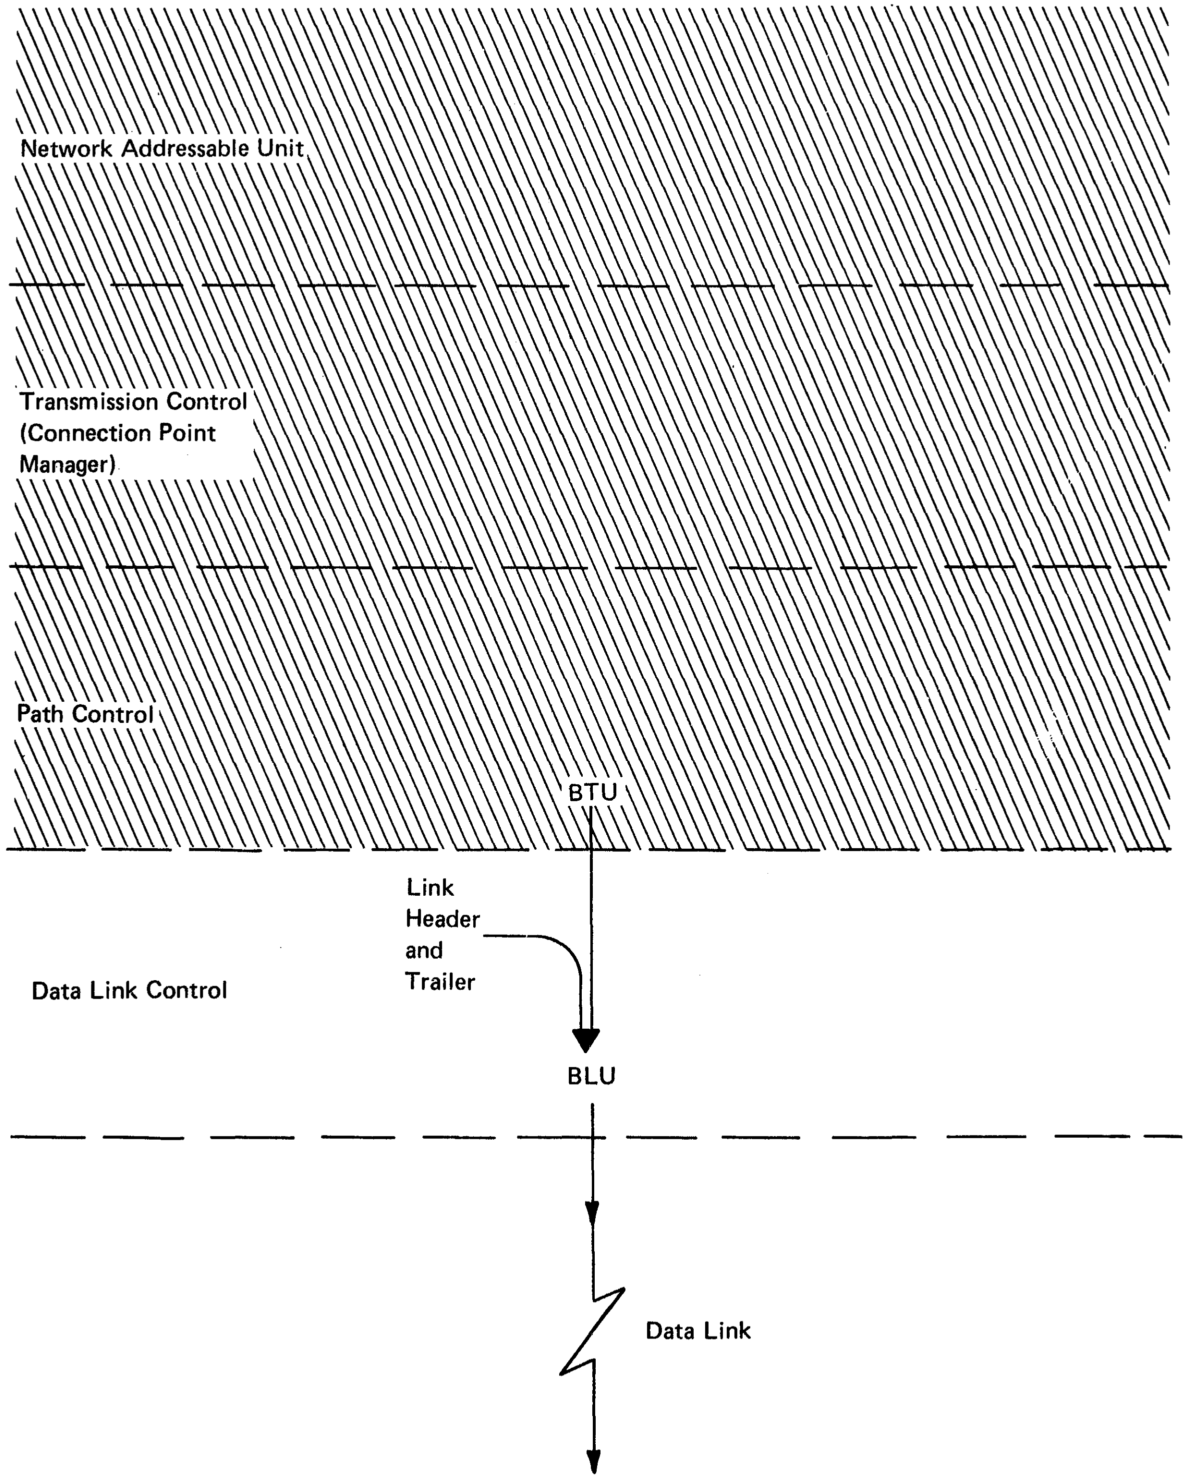

Figure 2-11. Basic Link Unit (BLU)

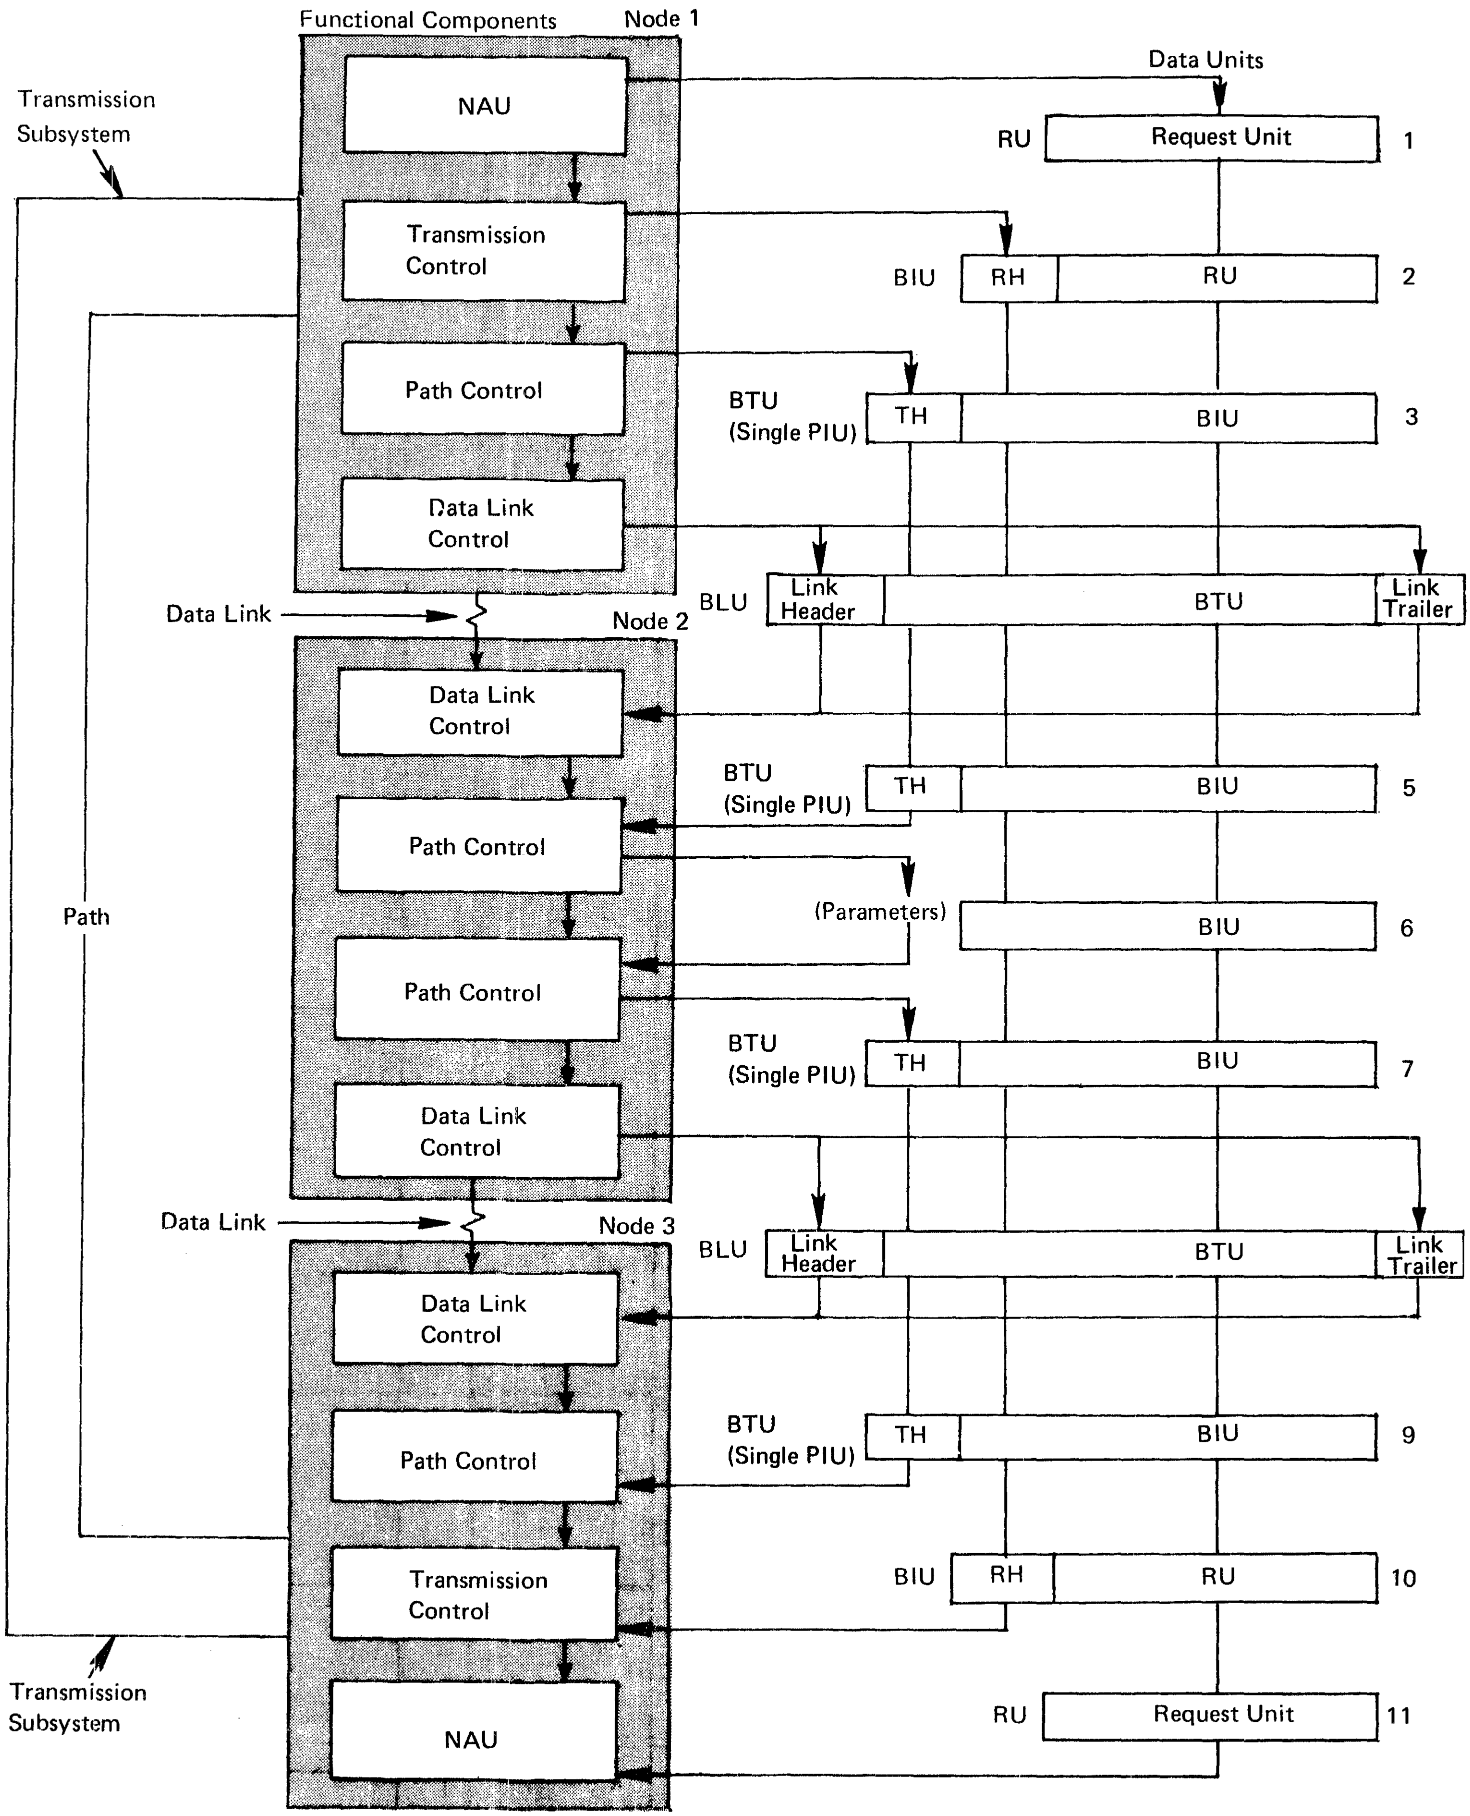

Figure 2-12. Transmission Subsystem Control Data Flow

The BTU is passed to data link control for transmission. Data link control appends link control information to form a *basic link unit* (BLU) and transmits this BLU (4) over the data link to the next node.

After an error-free transmission, the receiving DLC element strips the link control information from the BLU and passes the remaining BTU to the receiving path control element. Path control deblocks the BTU into individual PIUs (if necessary) and determines whether the destination transmission control elements are in the same node. If not, the PIUs are used to create one or more BTUs which are passed to a DLC element (7). The DLC element creates BLUs, and transmits them to the next node in the destination direction (8).

The procedure just described is repeated for each link in the path until the path control element adjacent to the final transmission control element is reached (9). The final path control element assembles a complete BIU, and passes it to the transmission control element at the destination location (10).

The transmission control element at the destination processes the control information contained in the RH and passes the remaining RU and/or parameters to the destination NAU.

The procedure is identical for requests or responses, except that the roles of the origin and the destination are reversed.

Function Management Services

Function management services provide for control of the data flow and for transformation of data presented to the network.

Data Flow Control

Function management services include an element that provides data flow control protocol support. The data flow control protocol supplied by IBM is implemented through a set of encoded requests, called *data flow control* (DFC) requests, exchanged between data flow control elements. These requests are used to handle data units and structures such as chains of related request units. They are also used to manipulate the state conditions, such as "send" or "receive" state, that are defined by the IBM-supplied data flow control protocol.

The data flow control protocol do not perform any transformation functions on end user requests (or responses), but assist the end user in controlling the flow of requests (or responses).

Presentation Services

For LUs, the data transformation elements are *presentation services* and *logical unit services*. The presentation services element of the logical unit (see Figure 2-13) provides support for communication between end users engaged in LU to LU sessions. Presentation services includes support for application programs as end users, as well as for specific hardware capabilities such as a printer/keyboard used by a terminal operator end user. The application program end user employs a set of requests to invoke presentation services. Such a set of functional requests constitutes a *presentation class*.

An important feature of the presentation class concept is that these functional requests need not depend on primitive device-specific characteristics or network configuration. Functional requests may be transformed by IBM- or customer

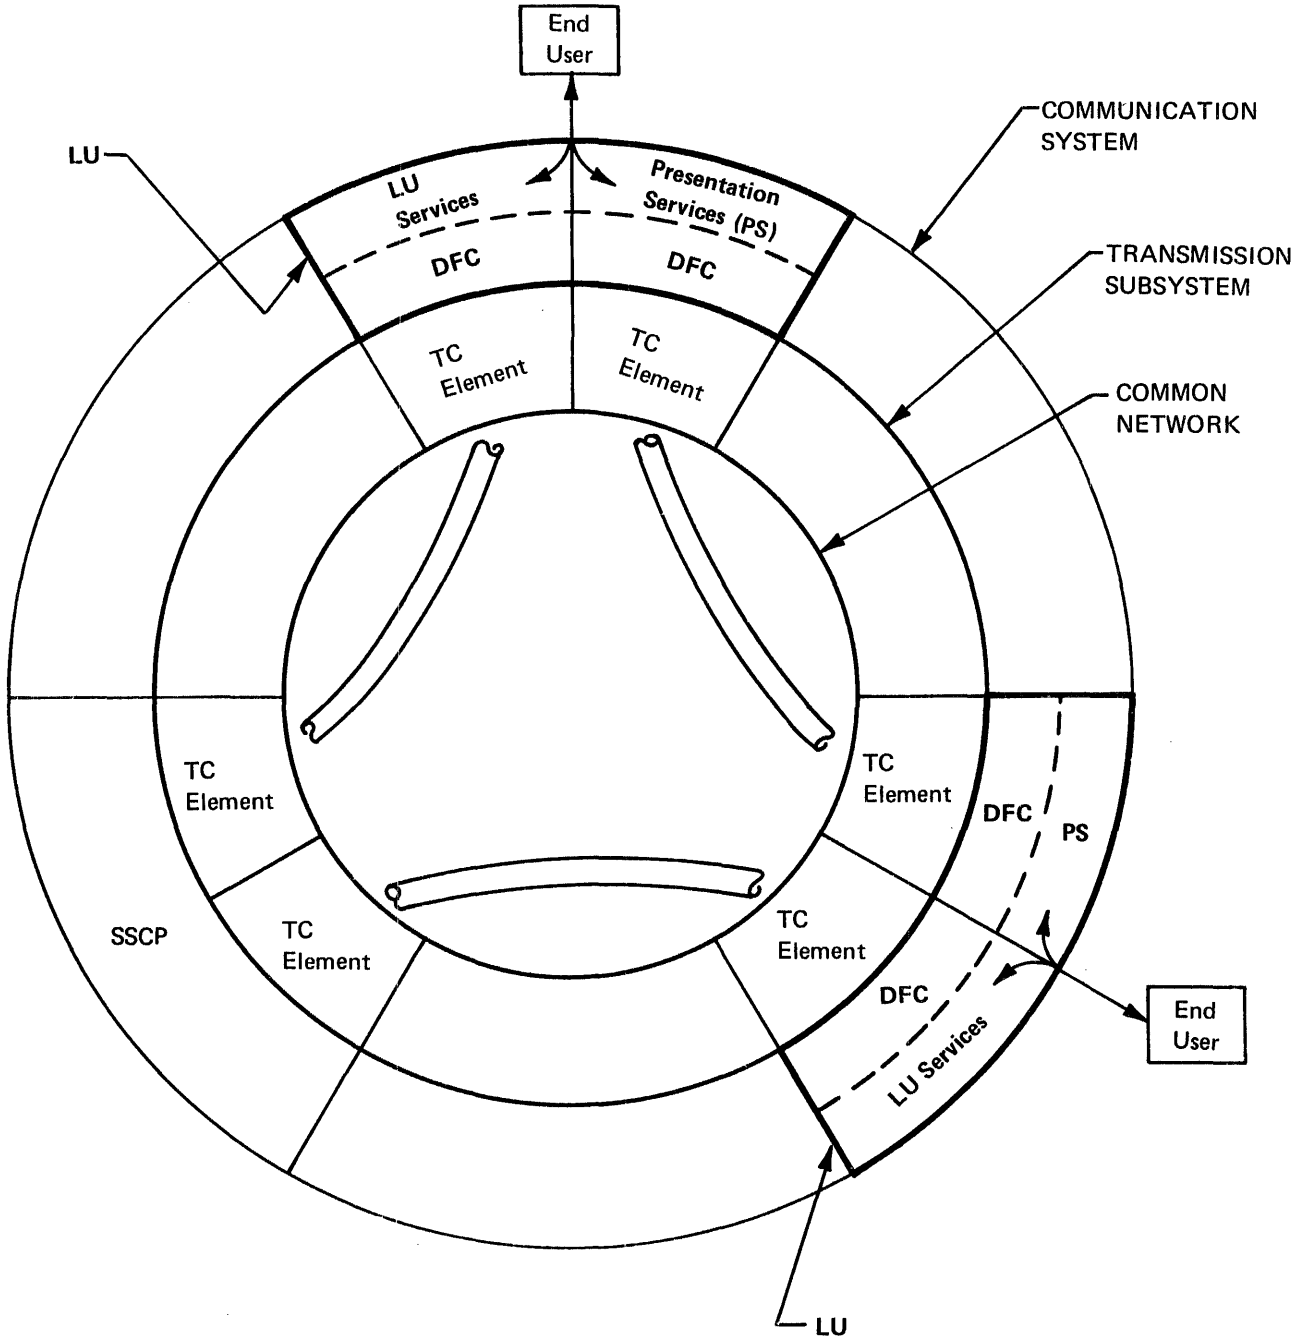

Figure 2-13. Logical Unit Function Management

supplied support programs at the function management layer into device-specific operations for a variety of devices.

The function management layer provides the end user with various levels of presentation function; he is free to choose the level of presentation function to be applied. For example, the end user may use one of the presentation classes provided by IBM. Alternatively, the end user may choose to handle a considerable amount of presentation function and ask only that the presentation services element provide control of data flow. The end user can also elect to perform both presentation service and data flow control functions. In this case, the end user works directly with the *transmission subsystem.* An example of presentation levels in the host node is depicted in Figure 2-14, showing IMS and CICS program products and VTAM support. A presentation class is a complete set of operations, which are performed on data units passed between end users. An end user may be an operator, certain physical device media, or an application program. The presentation class defines those operations available to each end user in communicating with the other end user. The presentation operations are defined separately from the format used to transmit the request or response between the presentation services pair. Error conditions are reported as errors relative to the class, and not as device-dependent conditions.

In addition to allowing a set of end user-invoked functions to be mapped onto different devices, the presentation class establishes the environment that allows distribution of the presentation class operations.

Logical Unit Services

The function management services used to support LU-to-SSCP sessions are provided by a set of command processors. These *LU services command processors* receive commands from the end user, and may in turn issue replies or commands to the SSCP, in order to perform the requested functions.

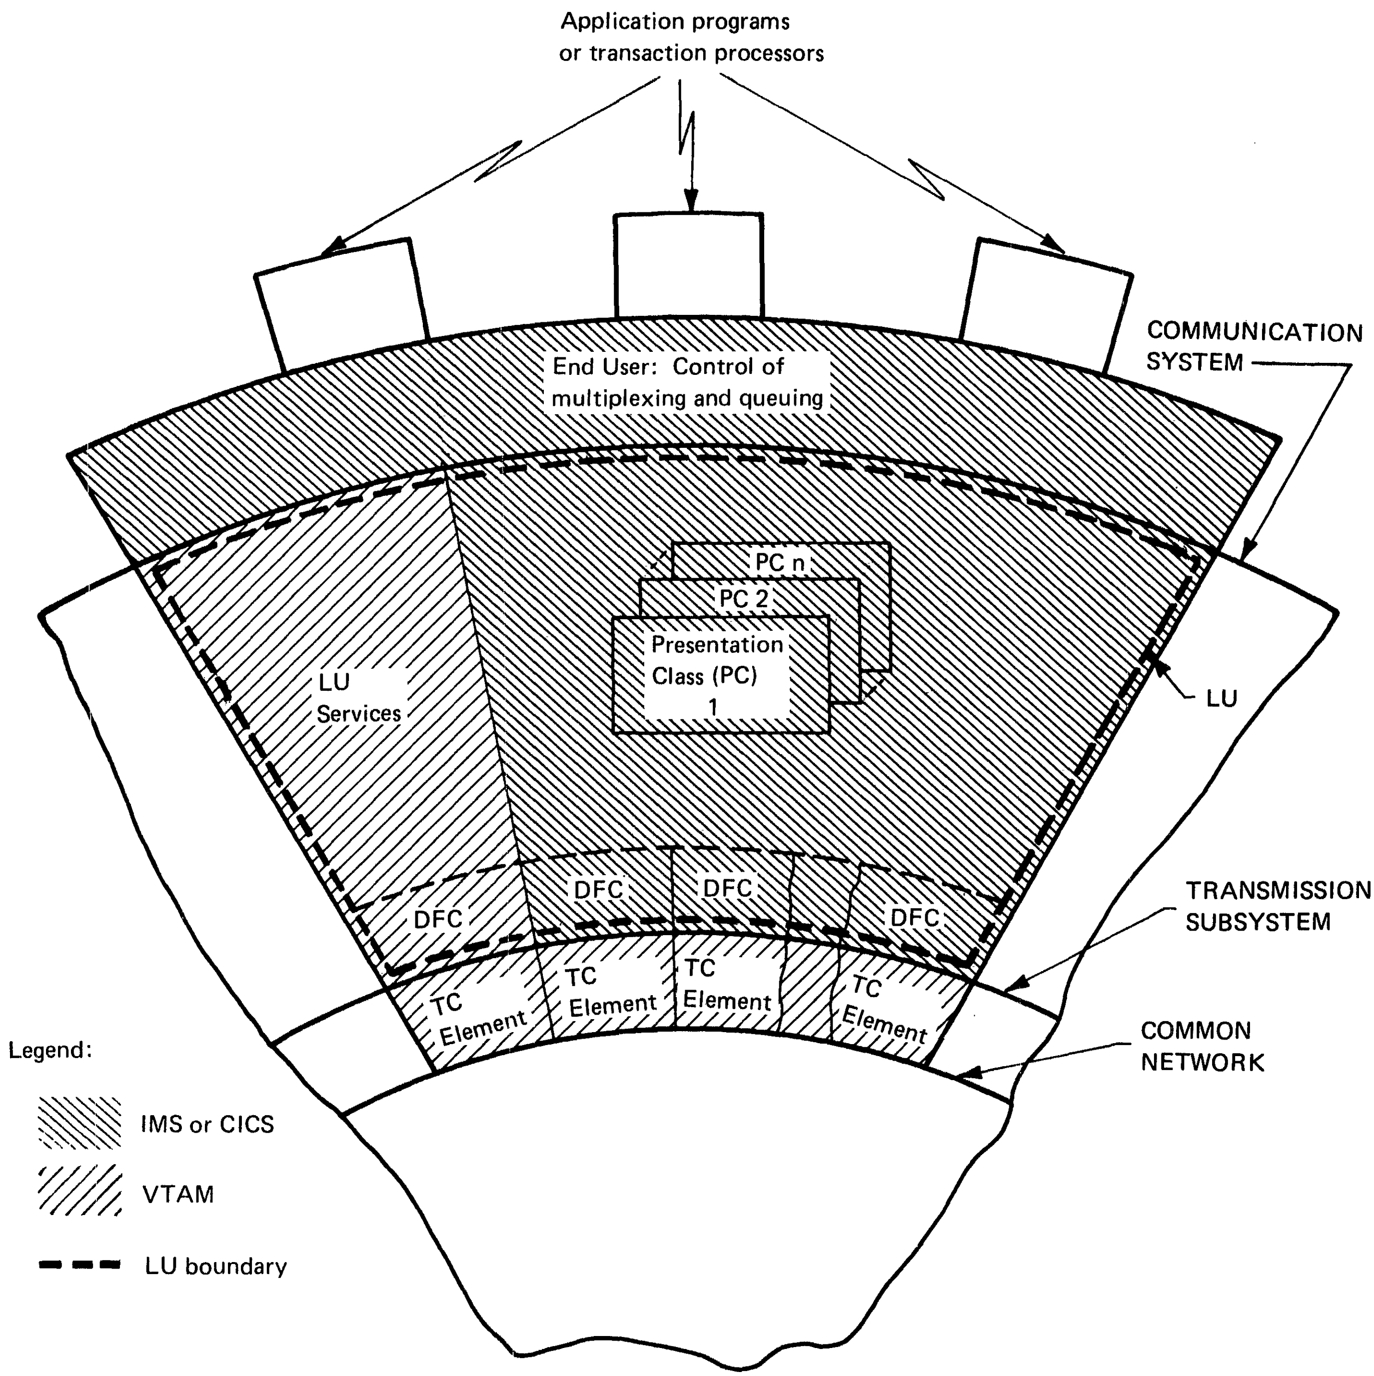

Figure 2-14. Example of Presentation Levels (Host Node)

System Services Control Point

SSCP-Logical Unit Sessions

The functions provided to end users by the *system services control point* (SSCP) are performed by a set of command processors. These command processors in the SSCP receive commands from the LU services elements in various logical units and may in turn issue commands or replies to other NAUs, in order to perform the requested function (see Figure 2-15). The LU-to-SSCP session supports LU-related control and use of the communication system. Like other FM services, LU services uses the transmission subsystem to communicate with the SSCP. The SSCP command processors supporting network users are generically called *network services.*

The LU services element allows the end user to enter commands for SSCP functions and provides the end user with appropriate replies to these commands.

An SSCP command or reply may be "formatted" (field-formatted) or "unformatted" (character-coded). The field-formatted form is used for commands from a logical unit supporting application program end users. The character-coded form is used by logical units supporting terminal operator end users. The terminal operator submits the network services command as a character string, usually by keying the appropriate format on the terminal keyboard.

When the entered command is character-coded, a translator at the system services control point transforms the command into a field-formatted command. Field-formatted commands sent to the SSCP are routed directly to the formatted command preprocessor for submission to the appropriate network services command processor.

SSCP-Physical Unit Sessions

In addition to the services for users of logical units, the system services control point provides services for each physical unit in the configuration, as well as services to support the system operators or administrators who control the configuration.

Each physical unit in the configuration has a session with the SSCP. A session also exists between the SSCP and the PU for the host node in which the SSCP is located.

These sessions are used for control of the physical configuration and for control of individual nodes and their resources (e.g., data links). The element of the SSCP that supports the PU-to-SSCP session is network services. The network configuration services element processes commands and replies used for startup and shutdown of the physical configuration as well as individual physical units, and also handles commands and replies used for recovery and resynchronization after a failure.

The SSCP also provides a control capability to the network operator/ administrator. Significant changes in the status of the various elements of the communication system are reported to the appropriate network operator/administrator.

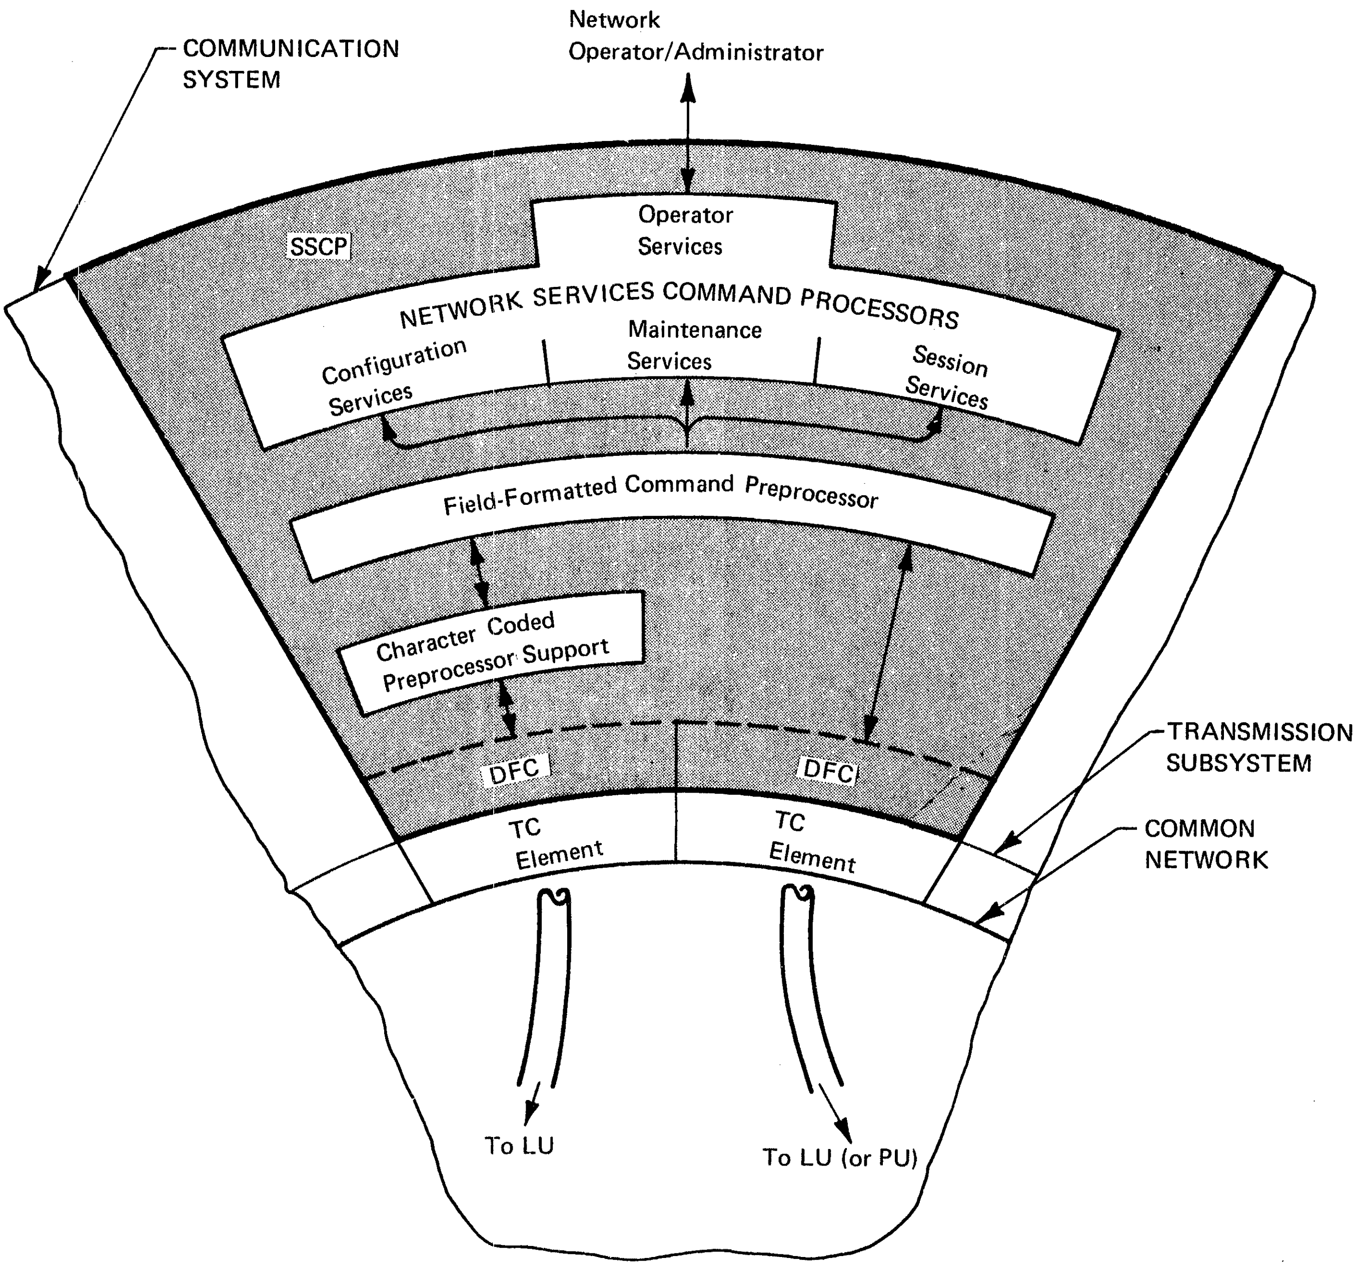

Figure 2-15. System Services Control Point (SSCP)

SSCP Network Services

SSCP network services include management of the nodes and links that make up the network. SSCP network session services deal directly with each logical unit for initiating and terminating sessions between logical units.

The network services provided by the SSCP include:

1. *Configuration services*: Support the activation and deactivation of physical and logical units and data links, and maintain the status of these elements. These network elements may be activated and deactivated by the network operator. Configuration services also support startup, shutdown, and restart of these network elements.

2. *Maintenance services*: Provide for testing the network facilities, as well as collecting and recording error information.

3. *Session services*: Provide facilities for a logical unit or network operator to request that the SSCP establish or terminate sessions between logical units. Session establishment involves:

   - Transforming the network names (that appear in the session initiation requests from NAUs) into network addresses.
   - Establishing the operating protocol to be used during the session.
   - Establishing the structure of the data to be exchanged.

Chapter 3. Communication System Configuration

Network Names and Addresses

A network name is associated with each physical unit, link, and logical unit in the configuration.

An end user requests a function or a service be provided by some element of the configuration by identifying that element's network name. The system services control point maintains a directory of network names, and transforms network names of physical units and logical units into network addresses. Names of links are also transformed into addresses. Network names are used by terminal operators, application programs, and network administrators. Network addresses are used within the transmission subsystem. They identify the origin and destination of information units flowing in the communication system.

The full network address is 16 bits long and consists of two parts: one identifying a *subarea*; and the second identifying an *element* within a subarea (see Figure 3-1). The boundary between the subarea and element parts of the network address is not specified, but for any one configuration the location of the boundary must be selected at system generation and remains constant for all addressable entities in the configuration.

SNA also defines transformations to address forms which are local to nodes. These local addresses are used by nodes with limited capability that have no need to deal with the full 16-bit network address.

Certain nodes of a network are responsible for a subarea. The specific subarea is identified by the subarea address associated with the node. These nodes handle transmission headers whose destination and origin address fields contain full network addresses. They also properly forward path information units (PIUs) for other subareas within the configuration. If the *destination address field* (DAF) contains a subarea address which identifies one of the subareas for which the node is responsible, the node uses the element address of the DAF to route the PIU to its destination.

If the final node cannot handle full network addresses, the node responsible for the subarea converts full network addresses to the local address form used by the final node. The subarea node converts the local address to a full network address when the PIU is flowing in the other direction.

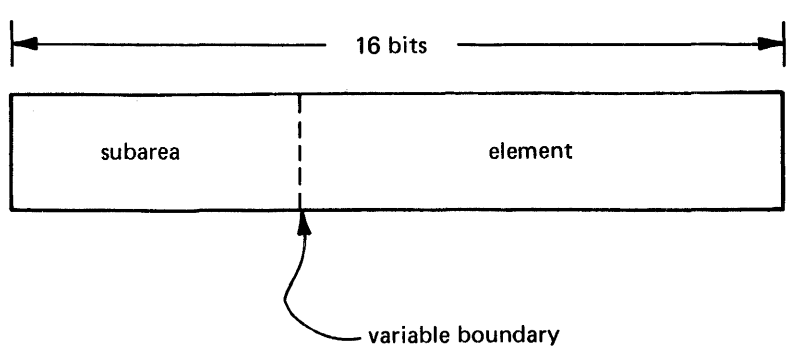

Figure 3-1. Network Address

Logical Configuration

An abstract view of the subarea configuration of a typical network is shown in Figure 3-2. Figure 3-3 shows a more concrete view. The SSCP provides control for both a fixed and variable portion of the logical configuration. The fixed portion has a static physical configuration, and the location of the various elements (physical units, logical units, and so forth) does not vary; the network addresses of these elements are fixed. The fixed portion of the logical configuration may consist of one or more subareas and associated physical and logical units. The units making up the fixed portion are connected by nonswitched communication lines.

The SSCP also provides control for a variable portion of the logical configuration through fixed access points. The location of these variable portions may vary from one "connection" to the next. The units of the variable portion connect to the fixed access points through switched communication lines. The network addresses associated with the elements within the variable portion may not be assigned until the physical connection is made.

A form of the variable logical configuration exists for application programs managed by a host operating system. These application programs use fixed access points provided by the operating system. The application program may "connect" to the logical configuration at different access points from one connection to the next.

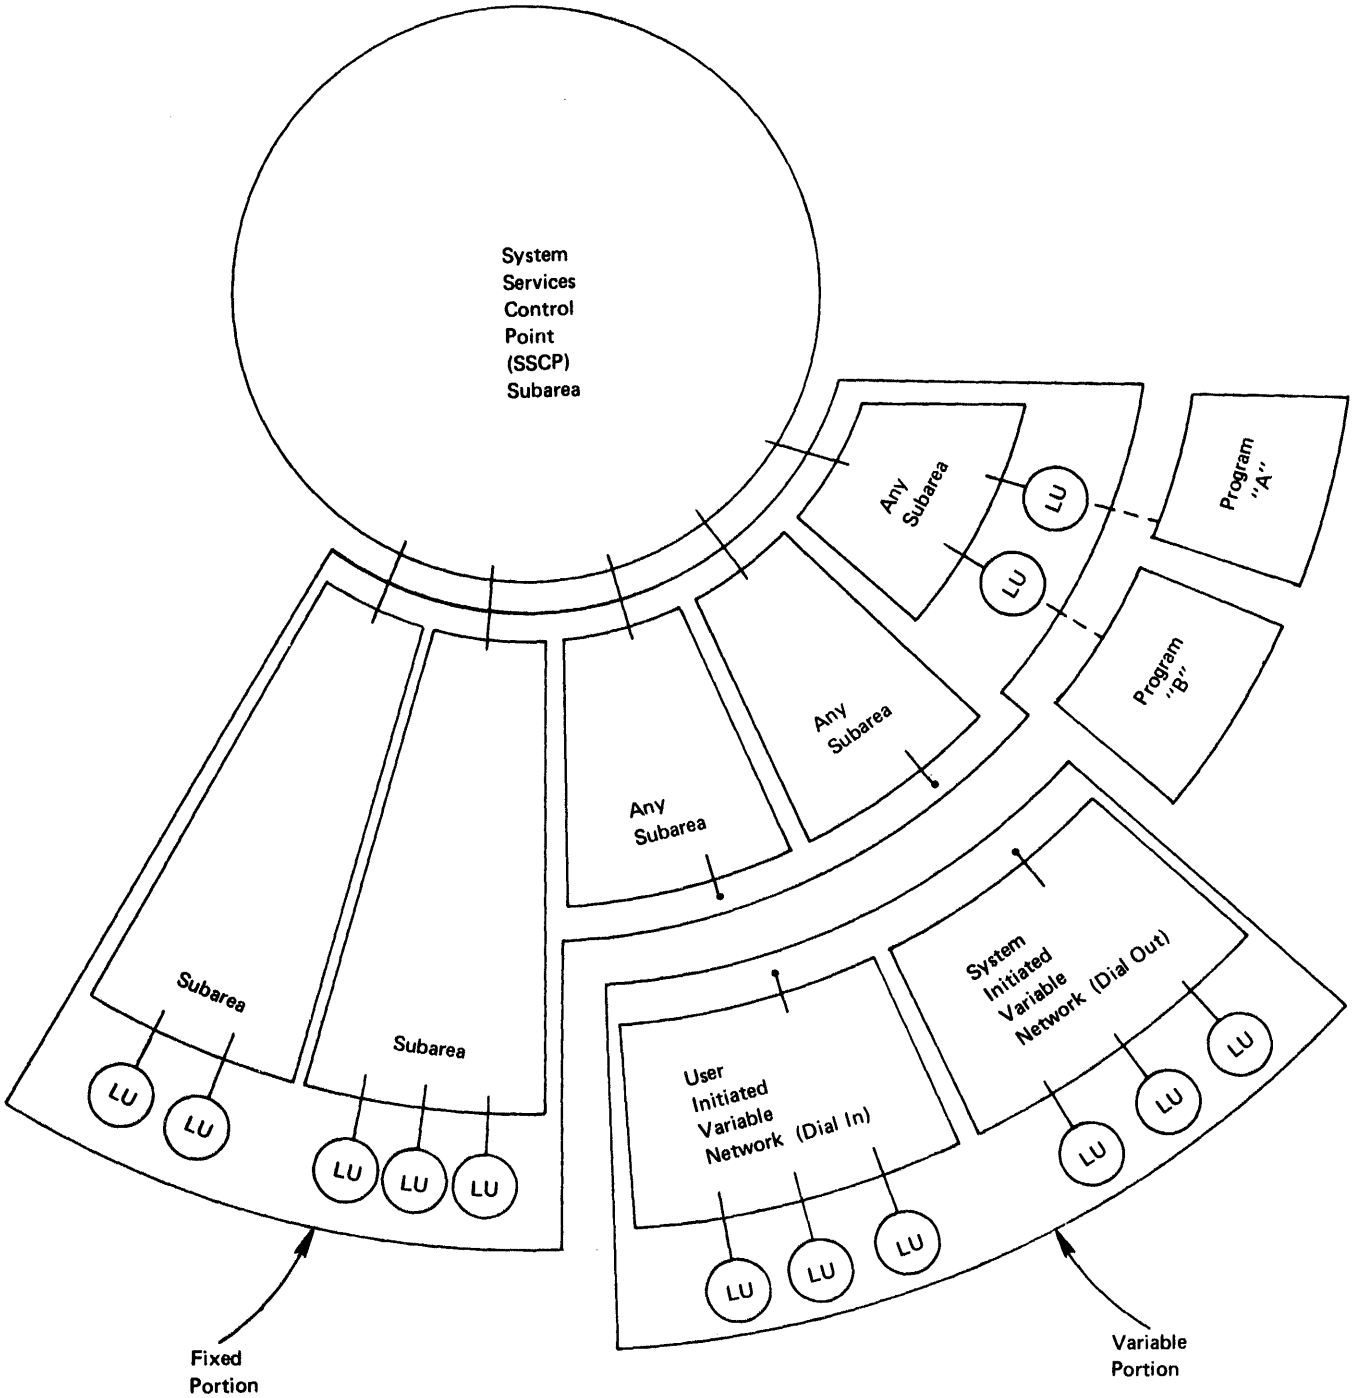

Figure 3-2. Logical Configuration of Subareas

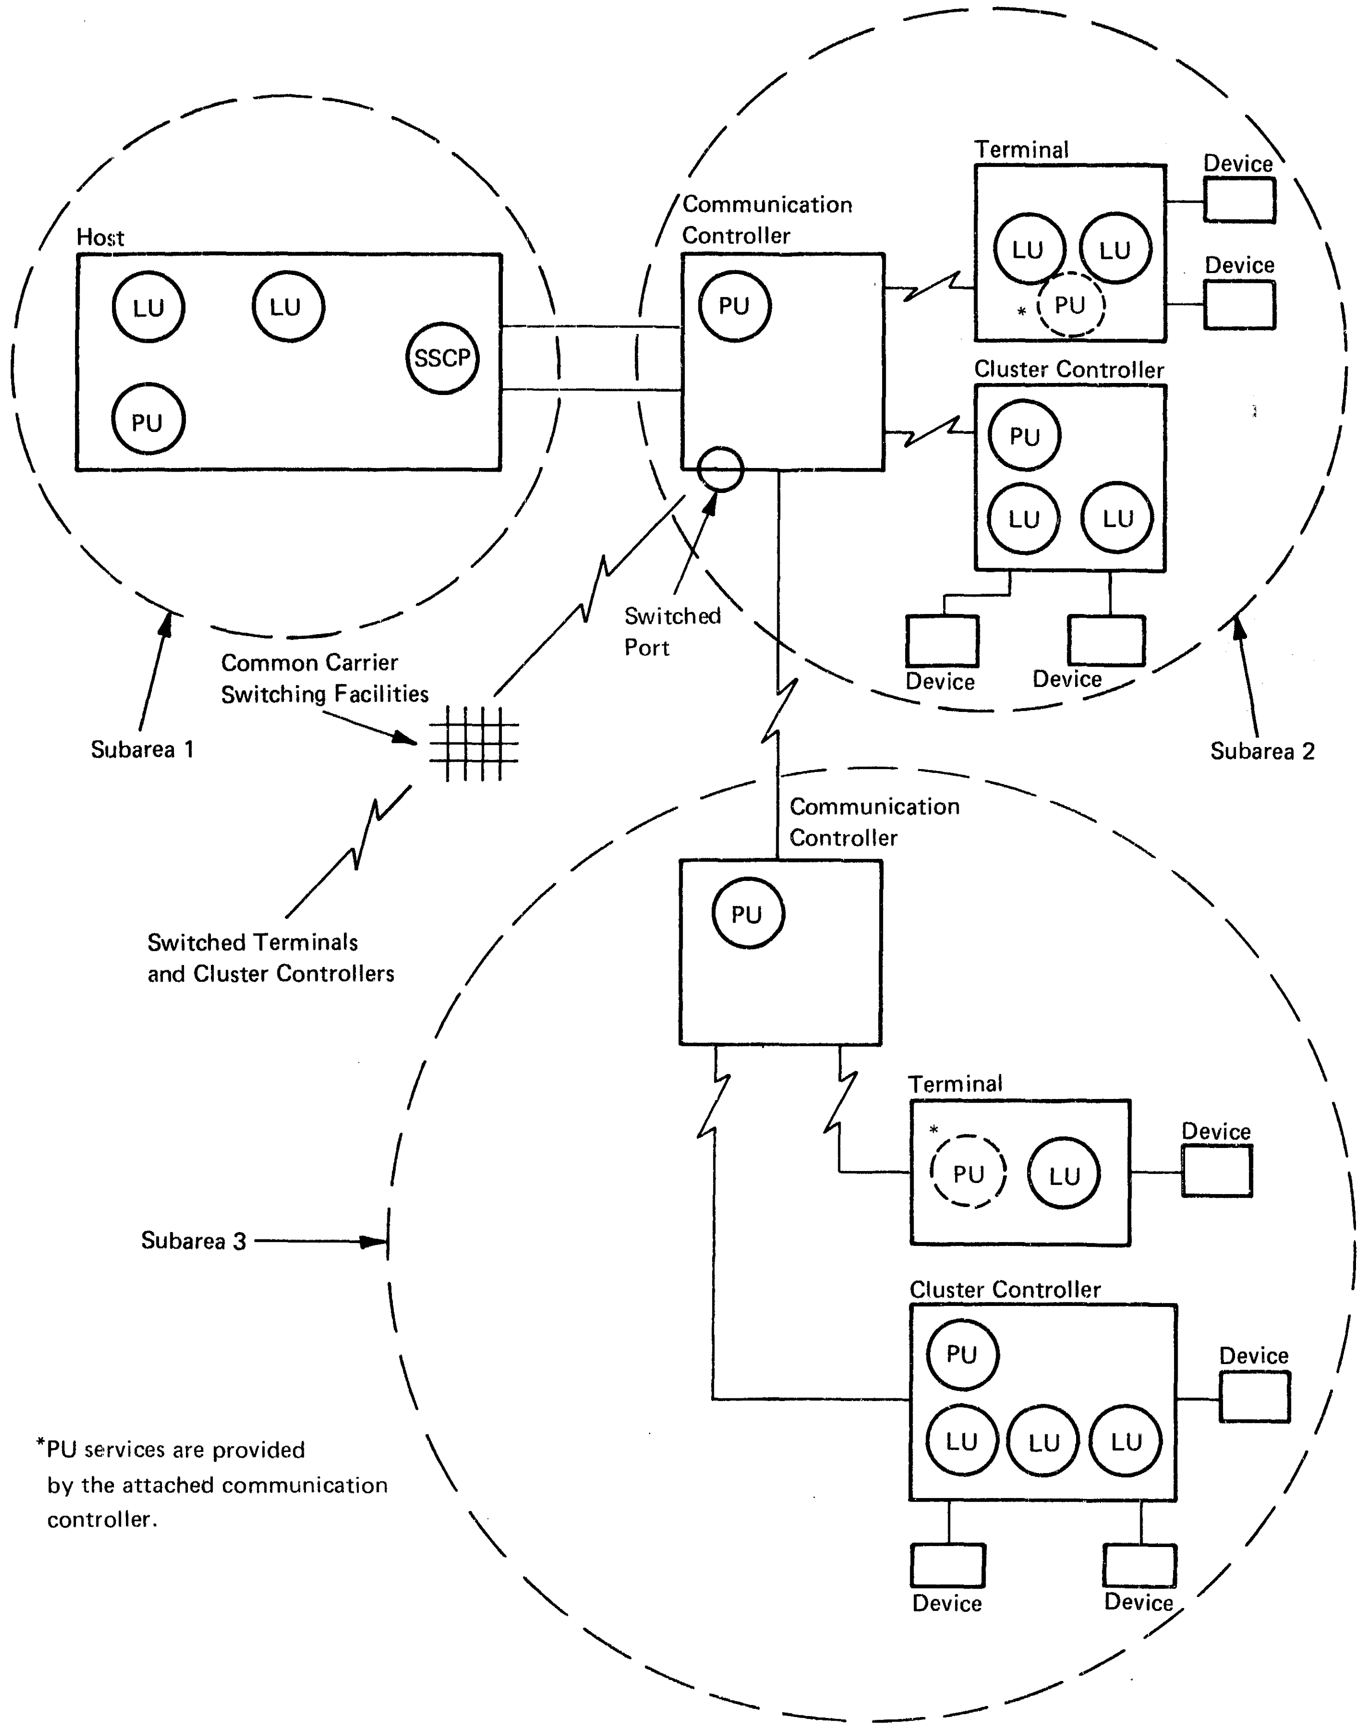

Figure 3-3. Physical Configuration of Subareas

Physical Configuration

Types of Network Nodes

The physical configuration of a network is defined in terms of four node types: host, communication controller, cluster controller, and terminal.

It should be noted that *devices* are physical entities that exist in the physical configuration. However they do not have network addresses and do not exist in the configuration information used by the SSCP. Devices are physical resources that are controlled and allocated by the cluster controllers or terminals.

Figure 3-3 shows the types of connections defined. Host nodes and communication controllers can control a subarea. Figure 3-3 shows three subareas.

A *host* node is a multiple-purpose facility that houses the system services control point in addition to executing application programs, managing data bases, and so forth. Examples of a host node are System/370 computers with VTAM and DOS/VS, OS/VS1, or OS/VS2.

A *communication controller* node is dedicated to the task of controlling communication lines (and related resources such as buffers) in addition to performing the functions related to supporting one or more subareas. Examples of a communication controller node are the IBM 3704 or 3705 with NCP/VS.

Cluster controller nodes support the attachment of a wide spectrum of devices to satisfy the requirements of a broad range of end users. Cluster controllers have less network management capability than host nodes or communication controllers. Cluster controllers use a local form of addresses. Examples of a cluster controller node are the IBM 3601 and 3791.

Terminal nodes have the least network management capability of all network nodes and use a local form of addresses. An example of a terminal node is the IBM 3767.

A communication controller node may provide two types of facilities: *intermediate functions* and *boundary functions.* A communication controller node providing an intermediate function routes messages to other subareas based on full network address processing. A communication controller node providing a boundary function converts a full network address (outbound from the host) to a local address form for adjacent cluster controller or terminal nodes. Cluster controller and terminal nodes depend on the node to which they are attached for support in scheduling the flow of data within a session.

Any communication controller node to which a terminal node attaches also provides rudimentary physical unit services for the terminal node.

Data Link Configurations

The physical configuration of a network is also determined by the data links used to connect the nodes of the network. The data links are of two types: the System/370 data channel, and SDLC communication channels. SDLC, in particular, allows a rich selection of communication configurations, such as point-to-point, multipoint, and loops. For further details, consult the IBM SDLC General Information manual (GA27-3093).

Chapter 4. Communication System Protocol

Figure 4-1 illustrates the relationships of the SNA elements in support of end user communication. It shows the SNA-defined facilities used to transmit control information and data units between various elements of the communication system. This figure illustrates a fundamental concept of SNA: the symmetry of functional capabilities at each end of a session. SNA defines the mechanisms for communication between functional elements to support end user to end user communication. The flow of requests and responses (RUs) takes place between paired *function interpreters*—one at each end of the session. The connection point manager recognizes four types of function interpreter pairs: the function management (FM) pair, the pair associated with data flow control (DFC), the pair associated with session control (SC), and the pair associated with network control (NC).

A function request is generated in the form of an RU at one end of the session by a function interpreter, and sent to the other end. There it is given to the appropriate function interpreter, which interprets the request and manages its execution. Based on end-to-end control information that accompanies the request, the executing function interpreter may send a response back to the originating function interpreter.

During a session between NAUs, requests and responses may flow in either direction or in both directions concurrently.

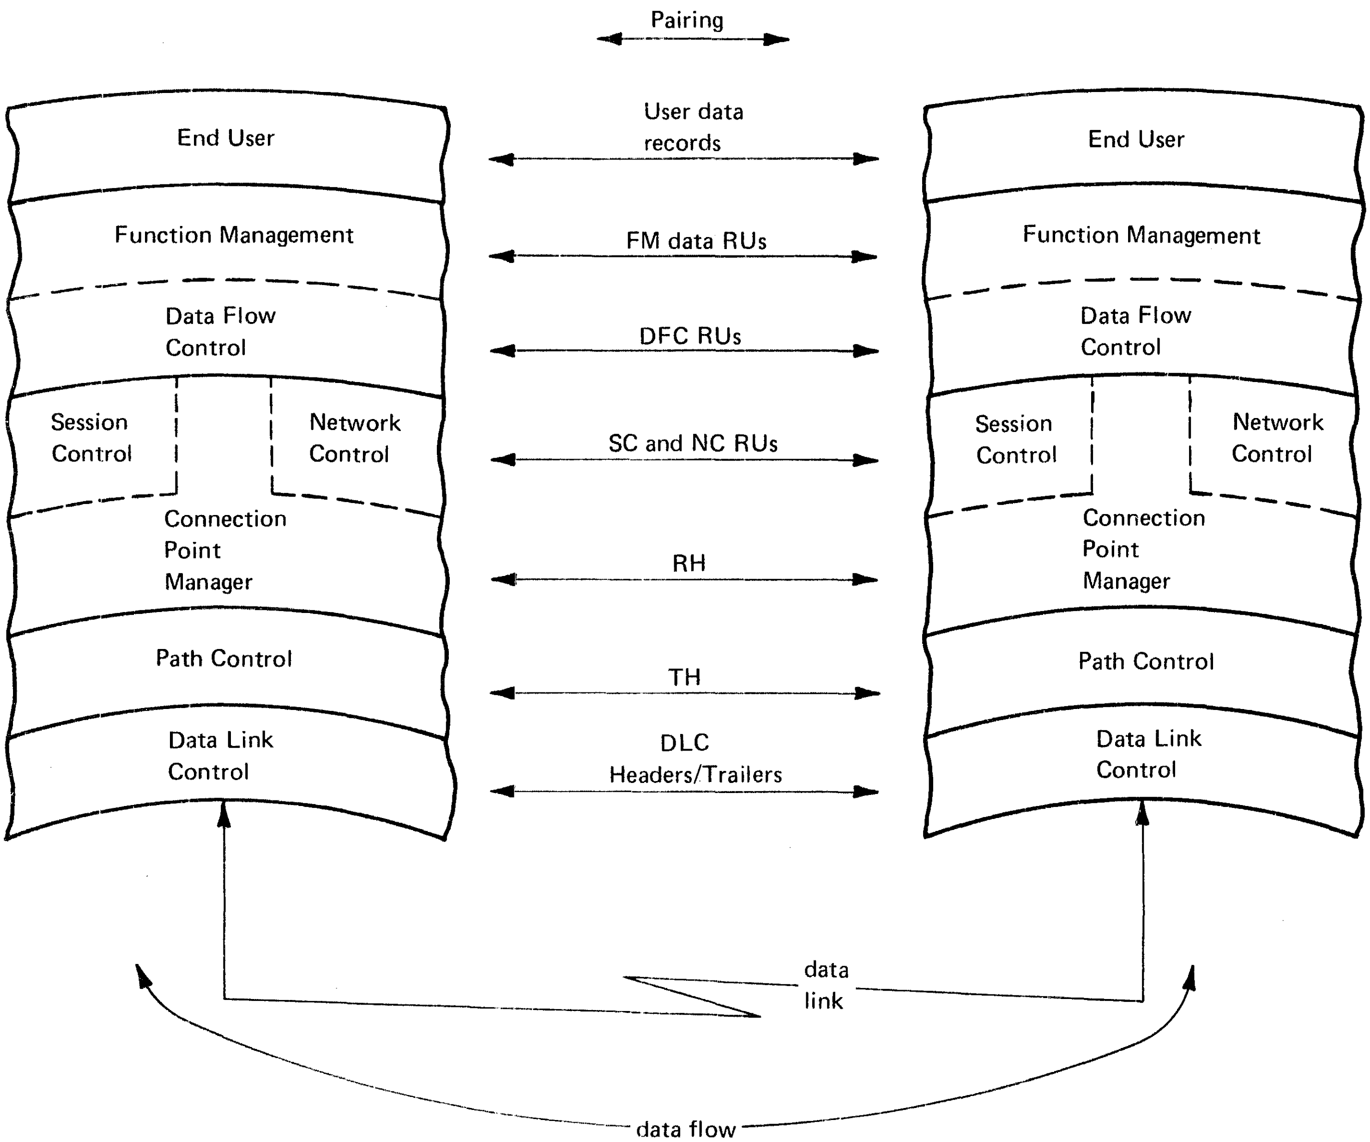

Figure 4-1. SNA Internal Communication

Function Management Protocol

The elements of the transmission subsystem provide for interactive or batch exchange of information between function management elements. SNA describes particular modes of operation available to FMs to assist in meaningful dialog. The concepts of immediate control, delayed control (with immediate or delayed request), and immediate and delayed response are defined. Also defined is the concept of chains of RUs.

Immediate and Delayed Control Mode

The terms *immediate* and *delayed control* denote the type of network synchronization expected by the issuer to transmission services. Immediate control mode means that the issuer will send a single RU and wait for a response before sending another. Delayed control mode means that the issuer may send many requests before waiting for a response. Two options are available within the delayed control mode: *immediate request* and *delayed request.* Immediate request denotes that the issuer may send a number of requests; however, only the last of these requests may indicate a *definite response.* That is, once a request requiring a definite response has been sent, no more requests may be issued until the definite response has been received. Delayed request mode allows the issuer to send multiple requests, each soliciting any form of response, without waiting for intervening responses.

Immediate and Delayed Response Mode

Immediate and *delayed response* modes denote the manner in which the receiver of a request returns a response. Immediate response mode indicates that responses are to be returned in the same order in which the requests were received. Delayed response mode indicates that the receiver may accept several requests before responding, and that the responses to these may be returned in an order different from that in which the requests were received.

Forms of Response

An FM interpreter may elect to be notified of the state of its requests within the receiving FM interpreter. The FM interpreter can specify that a *definite response* be returned. The specification of *exception response* enables the FM interpreter to ask that responses be returned only on detection of error conditions. These capabilities are implemented via control bits in the RH portion of a BIU (RH/RU).

Chaining

The issuer of the request passes requests into the network in the form of request units. These units of data are delivered by the transmission subsystem to the designated function interpreter, which may send back a response (response unit) to the function requester. The issuing function requester may specify a relationship, called *chaining,* for consecutive RUs. A chain may consist of a single RU or may consist of several consecutive RUs. Every RU belongs to one and only one chain, which has a well-defined beginning and end indicated via control bits in the RHs for the RU chain. Chains for a given function requester may not intermingle; one must stop before the next begins. The fundamental characteristic of a chain is that if one of its RUs cannot be processed, then the entire chain must be discarded. Error recovery action is done using the chain as the basic recoverable unit.

Potential ways to use chaining are:

1. RU consists of a line of keyboard input: a chain consists of several such lines, which together form a transaction input.
2. RU consists of a line of printer output: a chain consists of several such lines, which together form a message.

Function management chains are processed by IBM data flow protocol as if they were independent of each other. That is, a chain found to be in error is discarded. The next chain received is passed to the FM interpreter for processing. If the successive chains must be sequentially processed, it is necessary to synchronize at the end of each chain with a request for a definite response and to run in immediate request mode.

Data Flow Protocol

Data flow control elements provide a set of control functions that the FMs can use to control the flow of FM data within a session. A set of commands enables FMs to manage aspects of the data flow such as quiescing the flow under various conditions, and assisting in the recovery procedures between NAUs.

Session Control Protocol

Session control is responsible for the resource allocation necessary for a session between two NAUs. A set of commands supported by session control enables two NAUs to establish a session, assists in error recovery, and in reestablishing a session after a catastrophic error, and supports requests from the SSCP to initiate a session.

Network Control Protocol

Network control functions are special functions used to communicate between adjacent CP Managers without formally establishing a session.

Chapter 5. Transmission Subsystem Data/Control Flow

The protocol for a session allows independent two-way flow of data between NAUs. The definition of the transmission subsystem data and control flows encompasses the FM pair (FM/FM), and control function pairs (DFC/DFC, SC/SC, NC/NC). The connection point manager coordinates the concurrent session management control and FM data flow.

Data Flow (FM Data Pair)

When a session is established between two NAUs, one NAU acts as a primary and the other as a secondary (data flow is symmetric; control flow is asymmetric). Two *data flows* are established: primary to secondary, and secondary to primary.

The primary NAU sends requests and receives responses (to its requests) on the primary-to-secondary flow. The secondary NAU sends requests and receives responses (to its requests) on the secondary-to-primary flow.

The physical implementation of the primary-to-secondary flow is concerned with transmitting data units from an origin to a destination. This physical view groups requests and responses issued by the primary in one flow and requests and responses issued by the secondary in a second flow. Another characteristic of requests and responses is that of processing order. Each transmission subsystem data flow (primary to secondary, secondary to primary) is defined such that requests (or responses) are processed in the same order in which they were entered into the transmission subsystem. The term *normal flow* is used to describe this requirement.

Control Flow (DFC, SC, NC Pairs)

The set of control functions defined by the DFC, SC, and NC element pairs consists of requests and responses that affect the FM data flow. Some control function requests and responses must be processed in the same order in which they are entered into the transmission subsystem with respect to the FM data requests and responses. These control functions therefore are transmitted on the normal flow and are processed in line with FM data requests and responses.

Other control function requests and responses must not be held up by preceding normal flow requests and responses. An independent flow, the *expedited flow*, is defined, whose requests and responses can be given priority over normal flow elements. A control indicator in the TH is used to identify the selected flow for an RU. The expedited flow is separate from and controls the normal flow. The state of the normal flow (e.g., active, quiesced) has no effect on the capability of control functions to flow on the expedited flow. This allows the CP Manager to process the expedited flow when there is a stoppage in the normal flow.

Sequence Numbers

The normal flows utilize a sequence number generator (one pair per session) located in the CP Manager for the NAU issuing the request. All normal flow *requests* for a session are assigned sequence numbers. Each expedited flow generally has a limited request flow and therefore requires only a simple form of identification. A unique identification number is used for outstanding requests in each expedited flow.

Pacing

Pacing is a mechanism that permits a receiving CP Manager to control the rate at which it receives normal data flow requests. This includes both function management (FM) data and control functions flowing on the normal flow. It is used in cases where the sending FM can generate requests either faster than the receiving FM can process them, or faster than the network can transmit them. When a CP Manager exists within the path between the two end CP Managers, staged pacing can occur; i.e., sending CP Manager to middle CP Manager, and middle CP Manager to receiving CP Manager.

The transmission subsystem allows the sending CP Manager to send n normal flow requests before needing a pacing response to indicate that the next n requests can be sent. The receipt of a pacing response by the sending CP Manager will allow the next n requests to be sent. To allow a steady flow of requests, without a long pacing response delay after n requests have been sent, another parameter m is defined. This allows the receiving CP Manager to generate and send its pacing response after m (m less than or equal to n) of the n requests have been received.

SNA Control Information

SNA conveys control information concerning the management and operation of the network via three basic mechanisms: the transmission header (TH); the request-response header (RH); and reserved parts of particular request-response units (RU). The following is a general discussion of the TH, RH, and RU.

Transmission Header (TH)

Each unit of information (basic information unit or segment of a BIU) that is transferred through the network must have associated with it a transmission header containing the control information required by path control for handling of the BIU. The following control information is encoded in the TH:

1. Format identification field (FID) defines the subsequent format of the TH and the type of TH fields involved with the transmission. FID 1 is used between the host and communication controllers and between two communication controllers. FID 2 is used between communication controllers with boundary function and cluster controllers. FID 3 is used between communication controllers with boundary function and terminals. See Figure 5-1.
2. Mapping, sequence number, and data count fields are used in support of segmenting BIUs into transmission units.
3. Primary to secondary and expedited flow fields enable the specification of the data flow for this transmission.
4. Destination address field (DAF) and origin address field (OAF) contain the network addresses or local addresses involved in the session to which a transmission applies.

Request-Response Header (RH)

The request-response header contains control information to assist the connection point manager in the routing of RUs. Some data flow control information is also passed in the RH.

The following control information is encoded within the RH:

1. Routing information (FM data, DFC, SC, NC).
2. Information concerning chaining.
3. Indicators for form of response desired (definite, exception, none).
4. Indicators to specify unique requests and responses relating to sense data and pacing.
5. Indicators for DFC information.

Request/Response Unit (RU)

The request-response unit normally contains user data but may contain control information to assist in the routing of particular types of messages through the network. Network RU formats are defined for all transmission control components and network services.

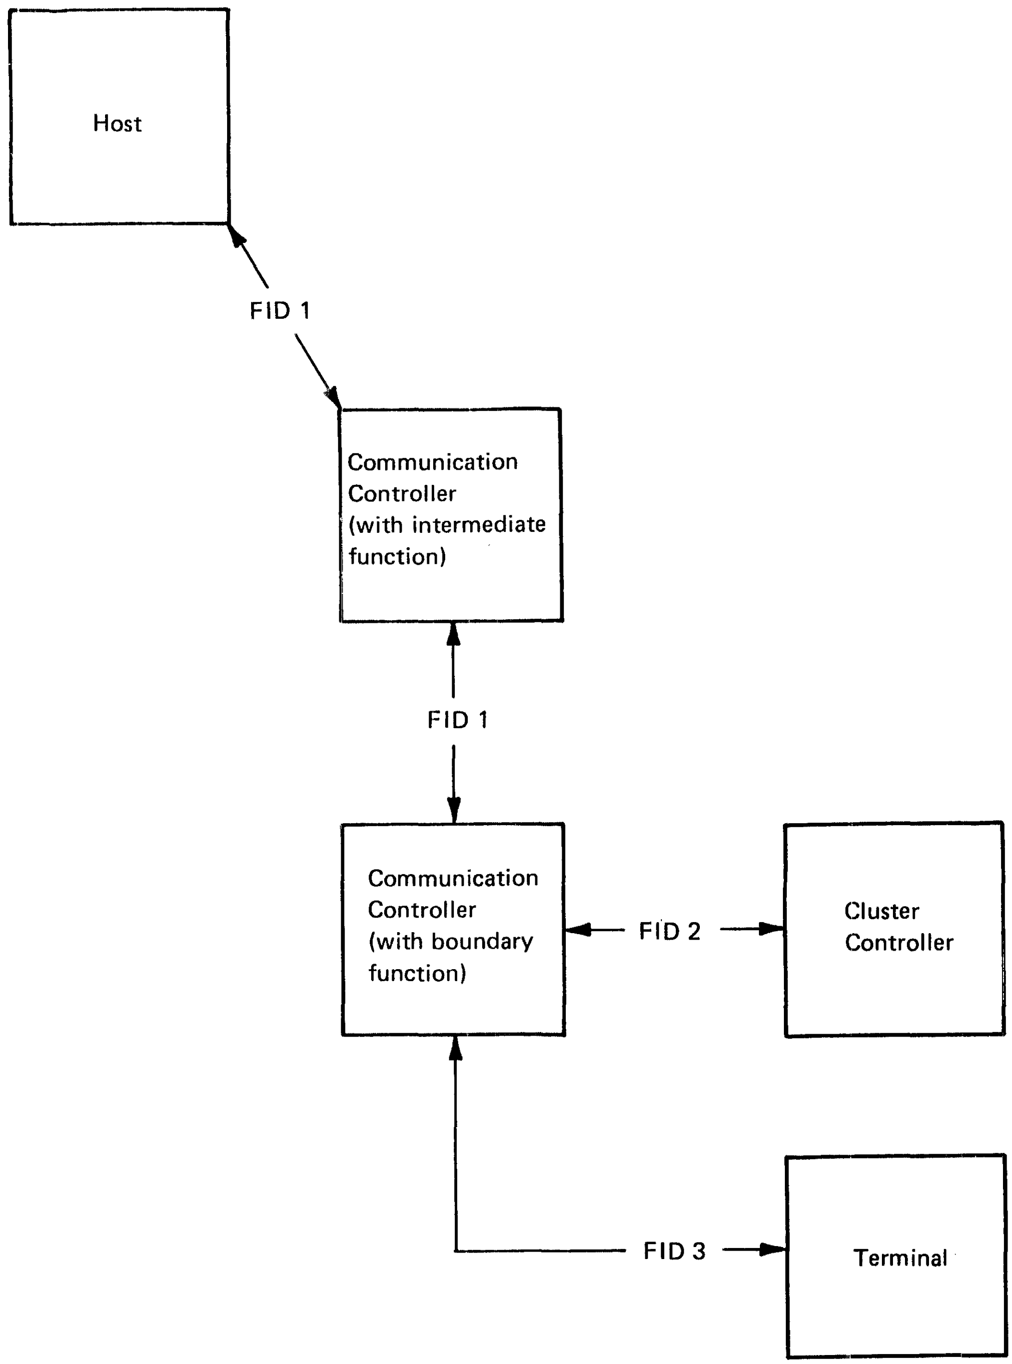

Figure 5-1. FID Conversion

Abbreviations

BIU	basic information unit
BLU	basic link unit
BTU	basic transmission unit
CP	connection point
DAF	destination address field
DFC	data flow control
DLC	data link control
FID	format identification
FM	function management
LU	logical unit
NAU	network addressable unit
NC	network control
NS	network services
OAF	origin address field
PC	path control
PIU	path information unit
PS	presentation services
PU	physical unit
RAS	reliability, availability, and serviceability
RH	request-response header
RU	request-response unit
SC	session control
SDLC	Synchronous Data Link Control
SNA	Systems Network Architecture
SSCP	system services control point
TC	transmission control
TH	transmission header

Index

System Network Architecture
General Information

READER'S COMMENT FORM

Order No. GA27-3102-0

Your views about this publication may help improve its usefulness; this form will be sent to the author's department for appropriate action. Using this form to request system assistance or additional publications will delay response, however. *For more direct handling of such request, please contact your IBM representative or the IBM Branch Office serving your locality.*

Possible topics for comment are:

Clarity Accuracy Completeness Organization Index Figures Examples Legibility

Cut or Fold Along This Line

What is your occupation?____

Number of latest Technical Newsletter (if any) concerning this publication: ______________________

Please indicate in the space below if you wish a reply.

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments.)

Cut or Fold Along Line

Your comments, please . . .

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. Your comments on the other side of this form will be carefully reviewed by the persons responsible for writing and publishing this material. All comments and suggestions become the property of IBM.

Fold Fold

First Class
Permit 40
Armonk
New York

Business Reply Mail
No postage stamp necessary if mailed in the U.S.A.

Postage will be paid by:

International Business Machines Corporation
Dept. E01
P.O. Box 12195
Research Triangle Park
North Carolina 27709

Fold Fold

System Network Architecture General Information (File No. S370-09) Printed in U.S.A. GA27-3102-0

IBM®

International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, New York 10604
(U.S.A. only)

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
(International)