GC30-3058-1
File No. S370-30

ACF/NCP/VS

Network Control Program
System Support Programs

Program Product

General Information

IBM

GC30-3058-1
File No. S370-30

ACF/NCP/VS

Network Control Program
System Support Programs

Program Product

General Information

Program Numbers: 5735-XX1
5735-XX3

Releases 2, 3

IBM®

Second Edition (May 1979)
This is a major revision of, and makes obsolete, GC30-3058-0. This edition applies to Releases 2 and 3 of the ACF/NCP/VS program product (Program No. 5735-XX1) and System Support Programs for ACF/NCP/VS program product (Program No. 5735-XX3) and to all subsequent releases and modifications unless otherwise indicated in new editions or Technical Newsletters.

Changes are periodically made to the information herein; before using this publication in connection with the operation of IBM systems, refer to *IBM System/370 Bibliography,* GC20-0001, for the editions that are applicable and current.

Publications are not stocked at the address given below; requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

This publication has been written by the IBM System Communications Division, Publications Development, Department E02, P.O. Box 12195, Research Triangle Park, North Carolina 27709. A reader's comment form is provided at the back of this publication. If the form has been removed, comments may be sent to the above address. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

Preface

This publication provides an introductory description of the types of data communication networks[1] that are accommodated by Advanced Communications Function for Network Control Program/Virtual Storage (ACF/NCP/VS), Advanced Communications Function for Telecommunications Access Method (ACF/TCAM), and Advanced Communications Function for Virtual Telecommunications Access Method (ACF/VTAM).

This publication summarizes the specific functions provided by Release 2 and Release 3 of ACF/NCP/VS and contains information on the System Support Programs for ACF/NCP/VS and planning for use of ACF/NCP/VS.

This publication is directed primarily to data processing managers and data communication network designers intending to install or upgrade an ACF/NCP/VS-based data communication network or to consolidate existing networks.

ACF/NCP/VS, ACF/TCAM, and ACF/VTAM are program product versions of their system control program (SCP) equivalents: NCP/VS, TCAM, and VTAM. The program product versions contain significant enhancements to the SCP versions.

The organization of this publication is as follows:

Chapter 1 reviews some basic elements of data communication networks based on IBM System/370 processors and describes the functions of network control programs in general.

Chapter 2 introduces some general concepts of Systems Network Architecture (SNA) as they apply to single-domain networks (networks with one control point).

Chapter 3 continues the discussion of SNA and ACF as they apply to multiple-domain networks (networks with more than one control point) and networking (the management of multiple-domain networks).

Chapter 4 describes in general terms the System Support Programs required for use of ACF/NCP/VS.

Chapter 5 gives a brief overview of each of the enhanced capabilities afforded by ACF/NCP/VS Release 2.

Chapter 6 gives a brief overview of each of the enhanced capabilities afforded by ACF/NCP/VS Release 3.

---

[1]The term *network* has at least two meanings. A *public network* is a network established and operated by common carriers or telecommunication Administrations for the specific purpose of providing circuit-switched, packet-switched, and leased-circuit services to the public. A *user-application network* is a configuration of data processing products, such as processors, controllers, and terminals, established and operated by users for the purpose of data processing or information exchange, which may use transport services offered by common carriers or telecommunication Administrations. *Network*, as used in this publication, refers to a user-application network.

Chapter 7 lists customer responsibilities and other information needed by those planning to use ACF/NCP/VS.

Chapter 8 contains preinstallation planning information in the form of questions to be answered by the installation manager (or other individual responsible for planning the development of an ACF/NCP/VS-based network) during the preinstallation phase of the project.

Appendixes list the communication stations and systems accommodated by ACF/NCP/VS; summarize some factors relating to use of ACF/NCP/VS in a remote communications controller and to use of the partitioned emulation programming extension; and provide a glossary of terms and abbreviations.

Prerequisite to use of this publication is a general knowledge of data communication. A knowledge of the basic concepts of Systems Network Architecture is also helpful. General information on the hardware aspects of the IBM 3705 Communications Controller in which ACF/NCP/VS is executed may be found in *Introduction to the IBM 3704 and 3705 Communications Controllers,* GA27-3051.

For general information on the capabilities of ACF/TCAM and ACF/VTAM, see *ACF/TCAM General Information: Introduction* (GC30-3057), or *ACF/VTAM General Information: Introduction* (GC27-0462), respectively.

For information on Systems Network Architecture, see *Systems Network Architecture General Information,* GA27-3102.

For an overview of Advanced Communications Function and its role in networking, see *Introduction to Advanced Communications Function,* GC30-3033.

A description of the communication control and problem determination facilities offered by three related products (Network Operation Support Program, Network Communication Control Facility, and Network Problem Determination Application may be found in:

- *Network Operation Support Program General Information*, GC38-0251
- *Network Communcations Control Facility General Information*, GC27-0429
- *Network Problem Determination Application General Information*, GC28-0942

Note: In this publication, long numbers are represented in metric style. A space is used, instead of a comma, to separate groups of three digits.

Contents

Figures

Summary of Changes

This second edition principally adds an overview of the major new network capabilities offered by Release 3 of ACF/NCP/VS. This overview appears as Chapter 6. The former Chapter 6, "Planning for Use of ACF/NCP/VS," is updated with information on Release 3 of ACF/NCP/VS and is renumbered Chapter 7. The former Chapter 7, "Preinstallation Planning," is renumbered Chapter 8.

Information on host processor requirements for ACF/VTAM and ACF/TCAM, and on the programming subsystems that these access methods support, is deleted from Chapter 1. This information is available in *ACF/VTAM General Information: Introduction* GC27-0462, and *ACF/TCAM General Information: Introduction,* GC30-3057.

To Appendix A, "Communication Stations and Systems," are added several new communication products.

Other minor changes, corrections, and clarifications appear throughout this edition.

Changes in the text or illustrations are marked with a vertical line to the left of the changes except in Chapter 6, which is a new chapter.

Chapter 1. Introduction to Networks and ACF/NCP/VS

This chapter reviews some basic elements of data communication networks based on IBM System/370 processors and introduces some concepts, such as distributed function and resource sharing, that are beginning to characterize networks based on the IBM Systems Network Architecture. The chapter then introduces three major program products—ACF/NCP/VS, ACF/TCAM, and ACF/VTAM—in the IBM Advanced Communications Function.

This publication discusses ACF/TCAM and ACF/VTAM only in terms of their relationship to ACF/NCP/VS; for detailed introductory information on these program products, see the *ACF/TCAM* and *ACF/VTAM General Information: Introduction* publications (see Preface for order numbers).

Regardless of the data processing application for which it is installed, a data communication network of the kind this publication describes is a system of interconnected terminals, controllers, and processing units used to transmit information between locations in the network. At various points in the network the information may be processed or its form changed, as dictated by application requirements.

Figure 1-1 shows an example of a simple data communication network having each of the kinds of elements typically found in System/370-based networks; a host processor (the System/370), an IBM 3705 Communications Controller, a group of terminals and the shared controller (called a cluster controller) that links them to the 3705, and two terminals directly linked to the 3705. In this simple network that Figure 1-1 illustrates:

- The host processor (A) contains (in part) an operating system, a communication access method, application programs, and logical units that represent the application programs to the network. (In this publication, *access method* always refers to a communication access method: ACF/TCAM or ACF/VTAM, or their non-ACF equivalents.) The access method supervises the operation of the network and interacts with the network control program in the communications controller to transmit message traffic between the host processor and the network.
- The communications controller (B), executing a network control program, manages the details of line control and routing of data through the network. The program can route data to the host processor and can route data to and from the cluster controller (C) or a terminal (D). (In a typical network the communications controller would route data to and from many cluster controllers, terminals, and other communications controllers.)
- The cluster controller (C) and the terminals (D) each attach one or more input/output units to the network. Cluster controllers and terminals differ primarily in the amount of processing capability they have and in the complexity of their configurations. Cluster controllers and terminals are mainly concerned with fulfilling application requirements rather than with performing network routing and control functions.
- The ultimate source or destination of information flowing in the network is called an end user (E). End users may be individuals such as terminal operators, physical device media such as tapes or cards, or application programs. End users are not identified as such to the network; they gain access to the network and its resources through logical units (F). Logical units are assigned names that uniquely identify them to the network elements (access

methods and network control programs) responsible for routing message traffic through the network.

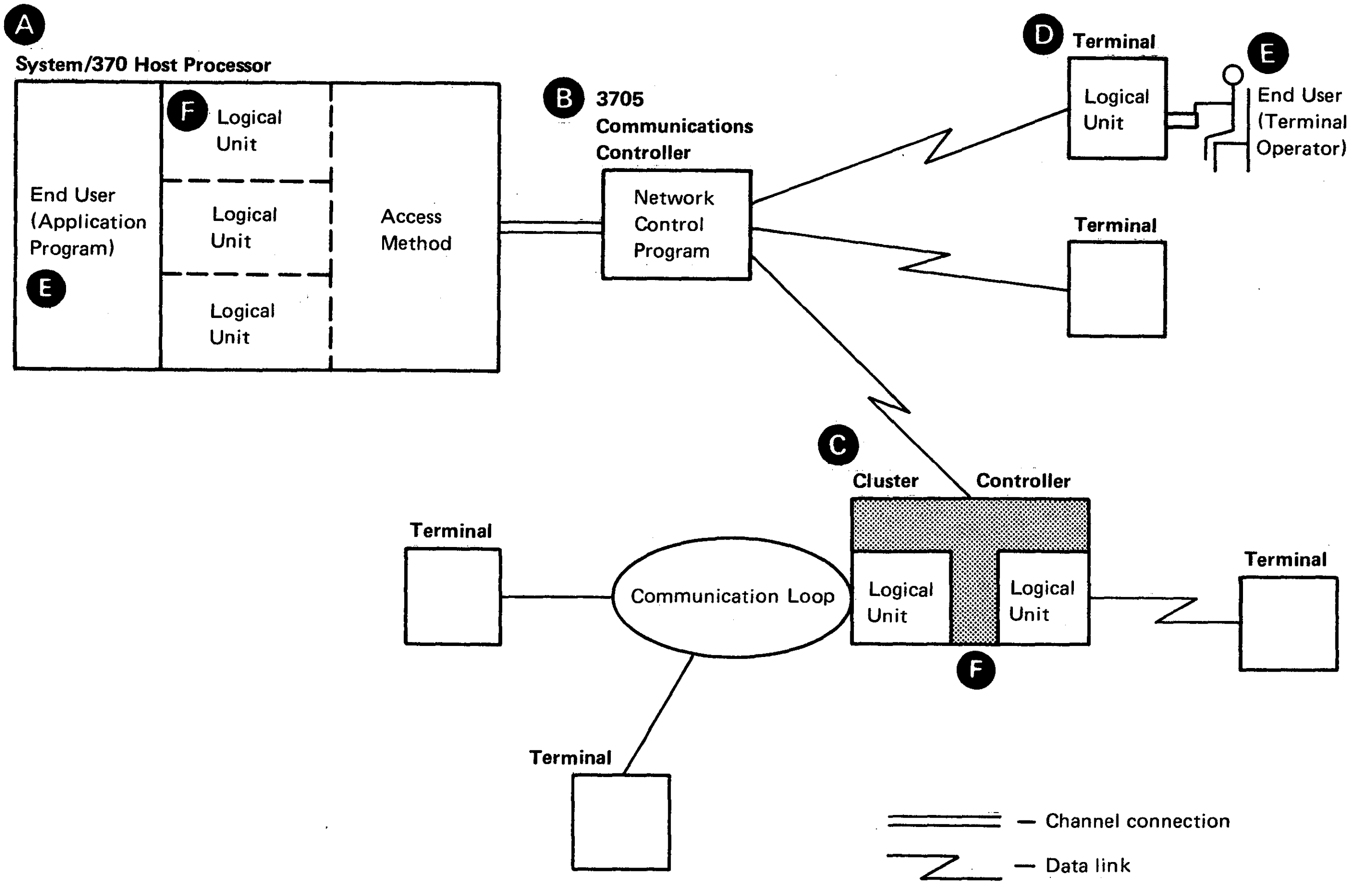

Figure 1-1. Example of a Data Communication Network

Distributed Function

An important concept of new networks is that of *distributed function.* This is the idea that processing of data can be done at many points in the network, in addition to the central site. Some network devices are programmable, thus allowing functions that were previously done at the central site to be performed at other points in the network. Distributing function in this manner is advantageous to the network in several ways:

Improved Response Time: In networks without distributed function, each message (request) must be forwarded to the host processor and a response returned to the originator of the message. The result could be slow turnaround time due to the distance that the request had to travel and the heavy processing load at the host processor. Distributed function allows many requests to be processed at their point of origin in less time. Therefore, the throughput of the entire network may be increased; it may exceed the throughput capability of the host processor.

High Availability: In a network without distributed function, a host processor failure incapacitates the entire network for as long as the host is not available. In a network with distributed function, programmable stations and cluster controllers may continue to operate effectively even when the host processor or communication line fails. Distributed function also increases availability of the network by giving each network component some capability for error recovery.

Less Duplication of Effort: In networks without distributed function, frequently the only backup capability when communication interrupted is manual handling of transactions at the remote location. When communication with the host processor is reestablished, the transactions have to be reentered in order to be permanently stored in the system. However, with distributed function, transactions that take place while the host processor or communication line is inoperative can be stored on an input/output storage device (if present) at the remote location for later transmission to the central site, thus eliminating duplicate operator effort.

Resource Sharing

Resource sharing is an important concept that allows the users of a network to share such resources as application programs, network control programs, and communication facilities. For example, stations attached to the same communication facility can be connected to different application programs. When an application program requests data transfer, the communication facility is used. Thus, the communication facility is shared among the stations and the application programs.

Communication Facilities in a Network

Data transferred between components of the network is transmitted by means of *communication facilities* which can be either switched, nonswitched, or a combination of the two.

A *switched* communication network establishes a temporary connection between a communications controller (or computer) and a remote station upon demand by the user. As in a telephone network, the actual path for a given transmission is not fixed but is automatically selected from a number of possible paths by telephone switching equipment. A switched line is point-to-point, since only one remote station on a line may be connected during a given call. Dialing can be performed at the host processor, communications controller, cluster controller, or at a remote station; the called station can answer either manually or automatically. Not all of these options are available for all types of line configurations and remote stations, or in all countries.

A *nonswitched* communication line (also called a *link)* connects the communications controller to the stations on the line for either continuous or recurring periods. A nonswitched line is *point-to-point* if it connects the controller to a single remote station, and *multipoint* if it is capable of connecting the controller to more than one station.

IBM's Systems Network Architecture (SNA) implements the concepts of distributed function among the physical components of a network and the sharing of resources among the users of the network. SNA is a design solution for a total network; it is an architecture that, when used in the design of devices and programs, allows application programs to be independent of devices and device attachments.

SNA networks accommodate a variety of SNA and non-SNA stations (terminals, cluster controllers, communications controllers and computers). An *SNA station* is capable of performing functions defined by the architecture. If the station is a terminal directly attached to the network, a logical unit within the terminal represents the terminal to the network control program and the access method. If the station consists of a cluster controller and one or more attached terminals, logical units within the cluster controller represent the terminals to the network control program and the access method.

A *non-SNA station* uses the start-stop or binary synchronous (BSC) line discipline to communicate with the host processor. These stations do not have logical units but the host processor and the communications controllers perform the functions that are necessary to communicate with these devices in an SNA network. The term station is used in this publication to mean a logical unit, a device that is represented by a logical unit, or a non-SNA device, unless a distinction is required for clarity.

Communications Controllers and the Network Control Program

The IBM 3705 Communications Controller is a programmed control unit designed to assume many of the line-control and processing functions for the network. The operation of a network with the 3705 communications controller is significantly different from one operating with an IBM 2701 Data Adapter Unit or with an IBM 2702 or 2703 Transmission Control. With the 2701, for example, the operation of the communication line for each station is controlled in the host processor; each line requires a separate subchannel at the host processor. With the 3705, a single subchannel can convey all data between the host processor and the communications controller; all communication lines to remote stations are controlled by the 3705.

Figure 1-2 shows a network in which all stations are directly connected to a single 3705 communications controller locally attached to a System/370 host processor by a data channel. More complex networks can include additional local communications controllers and can also include remote communications controllers. The latter are connected to a local controller by a communication line (SDLC link), rather than being channel-attached to a host processor. If the network configuration warrants, a 3705 can be equipped for both channel attachment to a host processor and remote connection to another 3705. This publication focuses on the functions the network control program serves in a network, independent of whether it is being executed in a local or a remote 3705.

In general, the descriptions of these functions assume program execution in a local 3705, since all SNA networks include one or more local 3705s. Differences between NCP functions in a local and a remote 3705 are summarized in Appendix B. (For information on the details of local and remote attachment of a 3705, see *Introduction to the IBM 3704 and 3705 Communications Controllers,* GA27-3051.)

SNA stations in the network use the synchronous data link control (SDLC) line discipline. Non-SNA stations may use either start-stop (asynchronous) or binary synchronous (BSC) line disciplines. The communication lines between the communications controller and the stations may be either switched or nonswitched, depending upon the configuration desired. There is no limit to the distance separating the stations from the communications controller. A nonswitched line is *multipoint* if it connects one or more stations equipped with the multipoint line control discipline to the communications controller. A nonswitched line is *point-to-point* if it connects a station equipped with point-to-point line control discipline to the communications controller. Regardless of the line discipline that a station uses, the 3705 transfers all data between the station and the host processor over a single data channel.

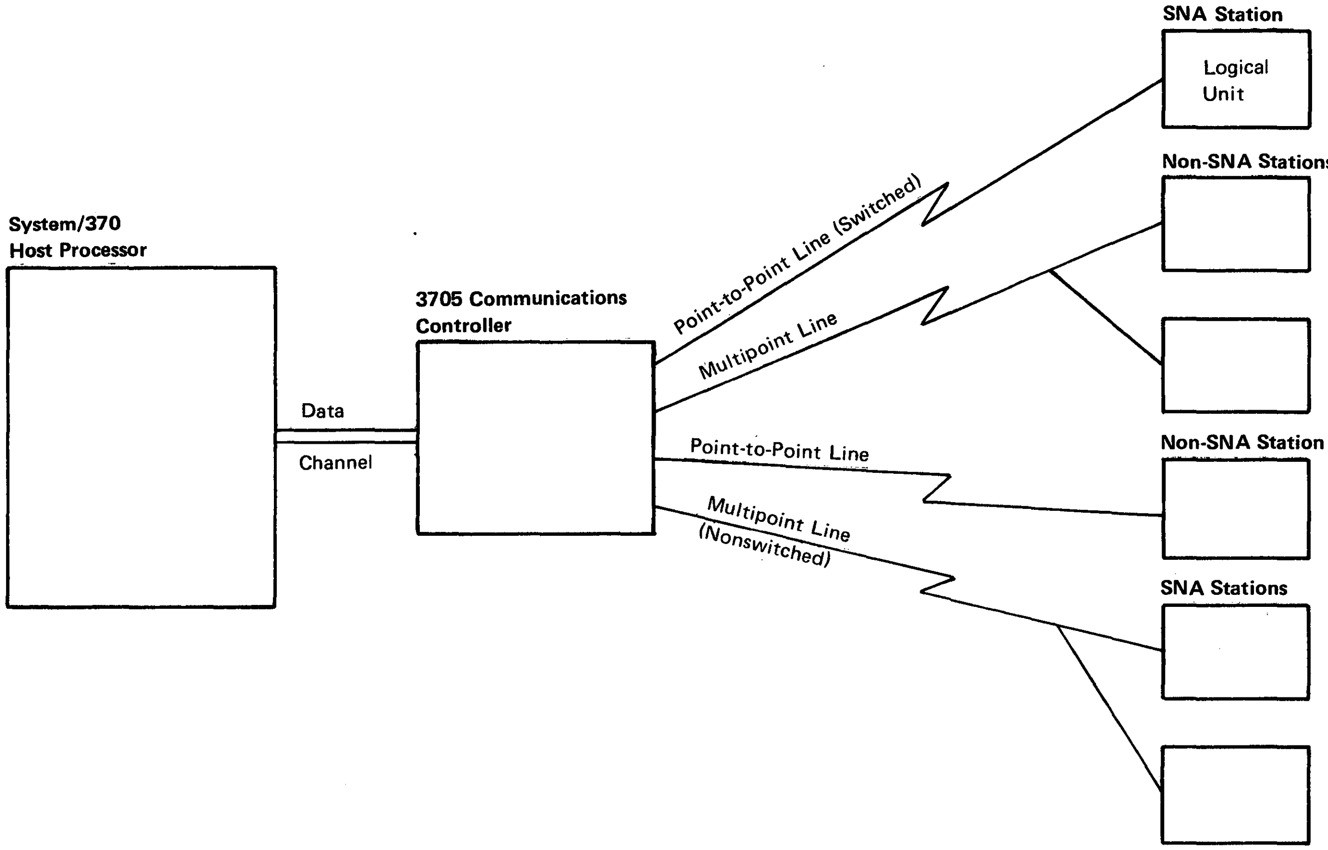

Figure 1-2. Network Showing Communication Lines and One Communications Controller

How the Network Control Program Communicates with the Access Method

The communications controller is physically attached to the host processor by a standard System/370 data channel over which the access method performs all communication with the rest of the network.

A basic operating premise of the relationship between the network control program (NCP) and the access method is that the NCP will be prepared to receive information from the access method as long as the NCP has buffers available, regardless of the availability of the line or station to transfer data at the moment. If the access method has buffers available, it reads data from the NCP each time that it writes data to the NCP. At other times, when the NCP has data for the access method, it presents an attention interruption over the channel to the host processor. The host processor cannot be forced to read from the communications controller, but it does acknowledge the attention interruption when ready to receive data by issuing a Read command to the controller. The NCP holds the data in its buffers until the access method reads it over the channel.

The NCP and the access method communicate over the channel by exchanging formatted information. This information contains control parameters and, optionally, message data. The control parameters direct the NCP to perform a specific operation. When the operation is completed, the NCP responds with corresponding control parameters, message data (if requested), and the completion status of the operation. Thus the access method directs the NCP, which in turn controls the network operation and provides the access method with its required data and resulting status information.

The NCP receives data from the access method and remote stations concurrently, as described above. The NCP collects the data in buffers and sends it to the intended destination when line availability and other conditions permit. Control parameters and data intended for the access method are sent to the host processor across the channel. Consequently, the access method directs the network at channel speeds, while the NCP and the communications controller are responsible for line control and data transfer.

Although freed of many realtime network responsibilities, the access method must still know the structure of all resources that it is to control in the network. This information is provided for both the access method and the NCP as a result of how the user specifies the network resources (lines, stations, etc.) when defining the network control program and the access method.

How the Network Control Program Operates the Communications Controller

Receiving its direction from the access method, the NCP operates the communications controller and transfers messages to their destinations. Because the controller is interrupt-driven, the NCP reacts to interrupts, which signal such events as:

- Data and/or control information arriving from the access method over the channel
- Data arriving from remote resources of the network
- Programmed interrupts and utility functions requesting attention
- Hardware errors occurring in the communications controller

To service the communications controller interrupts that signify data transmission events, the NCP uses a system supervisor, a channel adapter I/O supervisor, a communications interruption control program, error recovery facilities, and other services routines. These are the major programming components through which the NCP controls the internal operation of the communications controller. Because the controller is programmable, the NCP is able to manipulate data in the controller, rather than requiring that the host processor assume the burden of manipulating data.

The remainder of the network—communication lines, cluster controllers, terminals, terminal components, switching facilities, and so forth—all operate as the access method and the NCP direct. The means by which these programs effectively control message transmission between host processor and station include:

- *Connections:* establishing the physical path for communication and engaging in data transfers
- *Sessions:* allowing multiple concurrent sessions on each multipoint line by interleaving transmissions to separate stations
- *Switched network operation:* handling switching facilities for stations capable of dialing and answering
- *Polling and addressing:* executing line control sequences that result in data transfers

Other network services include load management, scheduling, and error recovery for remote resources and communication lines.

Functions of the Network Control Program

The network control program performs a wide range of functions for the data communication network. Certain functions are standard for any network control program; others are optional, selected as part of the program generation procedure.

Standard Functions

Standard functions of the network control program, working with the 3705 hardware, include those that any transmission control unit performs, such as recognizing and reacting appropriately to control characters, controlling communication-line time-outs, checking errors, and assembling characters received from a communication line into a buffer. Other standard functions that the network control program performs are as follows.

Communication Control

The network control program takes over most of the control of the communication lines from the communications access method. The standard communication control functions are:

- Polling and addressing of stations on multipoint communication lines.
- Dialing and answering stations over the switched communication network.
- Performing character service: The line attachment hardware in the communications controller interrupts the network control program either (1) when a character arrives over a communication line (the program then moves the character into a buffer), or (2) when an entire buffer is filled with characters from the line. The type of communication scanner in the controller determines whether the program is interrupted after each character is received or after the buffer is filled.
- Inserting and deleting control characters: The network control program inserts control characters at the beginning and end of each block of data when transmitting to a station and deletes them when receiving from a station.
- Translating character codes (BSC and start-stop lines only): As data arrives from a station, the network control program automatically translates it from transmission code into EBCDIC. Similarly, the program translates EBCDIC data into transmission code before transmitting the data to a BSC or start-stop station.
- Performing dynamic buffering: The network control program allocates buffers from controller storage as it receives data from a station or from the host processor. Upon accumulating an entire block of data (regardless of whether the last buffer is filled), the NCP transfers the data to the host processor. Blocking of data received from many points in the network maximizes utilization of the channel connection to the host processor.
- Selecting line speeds: Speed selection allows the network control program to change the transmission rate on a line equipped with IBM 3872 or 3875 modems. A command from the access method specifies whether the normal (high) rate or alternate (low) rate is desired. (A communication line whose performance has become too degraded for satisfactory transmission at the normal rate can often transmit satisfactorily at the lower rate.)

Error Recording and Diagnostics

The network control program maintains several types of error records and provides diagnostic capabilities.

- *Hardware- and program-check recording:* The program keeps a record of hardware and program checks, transferring the information to the host processor whenever possible. If transfer is impossible (for example, if the channel adapter fails), the program attempts to reset the error condition and allows automatic loading (IPL) of the 3705.

- *Permanent line error recording:* If normal error recovery procedures fail to recover from a transmission error, the network control program transfers to the host processor a record containing information about the error.
- *Statistics recording:* The program maintains for each station a count of the number of I/O operations that are executed and the number of temporary errors that occur for that terminal.
- *Dynamic panel display:* This function permits the operator to display storage areas, register contents, or control information on the control panel of the 3705.

Optional Functions

Many of the network control program functions are optional; they may be performed instead by the communications access method, or they may be omitted entirely, as specified by the user. The following functions are optional.

Block-Handling Functions

For binary synchronous and start-stop communications, the network control program can process blocks of data from either the station or the host processor via optional *block-handling routines.* These routines are:

- Date and/or time insertion
- Correction of text incorrectly entered from a station

Additional block-handling functions are possible through user-written routines. The user includes these routines in the program by coding a generation macro provided for that purpose. User-written routines must be assembled with the communications controller assembler.

Error Recovery and Diagnostics

The network control program can include the following optional error recovery and diagnostic functions.

- *Critical situation notification (BSC and start-stop stations only):* The network control program can notify stations when the host processor, channel, or local-remote fails via a user-written predefined message.
- *Address trace:* The operator can request through the 3705 control panel that the network control program record the contents of four variables (storage areas and/or registers) when a certain address in controller storage is accessed. This provides a dynamic trace facility for diagnostic purposes.
- *Online terminal testing:* The 3705 provides online terminal testing (OLTT) facilities. The network control program accommodates the OLTT functions by recognizing test requests from terminals and executing test routines constructed by an OLTT program in the host processor.
- *Online line testing:* Online line testing (OLLT) capabilities are available for SDLC links. The NCP executes test routines constructed by an executive program in the host processor.
- *Pause-retry:* When a transmission error occurs, the network control program tries to retransmit the data after a user-specified interval has elapsed. The user also specifies the maximum number of retries to be attempted for each station. This function is included for all stations unless the user specifies that no retries are to be made.
- *Switched network backup:* For certain kinds of BSC and start-stop stations, the user can specify an alternate path for communication over the switched

telephone network. The alternate path is used if the nonswitched point-to-point line that the station normally uses encounters an error condition from which error recovery procedures fail to recover.

- *Manual switched network backup:* This facility, an extension of the switched network backup facility, allows the console operator to call a station upon being informed that the regular nonswitched line to the station has failed. The operator enters a console message identifying the station to be contacted, and receives in response the identification of the switched backup line. The operator then places the call to the station; transmission to the station can then resume. Transmission continues over the backup line until the console operator reestablishes the regular nonswitched line connection. Manual switched network backup is for use when the equipment required for automatic calling is not available.

Miscellaneous Options

The network control program can include the following optional functions.

- *Channel attention delay option:* This option allows the user to specify an interval, in increments of 100 milliseconds, to be observed before the network control program presents Attention status to the channel. If an interval is specified, data arriving from the stations is stored in the network control program buffer pool until the interval elapses. Then all the stored data can be transferred across the channel with only one interrupt to the host processor, thus decreasing host processor overhead. If no interval is specified, each block of data is transferred as soon as it is processed by the network control program, requiring more frequent interrupts to the host processor. If the network control program receives enough data to fill all the allotted buffer space in the host processor before the specified interval elapses, it presents Attention status and status modifier to the channel immediately.
- *Verification of identification (ID) received from BSC stations:* This option is available for certain BSC stations that communicate over the switched communication network. The user provides a list of valid IDs for communication lines on which ID verification is to be used. The NCP compares the ID received against those in the list and allows the station to connect if the received ID matches one in the list. If it finds no match, the NCP can pass the information to the access method in the host processor, or it can break the connection, at the user's option. The ID verification option permits some IDs to be kept in the controller (for example, those of the more active stations) and some to be kept in the host processor (those of the less active stations). Or ID verification can be done entirely by the host access method.
- *Multiple terminal access (MTA):* This option, available for certain low-speed, start-stop terminals, allows the NCP to communicate with dissimilar types of terminals over the same switched communication line. When a terminal calls the 3705 over the MTA line, the MTA option identifies the type of terminal and the transmission code used. This option accommodates the following types of terminals:

  IBM 1050 Data Communication System
  IBM 2740 Communications Terminal (Basic)
  IBM 2740 Communications Terminal (Transmit Control)
  IBM 2740 Communications Terminal (Transmit Control with Checking)
  IBM 2740 Communications Terminal (Checking)
  IBM 2741 Communications Terminal
  Terminals using CPT-TWX (models 33 and 35) code (at a line speed of 110, 134.5, and 300 bps)

The terminal types, code combinations, and communication lines to be used for multiple terminal access are specified as parameters in the program generation language.

- *Manual dial operation:* This option is for use when automatic calling is not available. Upon receiving a command to contact a station, the network control program, via the access method, sends the console operator a message instructing him to make the call. After the operator establishes the call, he places the communication line in data mode. The program can then communicate with the station.
- *Carriage return delay:* This option, available for certain start-stop stations, causes the network control program to pause momentarily before starting a write operation that immediately follows a read operation from the station. This prevents random printing during the return motion of the station's printing mechanism by allowing time for the printing mechanism to return to the left margin.
- *Monitor mode:* When this option is selected for a communication line, the network control program monitors the line during input and output operations. Between commands for an Attention signal sent by the station or a disconnect condition, the NCP notifies the access method.

Dynamic Control Functions

The network control program recognizes commands from the host access method to dynamically change certain parameters of the networks. Some of the dynamic control functions are standard; others are included in or excluded from the program by specifying them in the program generation macro instructions. The dynamic control functions include:

- Activating and deactivating communication lines. Commands from the access method request the program to activate or deactivate one or more communication lines attached to the 3705.
- Displaying any 256 contiguous bytes of 3705 storage. The requested bytes are sent to the host processor.
- Requesting the status of a communication line.
- Replacing ID characters and polling and addressing characters for BSC and start-stop stations.
- Changing the order in which stations on a multipoint BSC or start-stop communication line are polled and addressed.
- Changing the number of consecutive times stations on a multipoint BSC or start-stop communication line can respond negatively to polling before the line is rescheduled for other operations.
- Altering the sequence of network control program commands for a particular station.
- Changing the block-handling routines for data associated with a BSC or start-stop station.
- Setting the time and date in the 3705.
- Changing the maximum number of data transmissions between the host processor and a station on a multipoint BSC or start-stop line before the network control program tries to service other stations on the line.
- Turning off the power at a remote communications controller by command from the access method.

Advanced Communications Function

Described in this section is the Advanced Communications Function of the IBM Systems Network Architecture as it applies to the network control program/VS and to the Virtual Telecommunications Access Method (VTAM) and Telecommunications Access Method (TCAM).

What is Advanced Communications Function?

Advanced Communications Function is a group of IBM program products that use the concepts of Systems Network Architecture and the capabilities of products, such as the IBM 3705 Communications Controllers and SNA stations, to accommodate online applications in data communication networks of any size. The principal program products for Advanced Communications Function are: (1) Advanced Communications Function for Virtual Telecommunications Access Method (ACF/VTAM), (2) Advanced Communications Function for Telecommunications Access Method (ACF/TCAM), and (3) Advanced Communications Function for Network Control Program/Virtual Storage (ACF/NCP/VS). These program products are significant improvements to VTAM, TCAM, and NCP/VS.

ACF/VTAM, ACF/TCAM, and ACF/NCP/VS provide a flexible approach to the design and installation of single-processor and multiple-processor networks. These networks can take advantage of such SNA capabilities as sharing of resources, distributed function, device independence, and configuration flexibility.

Some of the key features of Advanced Communications Function are:

- Continuous growth potential from small to large single-processor networks, and from single-processor networks to a multiple-processor network
- A choice of access methods: ACF/VTAM and ACF/TCAM
- Compatibility with existing application programs that use the non-ACF versions of VTAM and TCAM
- Coexistence with other access methods

Extendable Data Communication Facilities

VTAM and TCAM (in their non-ACF versions) enable application programs to communicate with a wide variety of stations. Either of these access methods offers the following general services:

- Establish, terminate, and control access between application programs and stations
- Transfer data between application programs and stations
- Permit application programs to share communication lines, communications controllers, and stations, thus improving the flexibility of network configurations and possibly reducing the number of communications lines and stations required
- Permit the operation of the data communication network to be monitored and altered

The non-ACF versions of VTAM and TCAM permit only single-processor network configurations. That is, a single copy of the access method (TCAM or VTAM) can reside in the operating system for each host processor (System/370) in a network. This single-processor configuration is referred to as a *single-domain network*. ACF/VTAM and ACF/TCAM will support a multiple-processor network, referred to as a *multiple-domain network,* in addition to supporting single-domain networks. A multiple-domain network allows application programs and stations in one domain to communicate with application programs and stations in other domains by means of cross-domain connections. Each domain in a multiple-domain network must contain either ACF/VTAM or ACF/TCAM as the access method. A multiple-domain network permits cross-domain resource sharing among two or more interconnected domains. This multiple-domain capability is referred to in this publication as *networking.*

Both ACF/VTAM and ACF/TCAM use the facilities of ACF/NCP/VS with the IBM 3705 Communications Controllers to control remote stations. ACF/VTAM uses the ACF/NCP/VS to control stations on communication lines in network control mode only. ACF/TCAM can use the ACF/NCP/VS with partitioned emulation programming (PEP) to control stations on lines in either network control mode or emulation mode with local communications controllers. Appendix C summarizes the emulation mode support provided by ACF/NCP/VS, which does not differ from that provided by the non-ACF version of NCP/VS.

The IBM facilities for Advanced Communications Function accommodate both single-domain and multiple-domain networks. Five separate facilities provide discrete functional capabilities:

ACF/VTAM
ACF/VTAM Multisystem Networking Facility (an optional feature)
ACF/TCAM
ACF/TCAM Multisystem Networking Facility (an optional feature)
ACF/NCP/VS

Because this publication describes the network configurations that are possible with ACF/VTAM, ACF/TCAM, and ACF/NCP/VS, the host processor for the single-domain networks is assumed to contain either ACF/VTAM or ACF/TCAM. The communications controllers are assumed to contain ACF/NCP/VS. For the multiple-domain network, each host processor must also contain the appropriate ACF/VTAM or ACF/TCAM Multisystem Networking Facility. The program requirements in ACF/NCP/VS are independent of which access method (ACF/VTAM or ACF/TCAM) is used in the host processor, or whether ACF/NCP/VS is to be used in a single-domain or multiple-domain network.

Communications Controller Requirements

ACF/NCP/VS can be executed in either the IBM 3705-I or 3705-II Communications Controllers. However, certain network configurations and capabilities described in this publication are dependent on hardware options available only with the 3705-II Communications Controller. The terms *3705 communications controller* and *communications controller* are used to mean either the 3705-I or 3705-II when no distinction is required. When identifying functions available only with the 3705-II Communications Controller, this publication uses a specific model designation.

Stations Supported

Advanced Communications Function supports two categories of stations: *SNA stations* and *non-SNA stations*. SNA stations can perform functions defined by Systems Network Architecture (SNA) and, when attached to a 3705, use the synchronous data link control (SDLC) discipline. Non-SNA stations include start-stop and BSC stations (Some SNA and non-SNA stations can be locally attached to a host processor channel.) Appendix A lists the types of stations and systems that ACF/NCP/VS can control in a SNA-based network.

Programming Subsystems

Advanced Communications Function supports various subsystems, such as CICS/VS (Customer Information Control System/Virtual Storage) and IMS/VS

(Information Management System/Virtual Storage). The two access methods vary in the support of these and other subsystems. For information on the subsystems ACF/TCAM supports, see *ACF/TCAM General Information: Introduction,* GC30-3057. For information on the subsystems ACF/VTAM supports, see *ACF/VTAM General Information: Introduction,* GC27-0462.

Chapter 2. ACF/NCP/VS in a Single-Domain Network

This chapter introduces some general concepts of SNA as represented in the design of ACF/NCP/VS and the ACF access methods, ACF/TCAM and ACF/VTAM, along with some SNA terms frequently used in discussing SNA networks.

The term *single domain* applies to a network that has a single controlling point, namely, the access method with which the ACF/NCP/VS communicates. SNA networks that result from combining previous separate networks will frequently have more than one controlling point (access method), each responsible for a portion of the network called a domain. Or a single network may be designed from the beginning, or later modified, to have multiple controlling points as a means of realizing the distribution of function that is a primary objective of Systems Network Architecture. In either case, if elements of one domain (such as stations and application programs) can communicate with elements of another domain (cross-domain communication), the entire network is called a *multiple-domain network.* (If a network has more than one domain but there is no cross-domain communication, the result in effect is multiple single-domain networks.)

Systems Network Architecture

Systems Network Architecture fulfills the concepts of distributed function among the physical components of a data communication system and the sharing of resources among the users of the system. SNA is not itself a system; it is a design solution for a total data communication network. SNA is an architecture that, when used in the design of devices and programs, allows application programs to be independent of device attachments. The networks described in this publication use the SNA concepts and represent an implementation of these concepts. For a detailed understanding of SNA concepts, see *Systems Network Architecture General Information,* GA27-3102.

Network Addressable Units

The origins and destinations of information units flowing in an SNA network are called *network addressable units* (NAU). Any transfer of information within the network involves an origin NAU and a destination NAU that are formally bound to each other during the information transfer. The duration of this association between two NAUs is called a *session.*

SNA defines three types of NAUs:

- *System services control point (SSCP):* The SSCP is responsible for the general management of the network, such as bringing up the network, establishing sessions, or recovering when a network component fails to maintain contact.
- *Physical unit (PU):* Each resource in the network (host processor, communications controller, cluster controller, SNA terminal) that has been defined to the SSCP is associated with a physical unit. (Communications controllers provide physical unit services for certain nonprogrammable stations.)
- *Logical unit (LU):* A logical unit is the port through which the end user gains access to the services of the network. One or more logical units can be associated with a physical unit.

Network Resources

The essential characteristic of an SNA domain is the existence of a control function in the access method called the *system services control point* (SSCP), identified earlier as one of the network addressable units. Since the SSCP is responsible for the general management of the network, a *domain* is the set of resources in the network that is known to and managed by the SSCP.

Each resource in the domain is identified to the SSCP by a unique symbolic name, assigned by the systems programmer to that resource when the access method and the network control program are defined. All references by end users (operators and application programs) to resources are by resource name. Thus the end user need not be concerned about the location of the resource in the network, nor with the path the SSCP and the network control program establish to reach that resource.

The access method uses internal tables to translate each reference to a resource name into a *network address.* This address has two parts: a subarea address and an element address (analogous to the exchange digits and the remaining digits that make up a local telephone number). Each network control program in a network is assigned a unique subarea address by the system programmer. Each resource directly controlled by that network control program is assigned an element address. A subarea thus constitutes the network control program and all of the resources (elements) that that network control program directly controls. Similarly, subarea addresses are assigned to the host processor and to sets of stations that are locally attached to the host processor and are directly controlled by the access method. (Host logical units assume the subarea address of their host processor.)

Network Layers

Information from one end user of the network to another end user flows through a series of network layers. Each layer has a specific and sharply delineated purpose in the network. SNA defines requests and responses to be used by each layer, as well as the protocols (sequences of requests and responses) by which layers communicate. A specific layer in one resource is mainly concerned with communicating with the equivalent layer in another resource. A layer is not concerned with what the intervening layers have to do to forward its communication through the network, and it is not aware of the protocols that those layers use.

SNA designates three functional layers: application layer, function management layer, and transmission subsystem layer. (See Figure 2-1).

Application Layer

The application layer represents the end user of the network; it includes user-written application programs, as well as IBM programs that take the place of application programs.

Function Management Layer

The *function management layer* is concerned with the presentation of data from one application layer to another. Three services of the function management layer provide this function:

- *Data flow control* helps the end user by controlling the flow of requests and responses in the network; this allows the end user to enforce the data flow protocols he desires.
- *Presentation services* provide support for communication between end users engaged in LU-LU sessions.
- *Logical unit services* support communication between a logical unit and the SSCP (SSCP-LU sessions) and aid in establishing and terminating LU-LU sessions.

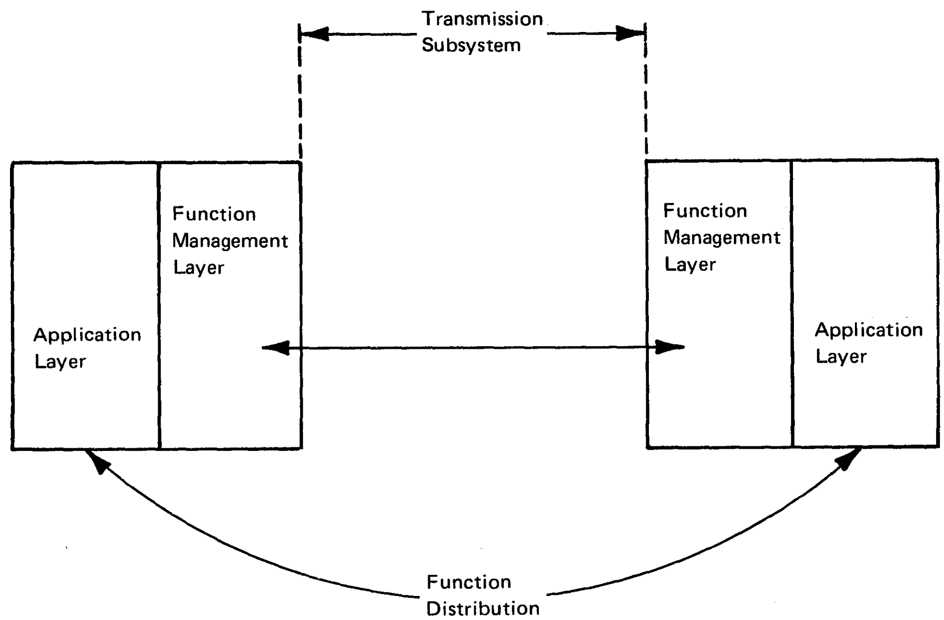

Figure 2-1. The Transmission Subsystem as the Innermost Layer of a Network

Transmission Subsystem Layer

Figure 2-1 shows that the *transmission subsystem* is the innermost layer of an SNA network. This layer is responsible for data from the time it leaves one function management layer until it reaches another. The transmission subsystem is concerned with routing and moving data between network addressable units (defined previously as the SSCP, physical units, and logical units). The network control program is part of a network's transmission subsystem.

The transmission subsystem consists of three sublayers: transmission control, path control, and data link control. (See Figure 2-2.)

- *Transmission control* consists of the following:
  - The *connection point manager* assigns and checks sequence numbers for each message transmitted or received, coordinates responses with requests and keeps them in proper order, paces data traffic, and creates or reacts to exception status and sense information.
  - *Session control* establishes and terminates sessions and obtains the resources required for a session. (Sessions are discussed later in this chapter under "SNA Sessions.")
  - *Network control* provides a means for adjacent connection point managers to communicate.
- *Path control* manages the routing of path information units to and from the access method.

- *Data link control* manages an individual data link. *(Data link* is a generic term for any physical facility used to carry data from one network component to another; examples are communication lines and communication loops.) Data link control in each node manages the data links attached to that node. Data link control in the access method manages the channel connection(s) between the host processor and the communications controller.

Request/Response Units and Headers

The combination of the original message (as entered by the end user) and the control information added by the function management layer is called a *request/response unit (RU).* Request/response units are the basic units of information entering and exiting the transmission subsystem layer. A request/response unit may be:

- Data
- Acknowledgment of data
- Commands that control the flow of data
- Responses to commands

Throughout this publication, a request/response unit is referred to as RU, or specifically as a request or a response.

Four types of request/response units (RUs) can flow through the network:

- Function management data RUs
- Data flow control RUs
- Network control RUs
- Session control RUs

The type of RU is identified in the request/response header, which is added by the connection point manager. Each type of RU is discussed below.

- A *function management data RU* contains message data sent to or from a network addressable unit. Each function management data RU corresponds to an access method buffer. If a message is larger than a buffer, the access method creates a chain of these RUs; control bits in the request/response header indicate whether an RU is the first, middle, last or only RU in a chain. A chain is the basic unit of information that can be recovered following an error; if one RU of a chain cannot be processed, the entire chain must be discarded or retransmitted.
- *Data flow control RUs* control the flow of requests and responses between logical units that are in session. These RUs are used to enforce protocols agreed upon when a session is established.
- *Network control RUs* are used by the access method to communicate between adjacent connection point managers without formally establishing a session.
- *Session control RUs* are used to establish and terminate sessions. Session control RUs are discussed under the topic "LU-LU Sessions" later in this chapter.

Figure 2-2 shows the step-by-step progress of message data through the network layers and the various headers that the layers affix to the data.

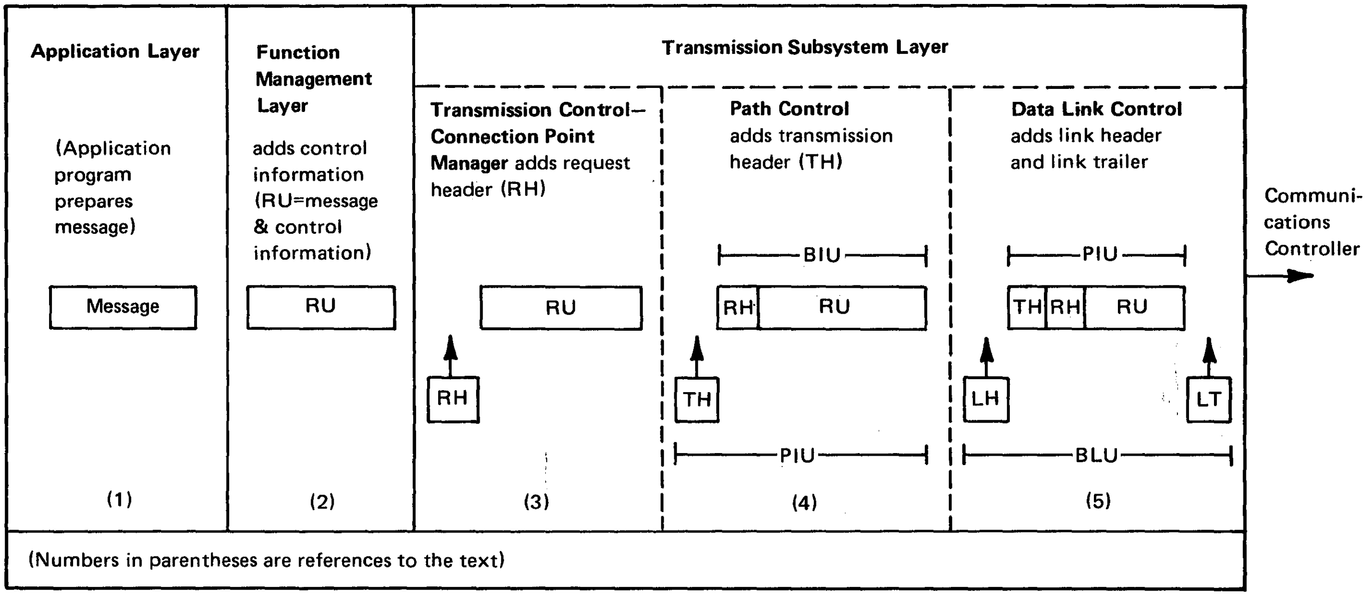

Figure 2-2. Data Flow through Network Layers

The steps are as follows:

1. An application program has a message ready to send to a remote station and passes it to the function management layer.
2. The function management layer adds control information; the message then becomes a *request/response unit* (RU). The function management layer then routes the request/response unit to the connection point manager.
3. The connection point manager uses the control information to build a request/response header (RH) for the message. The combination of a request/response header and a request/response unit is called a *basic information unit* (BIU). The connection point manager forwards the basic information unit to path control.
4. Depending upon the size of the basic information unit and the buffer size at the destination, path control may break the BIU into segments. (This is not shown in the figure.) In any case, path control places control information into a *transmission header* (TH); the transmission header includes information on addressing, mapping (segment indicators) and sequencing. Path control attaches a transmission header to each basic information unit (or if segmented, to each of its segments). After the transmission header is added, the basic information unit becomes a *path information unit* (PIU). Path control now has the option of blocking one or more PIUs into a *basic transmission unit.* Blocking is useful when several PIUs are ready at the same time to be sent to the same communications controller. PIUs within the basic transmission unit can be destined for the same or different resources reached through that controller.
5. Data link control adds a *link header* (LH) and a *link trailer* (LT) to the basic transmission unit and sends it to the communications controller. After the link header and trailer are added, the basic transmission unit becomes a *basic link unit* (BLU).

So, by the time that the message in step 1 is transmitted to the communications controller, it is called a basic link unit, whose format is:

|———————— Basic Link Unit ————————|

Link Header	Transmission Header	Request Header	Request Unit	Link Trailer

|—— Original Message ——| (Request Unit)

If the message is destined for a station (without a logical unit) connected to a communications controller, the reverse of this procedure takes place in the communications controller. So, the end user sees only what was sent in step 1. If the original message (from step 1) is destined for a terminal attached to a cluster controller or for a logical unit, the cluster controller or the logical unit removes the header and trailer.

SNA Sessions

Before information can flow between network addressable units (NAUs), the access method must determine which pair of NAUs is to communicate and must establish a session between them. Sessions can exist between:

- The SSCP and an outboard physical unit (SSCP-PU session)

  The SSCP-PU session is used to communicate information about the status of the physical network, commands or requests to change the status of the network, hardware test requests and results, etc.

- The SSCP and a secondary logical unit (SSCP-LU session)

  The SSCP-LU session is used to communicate (1) commands from the SSCP to activate or deactivate the secondary logical unit and (2) requests from the secondary logical unit to begin or end an LU-LU session. Operator commands can be communicated to the access method during this type of session.

- Two logical units (LU-LU session)

  The LU-LU session allows end users to exchange data and to control the flow of this data.

An access method parameter specifies the maximum number of concurrent sessions to be permitted at any given moment. For single-domain networks, this parameter is the maximum number of concurrent SSCP-PU, SSCP-LU, and LU-LU sessions.

SSCP-PU Sessions

Once the access method is active, the SSCP initiates a session with each physical unit in the network. This kind of session is called an SSCP-PU session. (The access method has the option of designating some physical units as initially inactive; these physical units can be activated later by operator commands.)

The SSCP-PU sessions remain established until the access method deactivates the physical unit. Deactivation occurs because of (1) an operator command to deactivate the physical unit, (2) closedown of the access method, or (3) occurrence of an error condition.

The first SSCP-PU sessions that the SSCP establishes are those with the network control programs in the domain. Once such a session is established, the network control program is active; the SSCP can then establish SSCP-PU sessions with physical units with which the SSCP communicates via that active network control program.

SSCP-LU Sessions

After an SSCP-PU session is established, the SSCP can initiate a session with the logical unit(s) associated with the active physical unit. (If the logical unit requested is not active, it can be activated by an operator command.) This kind of session is called an SSCP-LU session. An SSCP-LU session remains established until the logical unit or its associated physical unit is deactivated. As is the case for physical units, deactivation of a logical unit occurs because of (1) an operator command to do so, (2) closedown of the access method, or (3) occurrence of an error condition.

LU-LU Sessions

Once an SSCP-LU session has been established, the logical unit is eligible to participate in an LU-LU session with a host LU. (A host LU is either an ACF/VTAM application program or the part of ACF/TCAM that initiates and maintains a session on behalf of ACF/TCAM.)

The host LU can issue session control RUs and thereby establish and terminate the LU-LU session. Once that session is established, data may flow in either direction between the logical units until one of the logical units is deactivated.

Any active logical unit (that is, a logical unit that is in session with the SSCP) can initiate (request) a LU-LU session. However, only primary LUs can actually establish a session. (Session protocols require that one end of a session be a primary LU and the other end a secondary LU.) In sessions between a host LU and a station (logical unit), the host LU is always the primary LU.

Network Services

The SSCP uses a set of command processors to provide physical units and logical units with *network services*. These services include:

- *Configuration services,* which allow the use of operator commands to change the status of the network (for example, activating or deactivating a physical unit or logical unit)
- *Maintenance services,* which allow an end user to request (1) online tests of resources in the network, (2) network control program traces, (3) status of network resources, and (4) statistics on network resources.
- *Session services* (discussed earlier in this chapter under "SNA Sessions").

Single-Domain Network Configurations

A single-domain network configuration is shown in Figure 2-3. This configuration includes a single ACF/NCP/VS and a single access method and therefore constitutes a single-domain network.

Figure 2-4 shows a configuration that includes three access methods and a single ACF/NCP/VS. The resources (physical units, logical units, and non-SNA stations) served by the ACF/NCP/VS are divided among the access methods such that each access method communicates with only the resources in its domain. No resource in one domain can communicate with a resource in another domain. Because there is, in this case, no cross-domain communication, this configuration is in effect divided into three single-domain networks. This is an example of multiple single-domain networks sharing the same ACF/NCP/VS. If, in contrast, cross-domain communication did exist, the configuration would constitute a multiple-domain network; such networks are discussed in Chapter 3.

ACF/NCP/VS supports the multiple-channel attachment facility of the 3705 communications controller, as shown in Figure 2-4. This facility permits the concurrent sharing of a single 3705 and its network control program by more than one access method in the same or separate host processors. The number of concurrent channel attachments possible depends on the model of 3705 installed; see *Introduction to the IBM 3704 and 3705 Communications Controllers* (GA27-3051) for the attachment possibilities. While this multiple-channel capability adds flexibility to the possible configurations for single-domain networks, coordination is required between the network operators responsible for each domain.

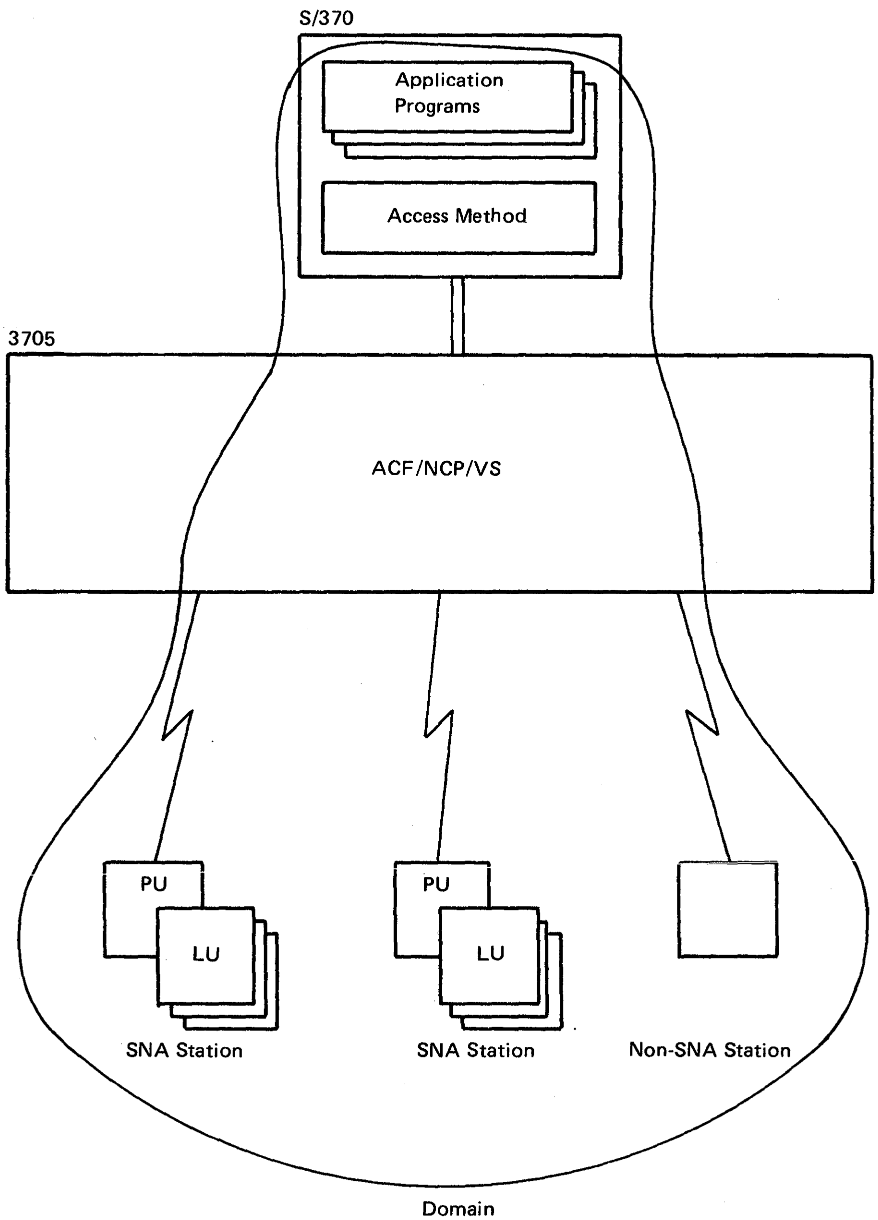

Figure 2-3. ACF/NCP/VS in a Single-Domain Network

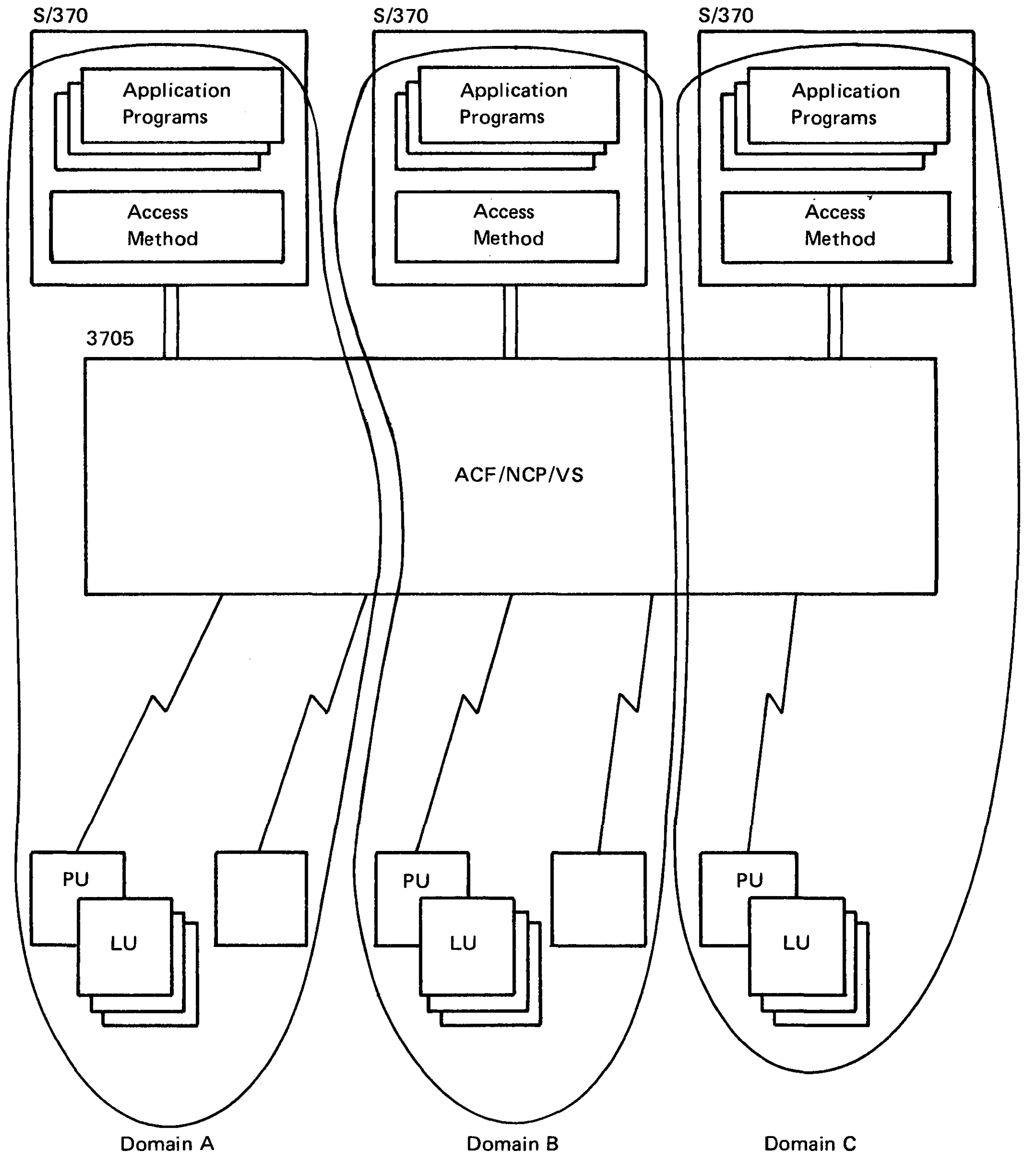

Figure 2-4. ACF/NCP/VS in Multiple Single-Domain Networks

Chapter 3. ACF/NCP/VS in a Multiple-Domain Network

In Chapter 2, the facilities of SNA single-domain data communication networks using ACF/VTAM, ACF/TCAM, and ACF/NCP/VS were described. The single-domain facilities accommodate multiple concurrent application programs and both SNA and non-SNA stations with extensive sharing of the network resources, such as stations, communications controllers, and communication lines, within the domain. While two or more of these single-domain networks may coexist, with their SSCPs even occupying the same local host processor and sharing the same communications controller, the networks remain independent of each other. The domain that constitutes each network cannot communicate with another domain, and resources cannot be shared among domains.

This chapter describes the capabilities that are available through ACF/NCP/VS and the Multisystem Networking Facilities of ACF/VTAM and ACF/TCAM, which allow the interconnection of two or more networks into one consolidated and coordinated network called a multiple-domain network. This consolidation of networks (domains) is called *networking.*

Some of the advantages of networking are:

- Increased access to information and host processing
- Extended sharing of resources across domains
- Elimination of redundant application programs in two or more domains
- Transparency, to the application programs, of the locations of stations involved in cross-domain communication
- Continuation of cross-domain operations following some failures in the network
- Possible decrease in communication line costs by utilizing cross-domain resources
- Backup capability for critical applications or devices

Networking Concepts

The key concept of networking is that the composite network is made up of two or more separate networks (domains) that have been interconnected. Domains are interconnected by cross-domain links or by use of the sharing capability of the ACF/NCP/VS. Cross-domain connections allow the network control program to be concurrently owned, or shared, by up to eight system services control points (SSCP). The SSCPs may communicate with the network control program over one or more host processor subchannels...

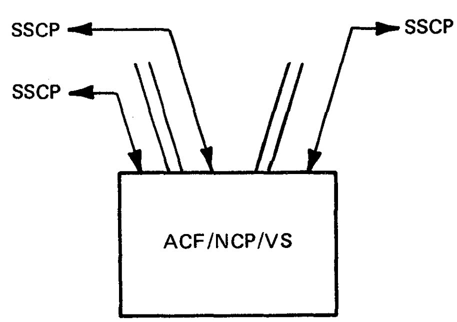

or over one or more SDLC links to other network control programs...

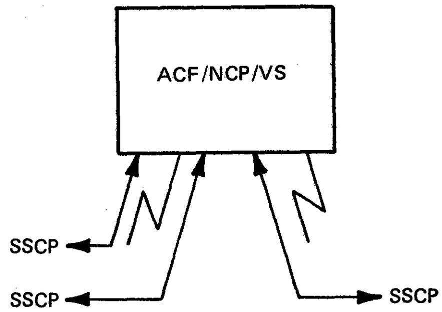

or over a combination of subchannels and SDLC links...

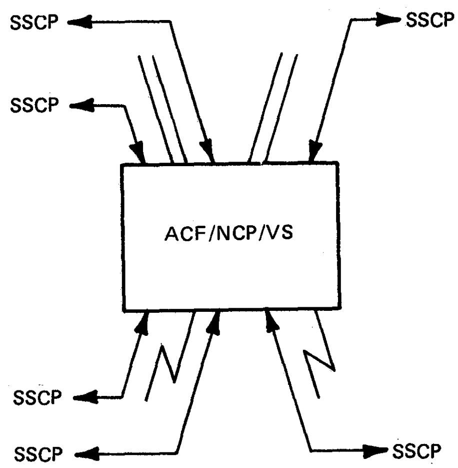

The specific configuration used will depend on the geographical locations of the host processor and communications controllers and upon such operational considerations as the need for backup data links and network control programs.

An ACF/NCP/VS may be a part of a single domain (that is, owned by a single SSCP), as shown in Figure 3-1; or it may be a part of as many as eight domains (that is, concurrently owned by as many as eight SSCPs). Figure 3-2 shows a single ACF/NCP/VS as part of three domains.

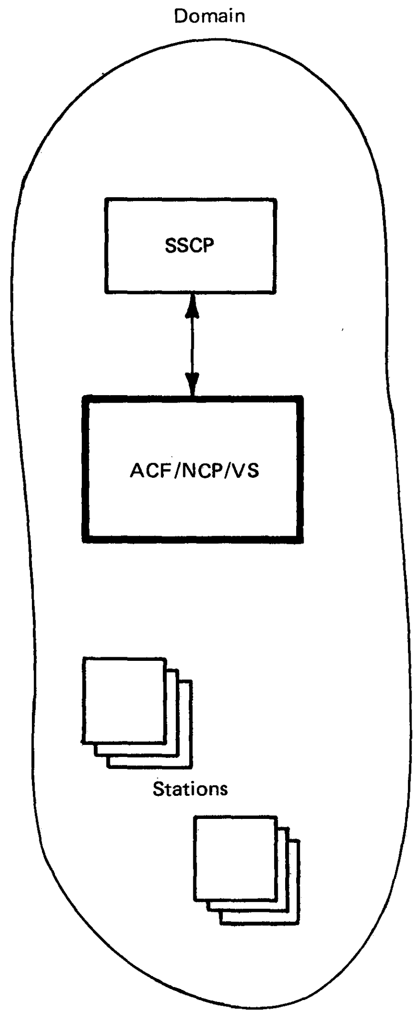

Note: SSCP may be in either ACF/TCAM or ACF/VTAM.

Figure 3-1. ACF/NCP/VS in One Domain

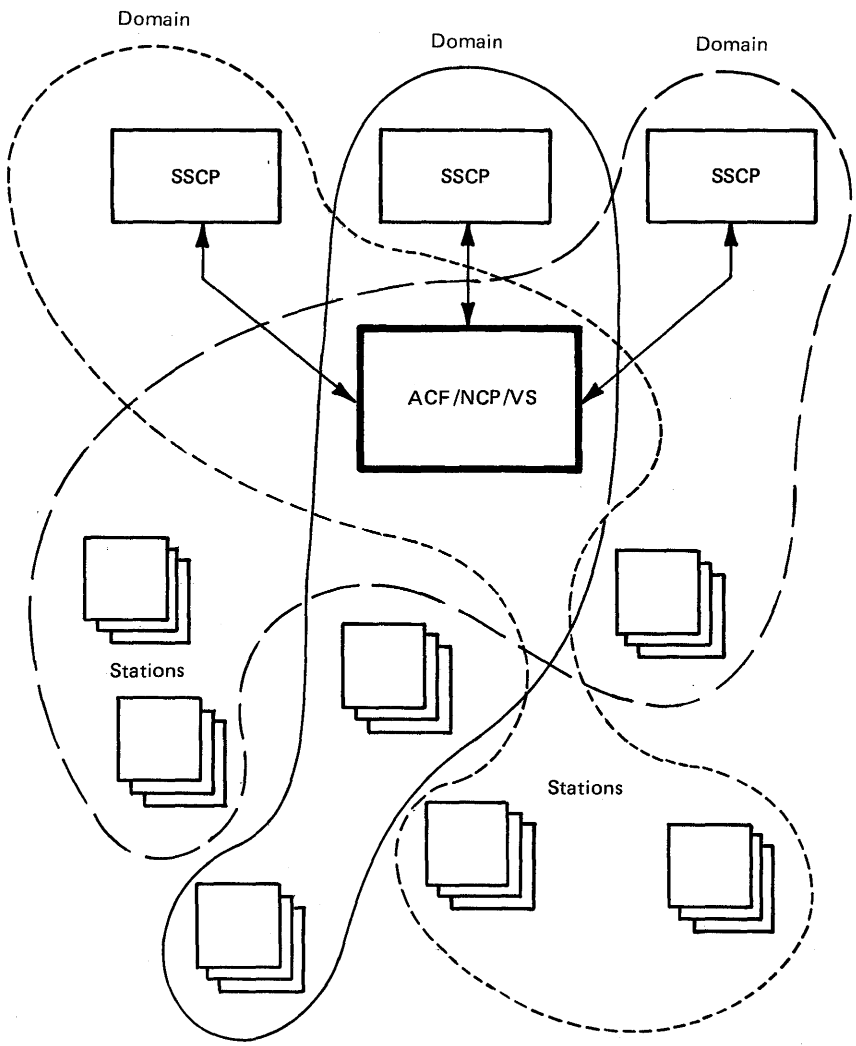

Note: SSCP may be in either ACF/TCAM or ACF/VTAM.

Figure 3-2. ACF/NCP/VS in Three Domains

The concept of shared ownership extends to the data links and stations in the network. The SNA stations and logical units (LU) attached to a multipoint SDLC link need not all be part of the same domain. Some can be in one domain, others in different domains. This capability makes possible extensive sharing of resources without regard to their physical location in the network. With this kind of sharing, ownership of the link is said to be concurrently shared by the SSCPs that control the domains. Stations and logical units on the link can be *serially* shared; that is, ownership of these resources can be transferred from one SSCP to another, but only one SSCP can own such a resource at any point in time. Transfer of ownership requires coordination between the network operators responsible for the domains involved.

Ownership of non-SNA stations controlled by an ACF/NCP/VS can also be divided among multiple domains, but ownership of all stations on any single multipoint start-stop or BSC line must be vested in the same SSCP. That is, ownership of the line cannot be concurrently shared. However, ownership of the line and its attached stations can be transferred as a unit from one SSCP to another. As in the case of SNA stations, such transfer requires coordination by the network operators responsible for the domains involved.

Shared ownership of data links, stations, and logical units does not depend on the paths by which the various owning SSCPs communicate with the network control program. Each SSCP may reach the network control program over a channel connection or a combination of channel connections and data links. However, an SSCP cannot own an ACF/NCP/VS (or resources attached to the ACF/NCP/VS) from which it is separated by more than one SDLC link. (An exception to this rule occurs when a local communications controller is reconfigured as a remote controller, as explained later in this chapter.)

A SNA-based network is not limited to eight domains. This limit refers to the maximum number of SSCPs that can concurrently share any one network control program.

The number of SSCPs that can own a particular ACF/NCP/VS does not directly depend on the type or number of physical connections between the communications controller in which the ACF/NCP/VS resides and other host processors and communications controllers in the network. For example, the controller may have a single channel connection to a host processor, as shown in Figure 3-3[a], and have one owning SSCP (Figure 3-3[b]). Or the controller may have multiple channel connections to host processors (Figure 3-3[c]) and have one or more owning SSCPs (Figure 3-3[d] and [e].

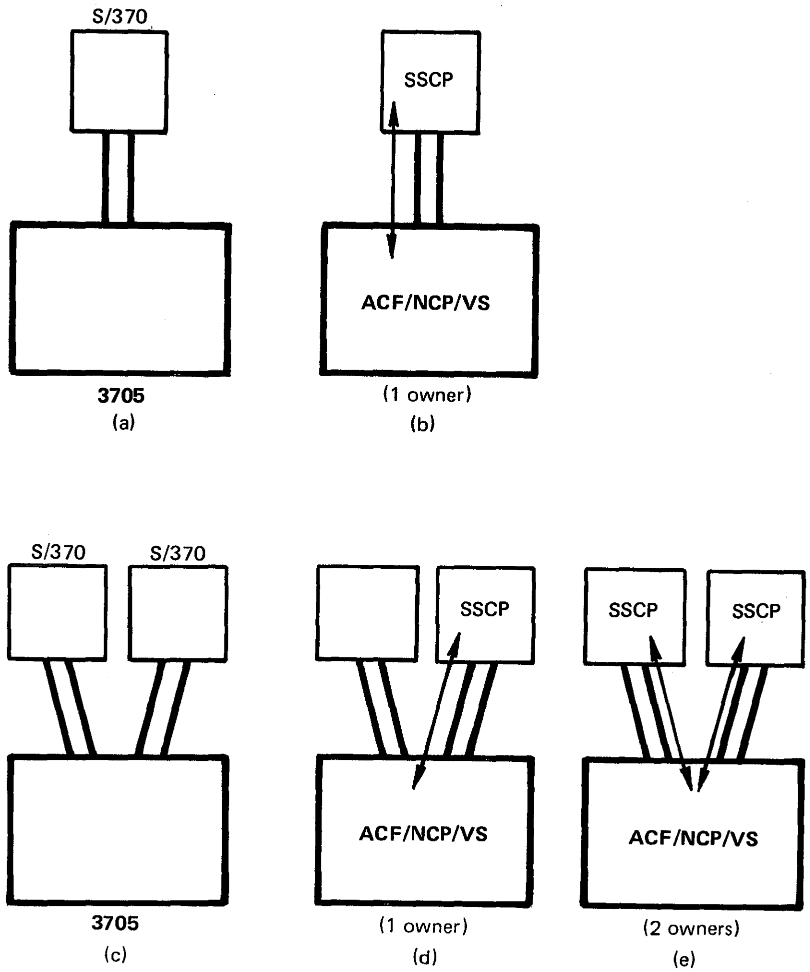

Note: SSCP may be in either ACF/TCAM or ACF/VTAM.

Figure 3-3. Examples of Shared Ownership of ACF/NCP/VS Via Channel Connections

Figure 3-4 shows examples of connections to one or more SSCPs over one or more SDLC links, which may be local-local or local-remote links. As in the case of channel connections, the number of SSCPs that can share ownership of an ACF/NCP/VS does not directly depend on the number of local-local or local-remote links attached to the controller.

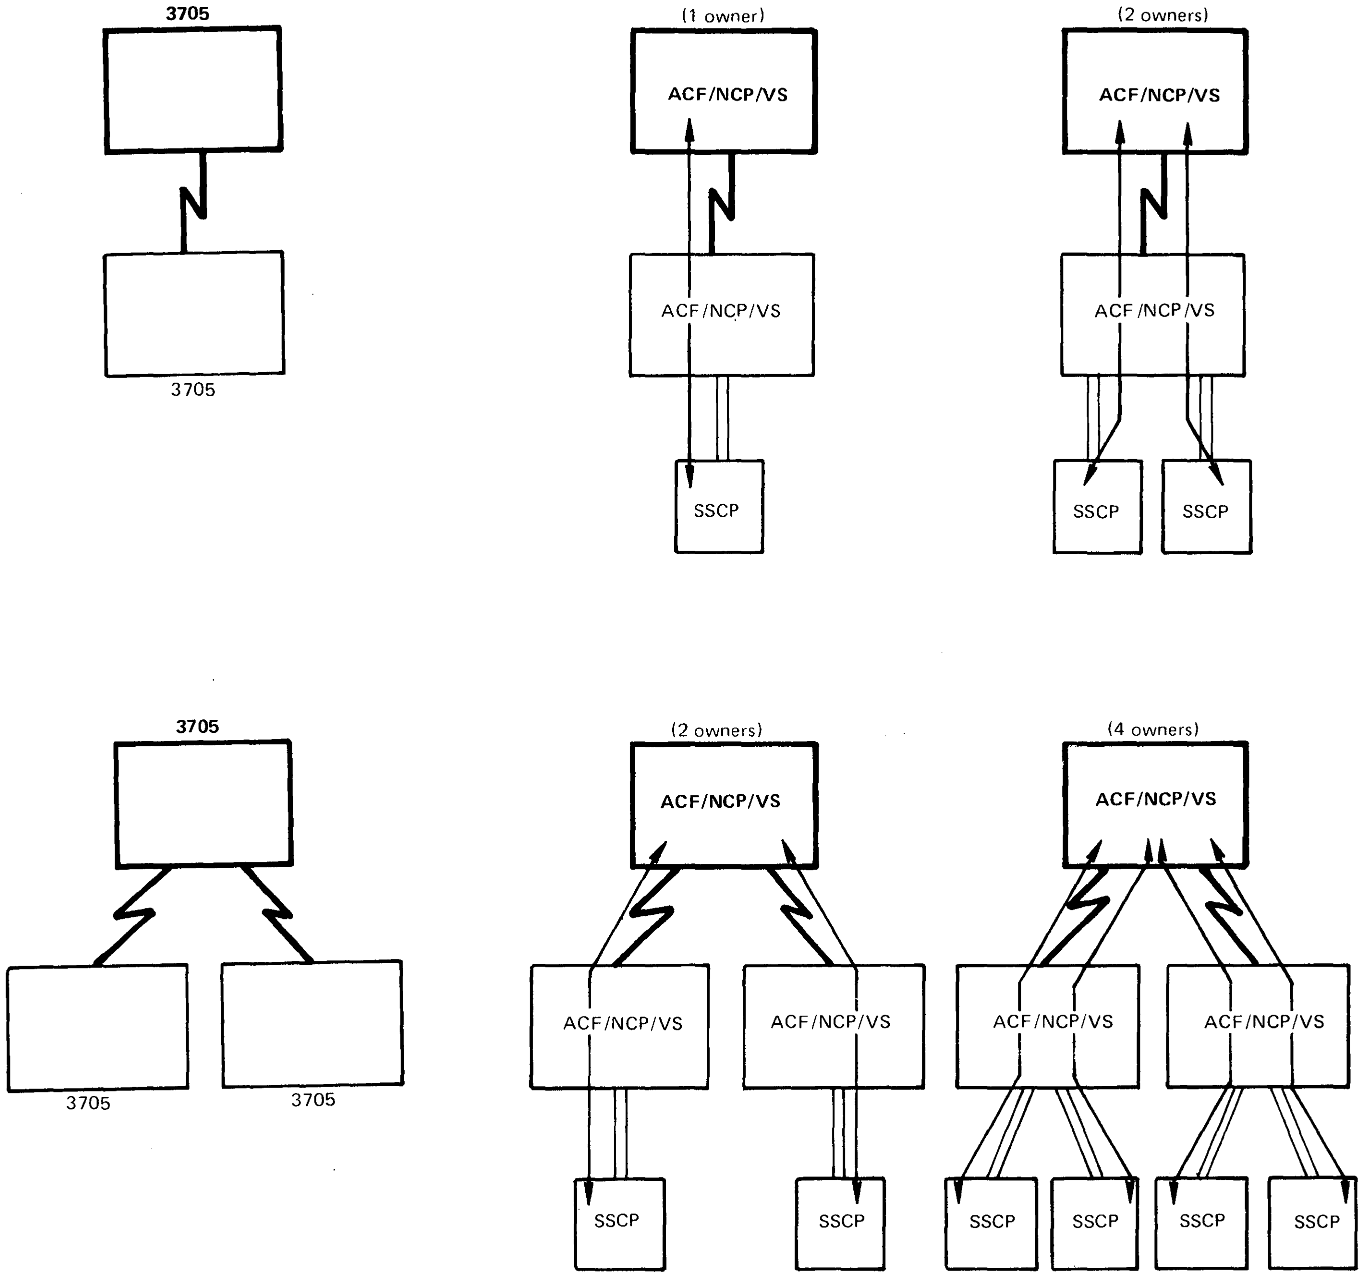

Note: SSCP may be in either ACF/TCAM or ACF/VTAM.

Figure 3-4. Examples of Shared Ownership of ACF/NCP/VS Via SDLC Links

Figure 3-5 shows some examples of the use of both channel connections and SDLC links to accomplish shared ownership of an ACF/NCP/VS by multiple SSCPs.

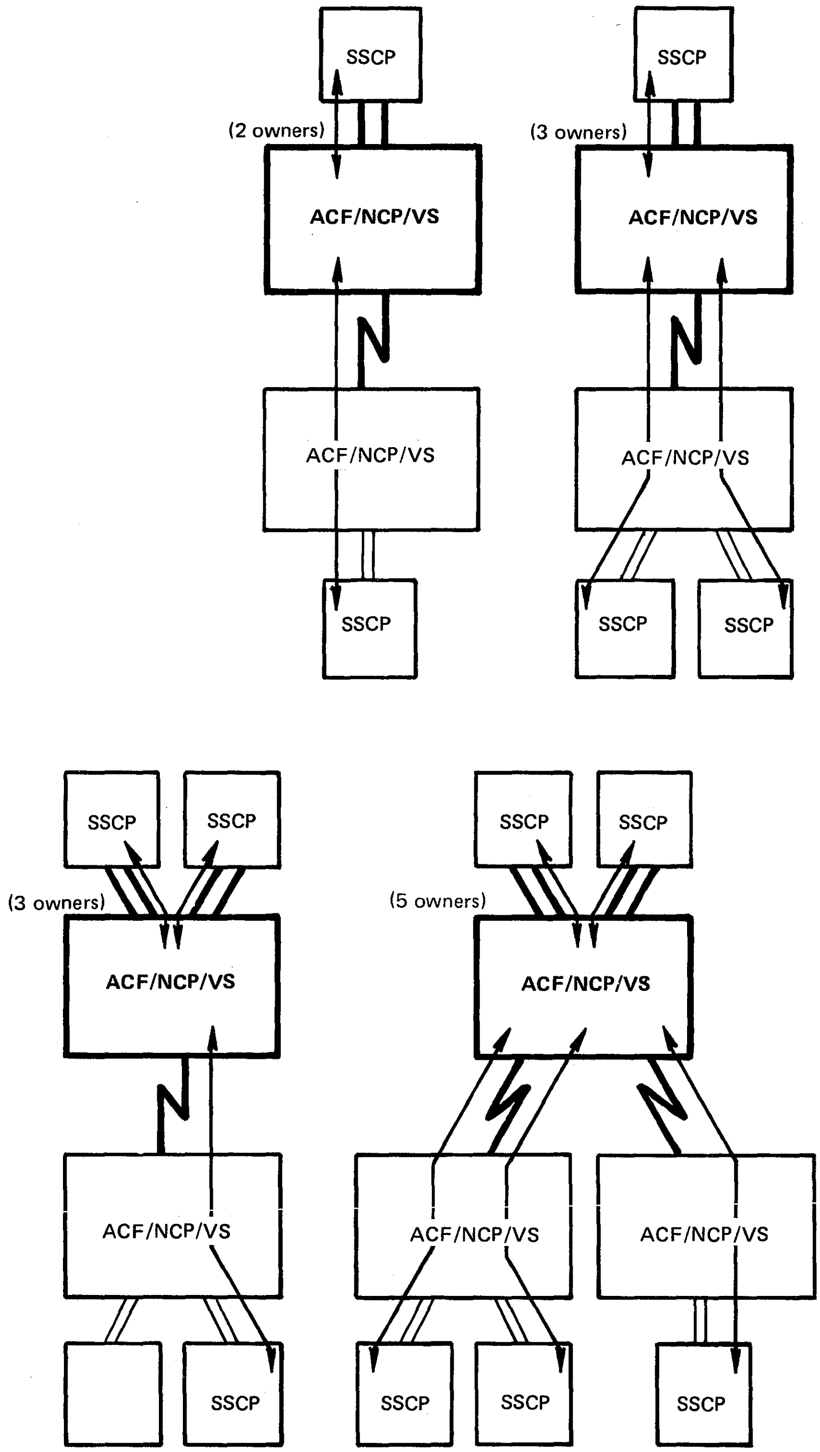

Note: SSCP may be in either ACF/TCAM or ACF/VTAM.

Figure 3-5. Examples of Shared Ownership of ACF/NCP/VS Via Both Channel Connections and SDLC Links

Although the combined channel and SDLC link connections to a communications controller do not directly govern how many SSCPs can share ownership of the ACF/NCP/VS, as Figures 3-3 through 3-5 show, there are, nonetheless, other aspects of the network configuration that can limit the number of possible owners. For example, as stated earlier, an SSCP cannot own an ACF/NCP/VS from which it is separated by more than one SDLC link except under certain backup conditions described later.

A Definition of Networking

Networking is the ability of an ACF/VTAM or ACT/TCAM application program in one domain to communicate with either another application program or a station in some other domain without involving host processing in intervening domains. This concept, although not restricted for use with SNA stations, is easiest understood in terms of the SNA concept of LU-LU sessions. Thus, unless a distinction is required between SNA stations and non-SNA stations, the networking concepts are explained in terms of SNA logical units.

Structural differences in ACF/VTAM and ACF/TCAM necessitate an understanding of the terminology used when referring to the ACF/VTAM and ACF/TCAM access methods. In most instances ACF/VTAM and ACF/TCAM can be used interchangeably. A distinction is made where required. If no distinction is needed, ACF/VTAM and ACF/TCAM are referred to as the *access method.* ACF/VTAM application programs provide some of the logical unit functions and all of the end user functions of an SNA-based network, while ACF/TCAM application programs are actually the end users and ACF/TCAM itself contains the logical unit associated with the ACF/TCAM application program. For simplicity the term *host LU* is used to mean either the ACF/VTAM application program or the part of ACF/TCAM that initiates and maintains a session on behalf of ACF/TCAM.

In networking, one end of an LU-LU session is always a host LU. Once an LU-LU session is established, data my flow in either direction using the networking facilities of the access method and ACF/NCP/VS. LU-LU session protocols require that one end of the session be a primary LU and the other end be a secondary LU. In sessions between a host LU and a station (or logical unit), the host LU is always the primary LU.

Cross-Domain Data Links

Cross-domain data links connect only local communications controllers (attached to the host processor by a data channel) as shown in Figure 3-6. For releases of ACF/NCP/VS prior to Release 3, there can be no cross-domain data links between remote communications controllers or between a local and a remote communications controller. (Release 3 of ACF/NCP/VS removes this restriction; see Chapter 6 for details.) The data links are normally nonswitched but can be manually dialed (switched). All cross-domain links use the SDLC discipline. Both domains joined by the link participate equally in the activation of the link.

Multiple cross-domain links are possible if multiple communications controllers are available. This capability is shown in Figure 3-6. This example shows three cross-domain links between domain 1 and domain 2. A possibility for a fourth link exists between L12 and L22 but is not shown. (In this and succeeding figures, H represents a host processor, L represents a local communications controller, and R represents a remote communications controller. LU represents a logical unit, C represents a cluster controller, and S represents a station.)

For releases of ACF/NCP/VS prior to Release 3, multiple links between domains do not provide alternate routing or parallel paths. However, multiple links do allow selection of path routing based on anticipated network traffic or other pertinent business considerations. A selected path must be the same in both directions (see the section on "Network Addressing and Routing"). See Chapter 6 for information on the parallel link and multiple routing capabilities of ACF/NCP/VS Release 3.

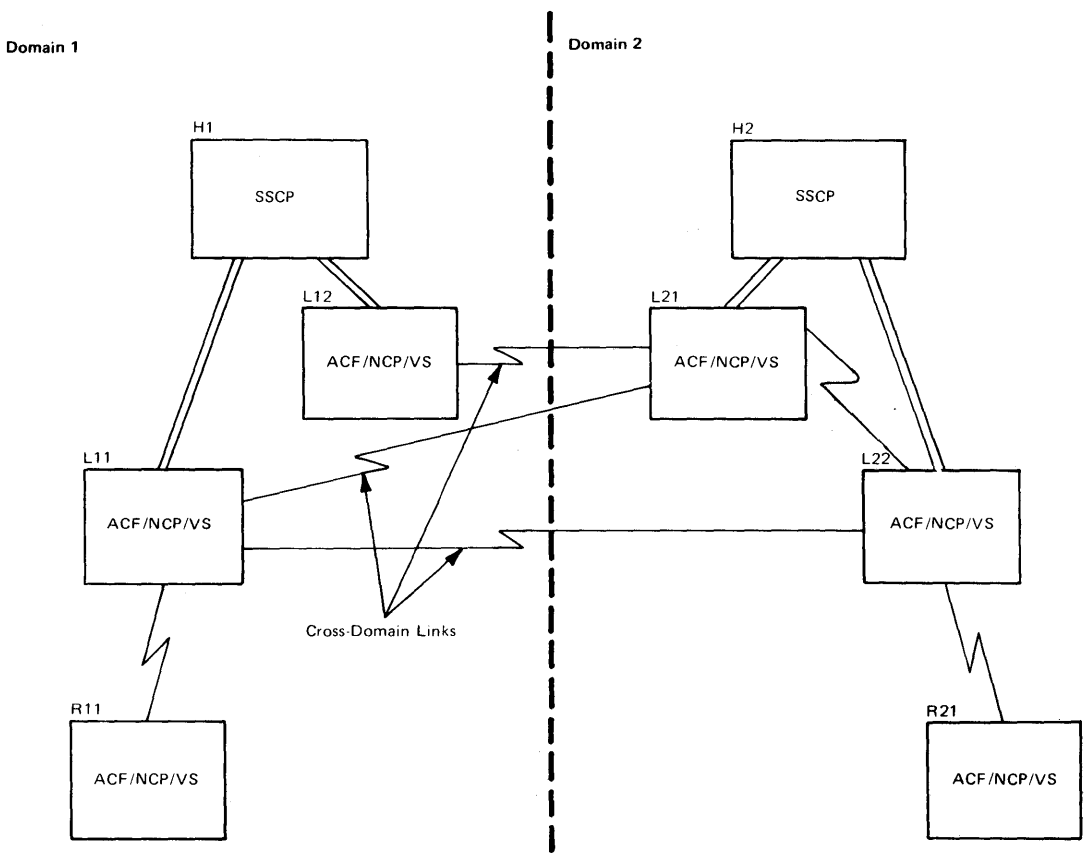

Note: SSCP may be in either ACF/TCAM or ACF/VTAM.

Figure 3-6. Multiple Cross-Domain Links

Cross-Domain Resources

The resources that can be shared directly across domains in the networking environment depend on the access methods involved in both ends of a cross-domain session.

If one or both ends of a session are in an ACF/VTAM-controlled domain or an ACF/TCAM-controlled domain, host LUs, remote SNA stations (LUs), remote BSC 3270 stations, and local (channel-attached) SNA 3270 and 3790 stations can be shared across domains. Local (channel-attached) SNA and non-SNA 3270 stations also can be shared across domains, as can (for ACF/VTAM) remote non-SNA stations supported by the Network Terminal Option. (Of the local 3270 stations, ACF/TCAM supports only the 3274.)

If both ends of the session are in ACF/TCAM-controlled domains, remote BSC and start-stop stations also can be shared across domains.

For either ACF/TCAM or ACF/VTAM, the ability to communicate with cross-domain resources is not influenced by access methods in intervening domains.

Initiating Cross-Domain SNA Sessions

As described in Chapter 2, LU-LU sessions can be initiated only after the SSCP has established SSCP-PU and SSCP-LU sessions. In a multiple-domain network, each domain contains an SSCP. Sessions are first established in each domain between its SSCP and the respective PUs and LUs in that domain. Requests for a cross-domain LU-LU session require that a session be established initially between the SSCPs in each domain by the access methods. As shown in Figure 3-7, the request for a session between LU 11 and a host LU in host processor H3 must pass through the host processor in both domains 1 and 3. After the LU-LU session is initiated, the data flow is directly between the LUs. Neither the session initiation/termination flow nor the data flow must pass through host processors in intervening domains.[1] The flows pass through only those network elements that are specified as the path between the end points. Cross-domain sessions can also be established with non-SNA stations; the LU functions for such stations are performed by the network control program.

Naming Resources

A key characteristic of the access methods for session establishment is that all references to network resources are by symbolic name. In single-domain networks, this characteristic frees the application program and network operator from concern about the location of the resources within the network. All address translation from the symbolic name to the actual network address is performed by the access method. Thus, the symbolic names used in each domain must be unique within each domain. The symbolic names of resources to be used across domains must be unique across all domains that may use these resources concurrently. The essential characteristic of a unique name is that it cannot concurrently define multiple resources to an access method. For example, the same name cannot be used concurrently to represent both a single-domain and a cross-domain resource.

Network Addressing and Routing

Subarea addressing is used to establish the path of data flow throughout the network. A unique subarea address is assigned to each host processor and each communications controller, as shown in Figure 3-8. Host LUs assume the subarea address of their host processor.

All remote logical units and stations, either switched or nonswitched, assume the subarea address of the ACF/NCP/VS by which the logical unit or station is controlled. Switched LUs and stations are assigned to a particular subarea only after connection is established and only for the duration of the connection.

---

[1]Exception: At the user's option, data may pass through intervening host processors on ACF/TCAM host-to-host data flows.

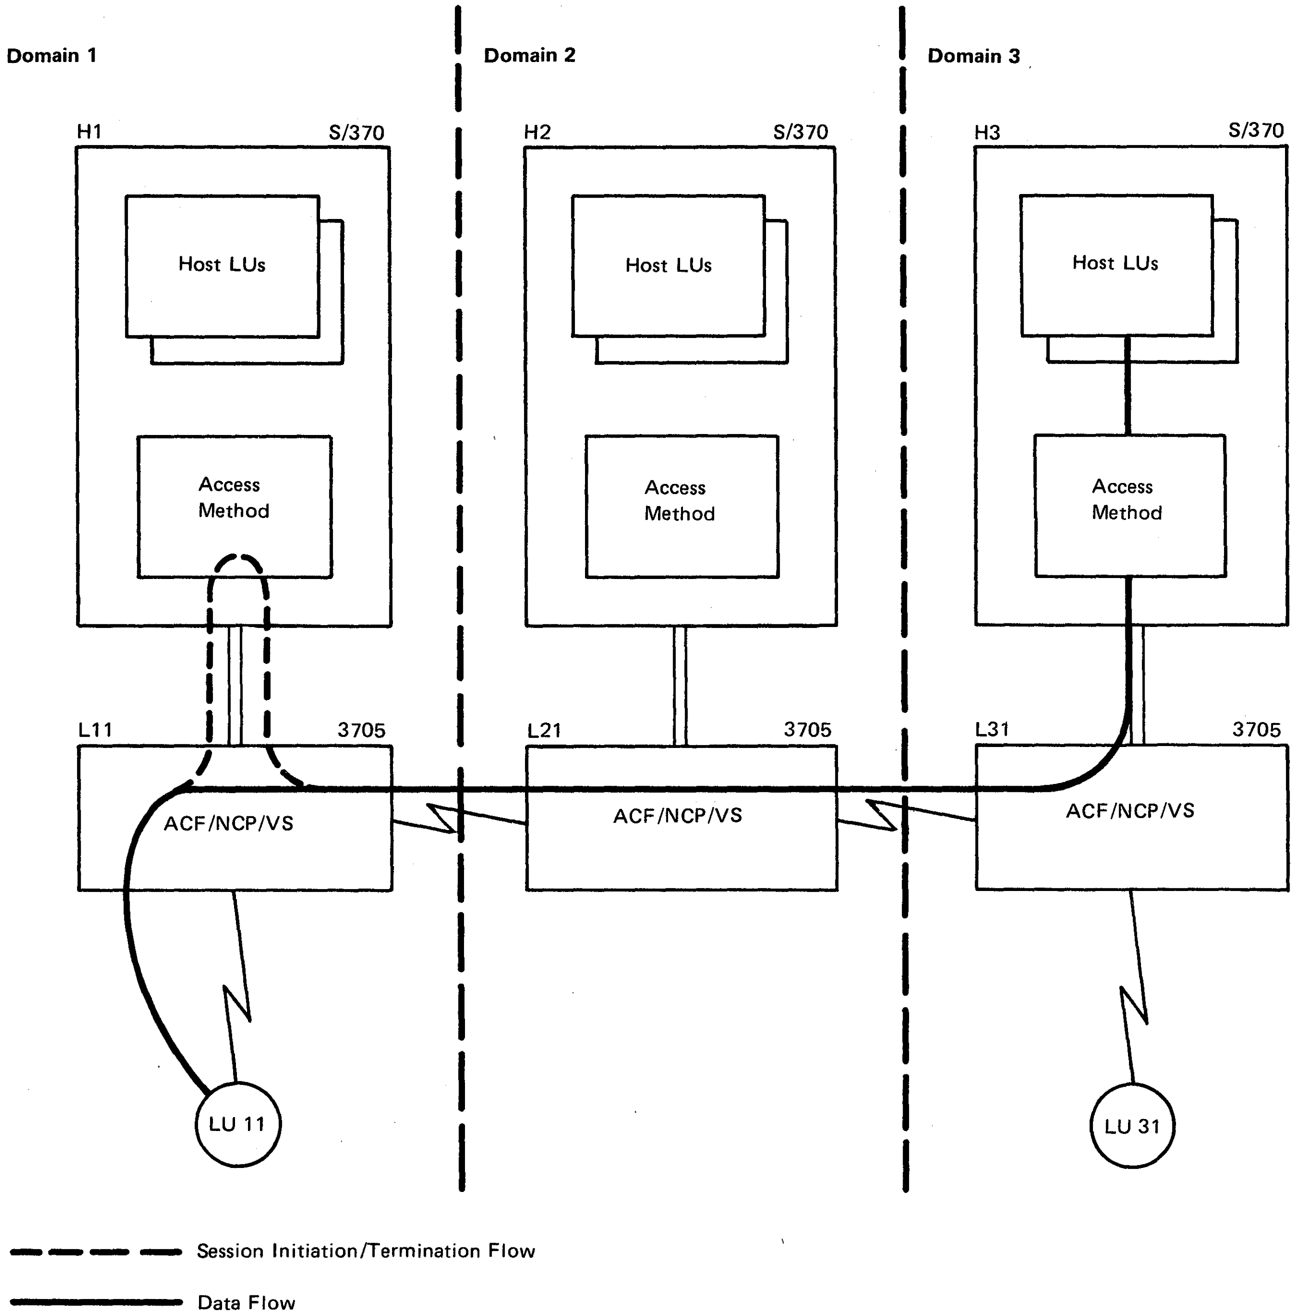

Figure 3-7. Session Initiation/Termination and Data Flows

Each host processor and ACF/NCP/VS must be able to select the appropriate adjacent subarea as a part of the path to any potential subarea address. Only one active path can exist between any two subareas in the network. (Release 3 of ACF/NCP/VS removes this limitation; see Chapter 6 for details.) This path must consist of the same elements (data links and physical units) in both directions between the subareas.

Intermediate Node Routing

In a multiple-domain network, data flows only through those network elements that are specified as the path between the two end points of a session. As shown in Figure 3-9, the ACF/NCP/VS is providing intermediate node routing for a cross-domain session. The ACF/NCP/VS can perform this function because it knows the subarea routing for each cross-domain data flow that it can receive. Upon receiving the data, the ACF/NCP/VS determines the appropriate adjacent subarea for the data flow and routes the data to that subarea. Only ACF/NCP/VS can perform intermediate-node routing; neither ACF/VTAM nor ACF/TCAM provides this function. (However, ACF/TCAM message handlers, in conjunction with ACF/TCAM host-to-host data flows, can be used to route data through an intervening host processor.)

Figure 3-10 shows a data link between two communications controllers in the same domain. Such a link, called a *same-domain link,* allows the ACF/NCP/VS in both controllers to provide intermediate node routing.

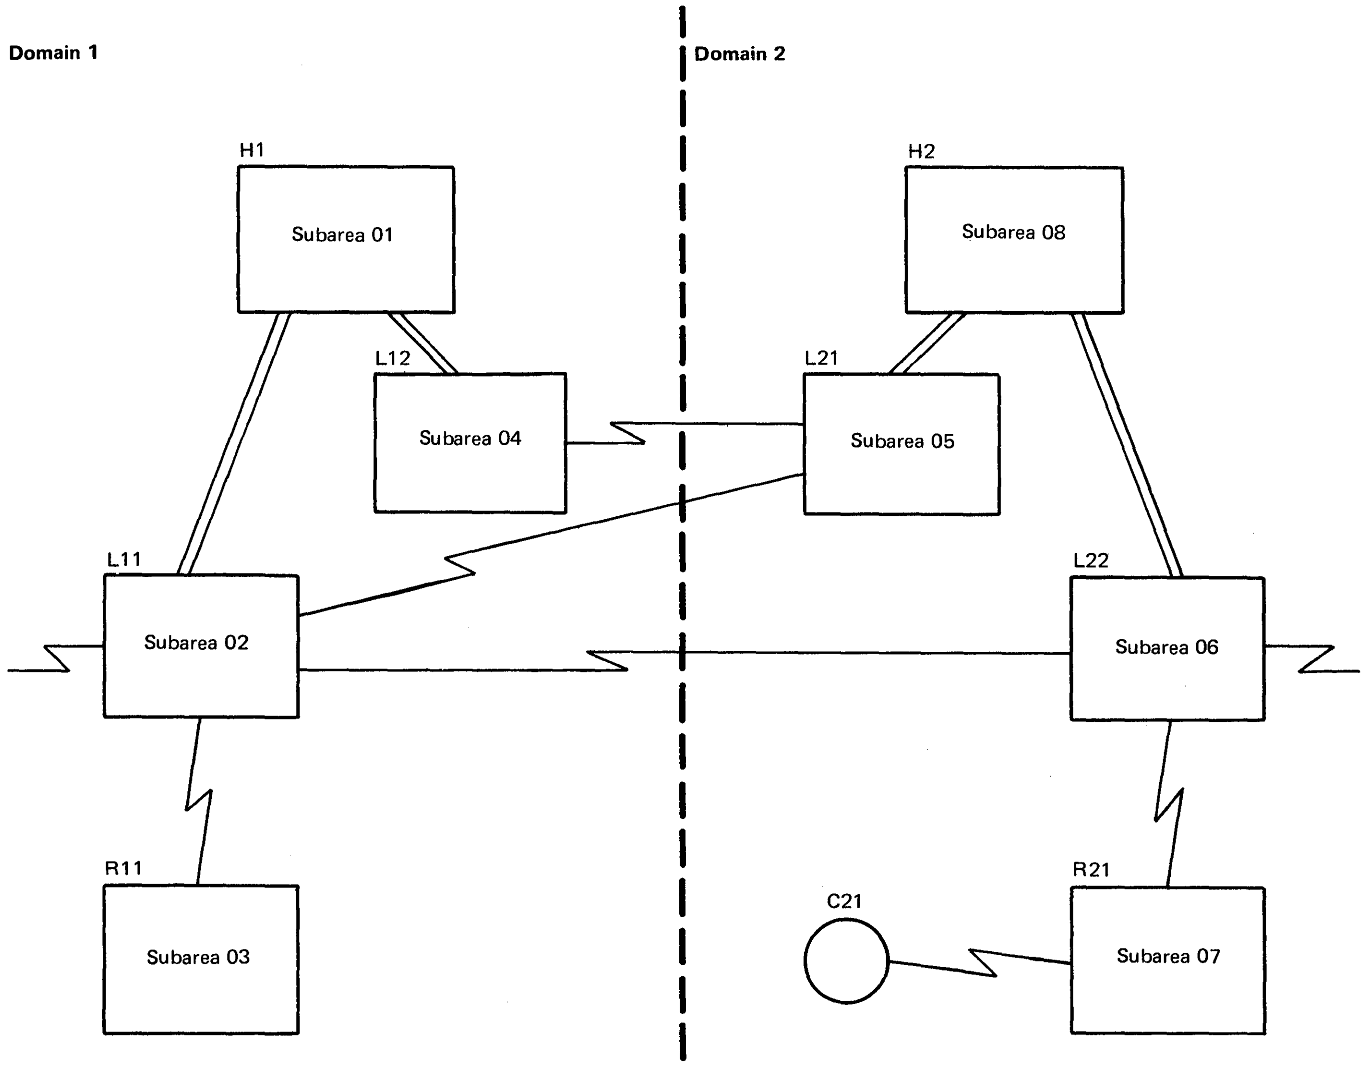

Figure 3-8. Configuration Having Multiple Paths between Domains

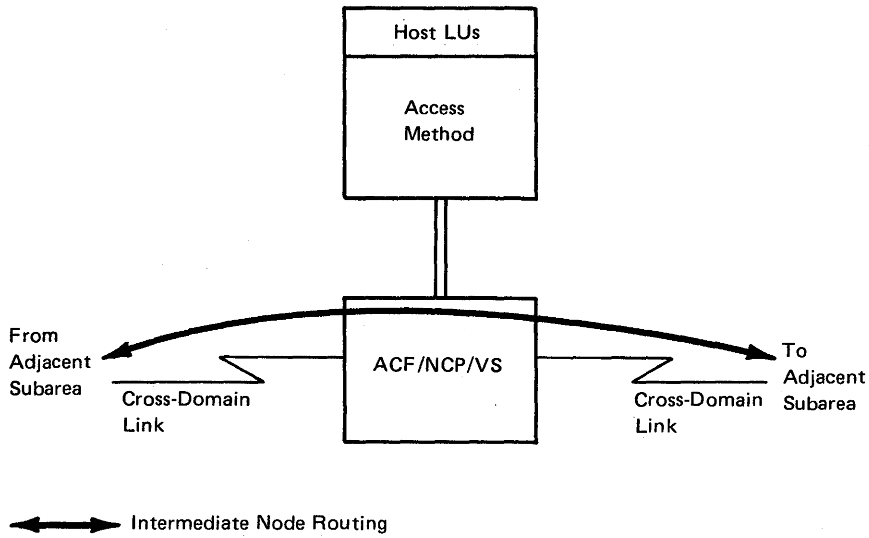

Figure 3-9. Intermediate Node Routing by ACF/NCP/VS

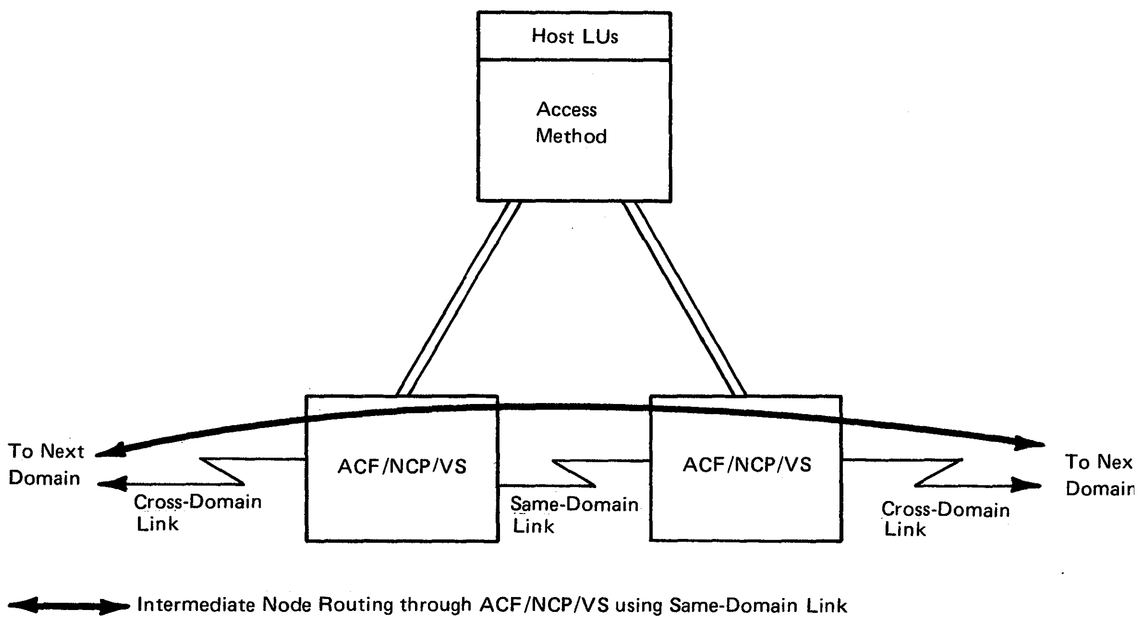

Figure 3-10. Intermediate Node Routing Between Local Communications Controllers in the Same Domain

Sharing Network Resources

The key capability provided by networking is extended sharing of resources in a multiple-host, multiple-location environment. For purposes of discussion in the following paragraphs, a four-domain, multiple-processor network, shown in Figures 3-11 through 3-13, is used to demonstrate cross-domain sharing. The configuration shown includes:

- A communications controller (L11) attached by data channels to host processors H1 and H2 using the multiple-channel attachment facility
- Cluster controller C32/41 on a switched line that can dial into, or be dialed by communications controllers L31 and L41

- Cluster controller C31 with two logical units, LU H and LU K, associated with the cluster controller.

The access methods in host processors H1, H2, H3, and H4 can be any combination of ACF/VTAM or ACF/TCAM with the Multisystem Networking Facility. The host operating systems can be either OS/VS1 or OS/VS2 MVS, or in the case of ACF/VTAM, DOS/VS. To use the multiple-channel attachment of communications controller L11, host processors H1 and H2 are at the same location. Host processors H3 and H4 can be at any locations.

Networking allows host LUs to communicate with other LUs in any part of the network, subject only to user-defined authorization. Networking can therefore eliminate redundant host LU functions and redundant special-purpose communication lines. Networking can also allow backup of critical host LU functions or devices.

How multiple-host LU functions can be avoided is illustrated in Figure 3-11. Host LU 2 in processor H2 can be shared by users in all domains: the function represented by LU2 need not be executed in each host processor. The heavy dashed lines in Figure 3-11 indicate the path required for session initiation and termination. After session initiation, data flows directly between the host LU and the station (logical unit).

Dedicated single-purpose data links can be replaced by shared links. For example, without networking, if logical unit LU B in domain 1 of Figure 3-11 needs to use a host LU in domain 3, a special connection for that purpose has to exist between domains 1 and 3. (In these descriptions, domain 1 is the domain whose system services control point [SSCP] resides in host processor H1, etc.) Similarly, if a user of logical unit LU F in domain 4 needs to use a host LU in domain 2, a special connection has to be provided. With networking, a set of common shared links are used rather than multiple special-purpose connections. The shared links permit complete access to host LUs in any domain by all cross-domain LUs in the network.

Figure 3-12 shows an LU-LU session between host LUs. Host LU 1 in processor H1 is in session with host LU 2 in processor H2. In this instance, the paths for session initiation and data flow are the same. Also shown in Figure 3-12, host LU 2 is in session with logical unit LU H in domain H3. The access method in processor H3 participates in session initiation between LU H and host LU 2, although the session initiation path is not shown. The heavy solid line in the figure shows the data flow after a session is established.

If a cluster controller contains multiple logical units, those logical units can be concurrently in sessions with host LUs in different domains. In Figure 3-12, logical unit LU H is in session with host LU 2 while logical unit LU K is in session with host LU 3.

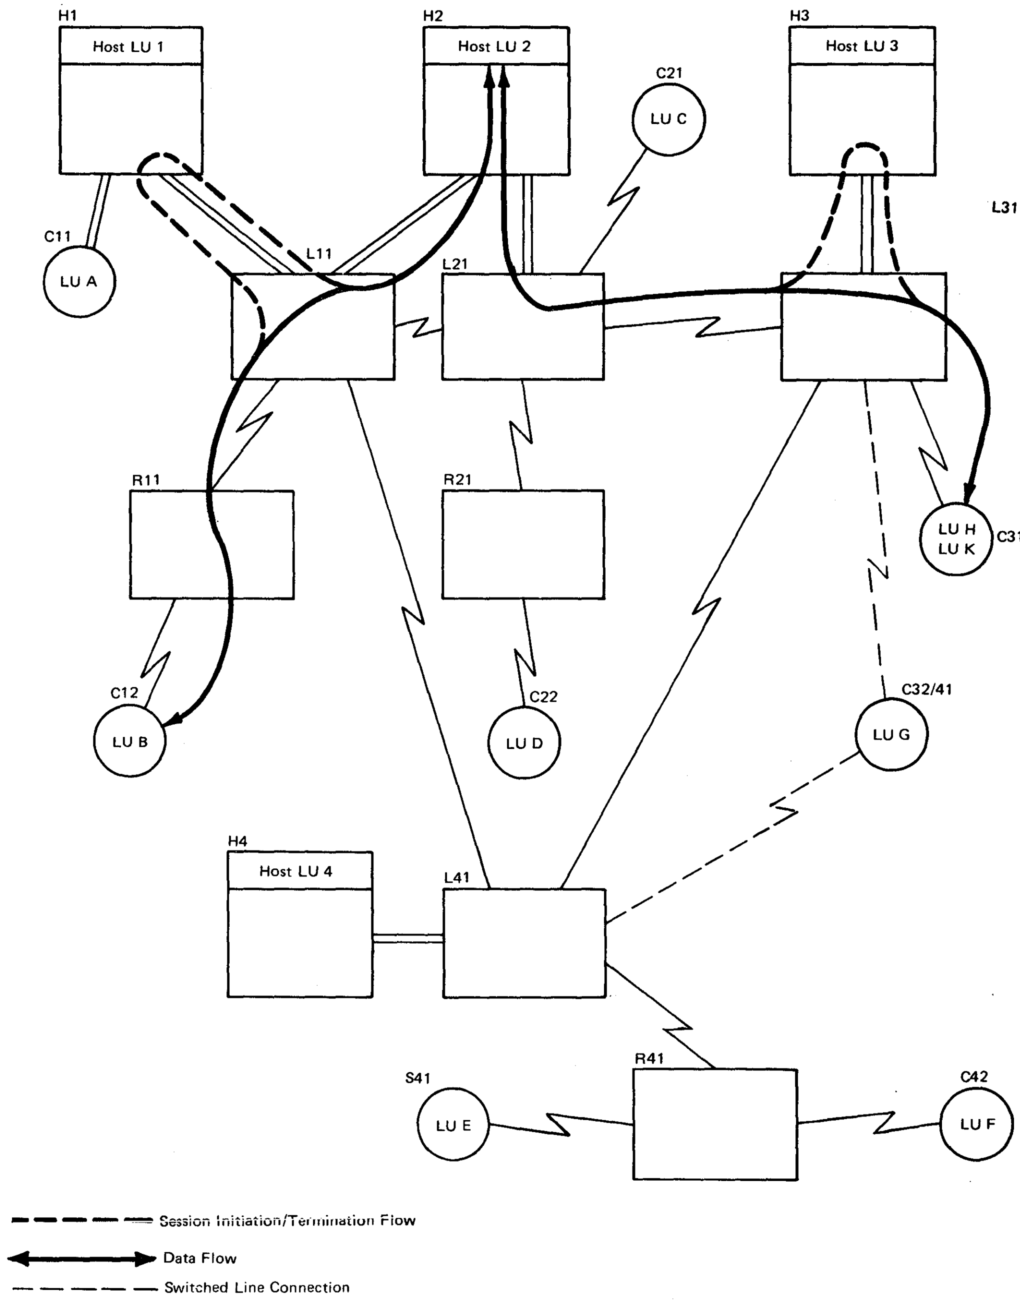

Figure 3-11. Sharing a Host LU

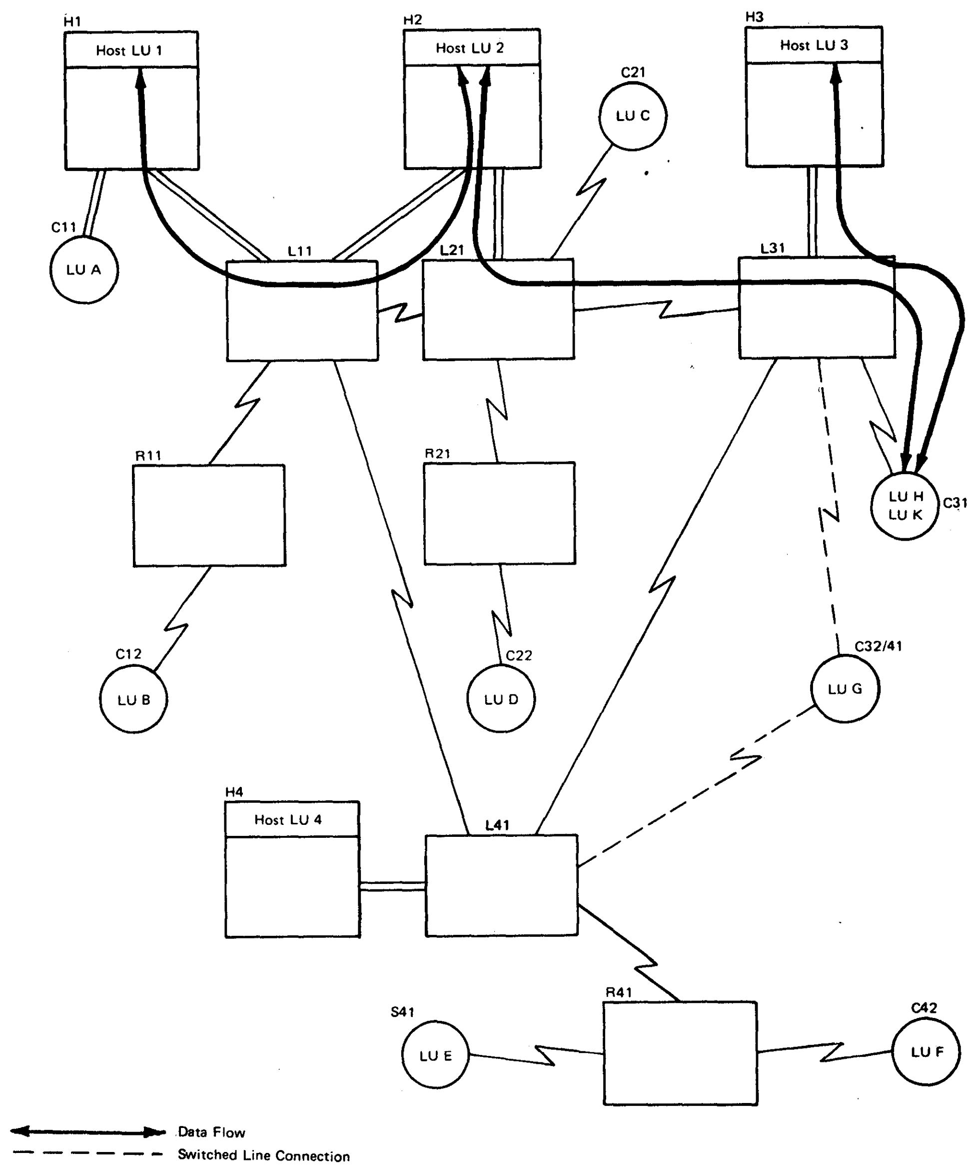

Figure 3-12. Access to Resources across Domains

A session can span intervening domains. For example, in Figure 3-13, a session between logical unit LU E in domain 4 and host LU 2 in domain 2 crosses intervening domain 3 using the intermediate node routing capability in communications controllers L41 and L31. Again, the heavy solid line shows data flow after session establishment. The light dashed lines between cluster controller C32/41 and communications controllers L31 and L41 indicate that the cluster controller is on a switched line and can dial into, or be dialed by, either communications controller L31 (domain 3) or L41 (domain 4). The cluster controller resources (PU and LUs) will be owned by the SSCP in whichever domain the cluster controller dials into.

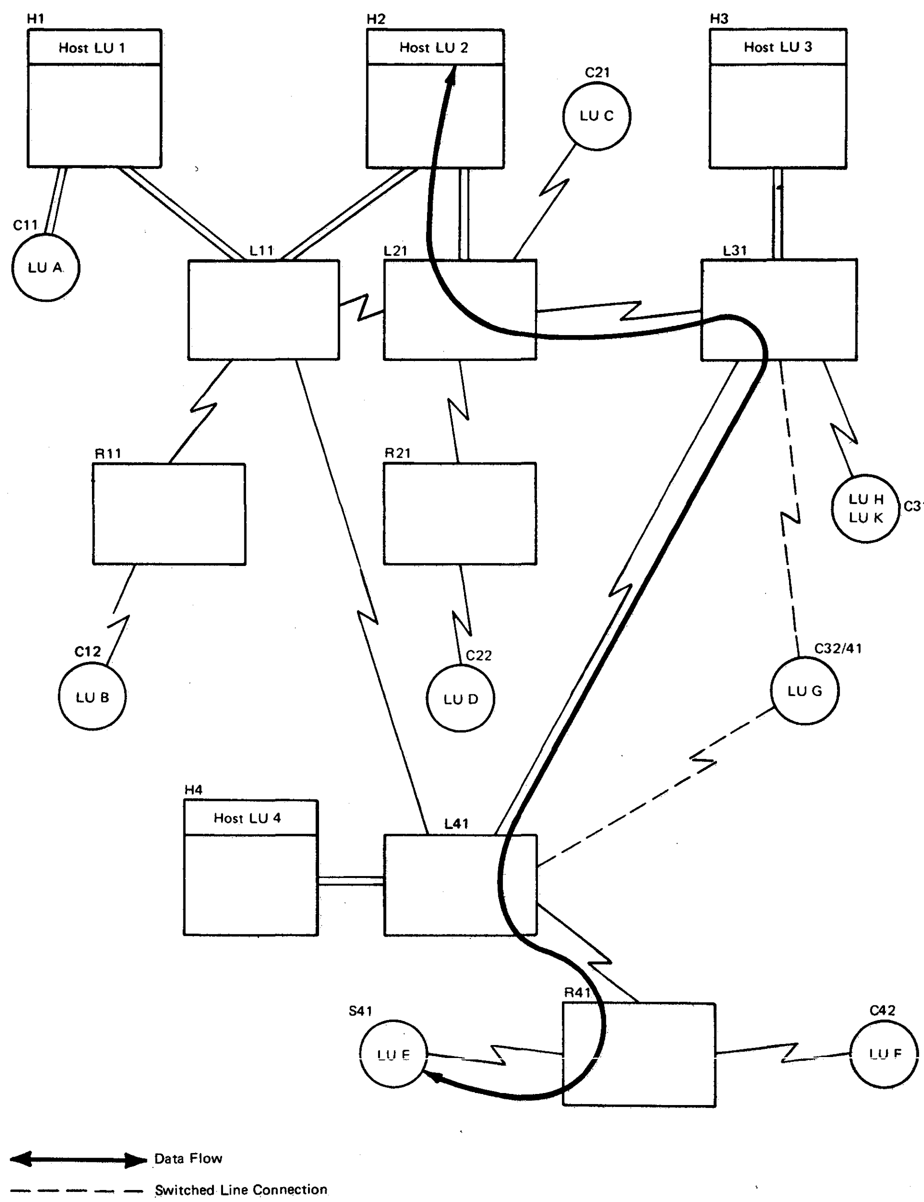

Figure 3-13. Data Flow through an Intervening Domain

Use of a Router or Message Handler

A user-written ACF/VTAM router application program or an ACF/TCAM message handler (router) can be used to provide cross-domain communication for stations not eligible for cross-domain sharing.

Ineligible stations for ACF/VTAM include local non-SNA stations and remote non-SNA stations. (Exception: Local 3270s and non-SNA stations supported via the Network Terminal Option *are* eligible for cross-domain sharing.) Ineligible stations for ACF/TCAM include local non-SNA stations and remote stations on lines operating in emulation mode. The ACF/TCAM message handler is insensitive to whether the routing is across domains or within the same domain.

An example of an ACF/VTAM router application program or an ACF/TCAM message handler is shown in Figure 3-14. The purpose of the router in host H1 is to relay messages between one or more stations and one or more host LUs in other domains (or in the same domain). The router in host H1 is concurrently in session with logical unit LU A in local cluster controller C11, remote non-SNA station S11 in domain 1, and host LUs in hosts H1 and H3.

The characteristic that differentiates this routing capability from networking is that the data flow from the non-networking station or LU is to a host LU (the router) in its own domain. The cross-domain session is provided by the router.

Implementing a Communication Management Configuration

Users having multiple System/370 computers controlling several data communication networks can consolidate many of the network management functions for all networks into a single processor, thereby freeing other processors at the same location for application processing.

Figure 3-15 shows a possible implementation of such a configuration with host H1 and communications controller L1 providing the communication management function. In this configuration all network resources other than host LUs and local stations belong to host H1. These resources include remote communications controllers, data links, and remote cluster controllers and terminals. Host H1 is responsible for such network management functions as activation, session initiation/termination, configuration changes, and deactivation. Hosts H2, H3, and H4 are attached by channels to communications controller L1 using the multiple-channel attachment facility of the 3705-II and the shared ACF/NCP/VS capability. Since most network management functions are performed by host H1, hosts H2, H3 and H4 can be dedicated primarily to application processing. The access methods in use can be any combination of ACF/VTAM or ACF/TCAM with the Multisystem Networking Facility. The access method in each of the hosts H2, H3 and H4 is required to provide the connection between its host LUs and the network resources using the shared ACF/NCP/VS capability of communications controller L1. The entire configuration can be interconnected with other domains, making possible full networking participation with the interconnected domains.

Since all host processors are System/370 computers with either DOS/VS, OS/VS1, or OS/VS2 MVS (as appropriate for the access method), any excess processing capability in any host processor is readily available for other applications. If required, the communication management function can be taken over by one of the application processors. As with the choice of operating systems, the model of computer required for the communication management configuration can be any model that is supported by the access method and is consistent with operational requirements.

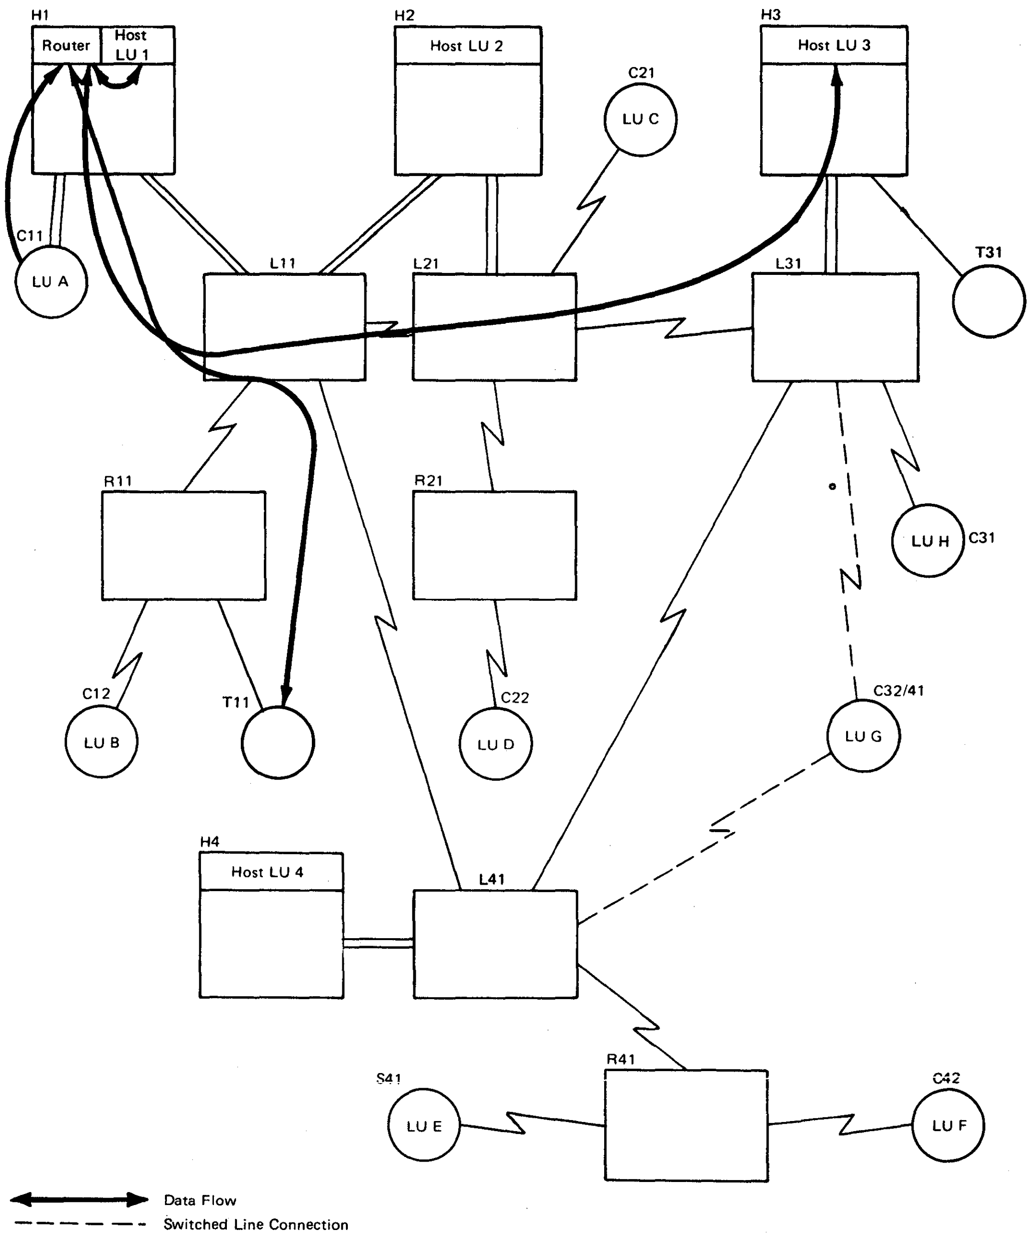

Figure 3-14. Use of a Router or Message Handler

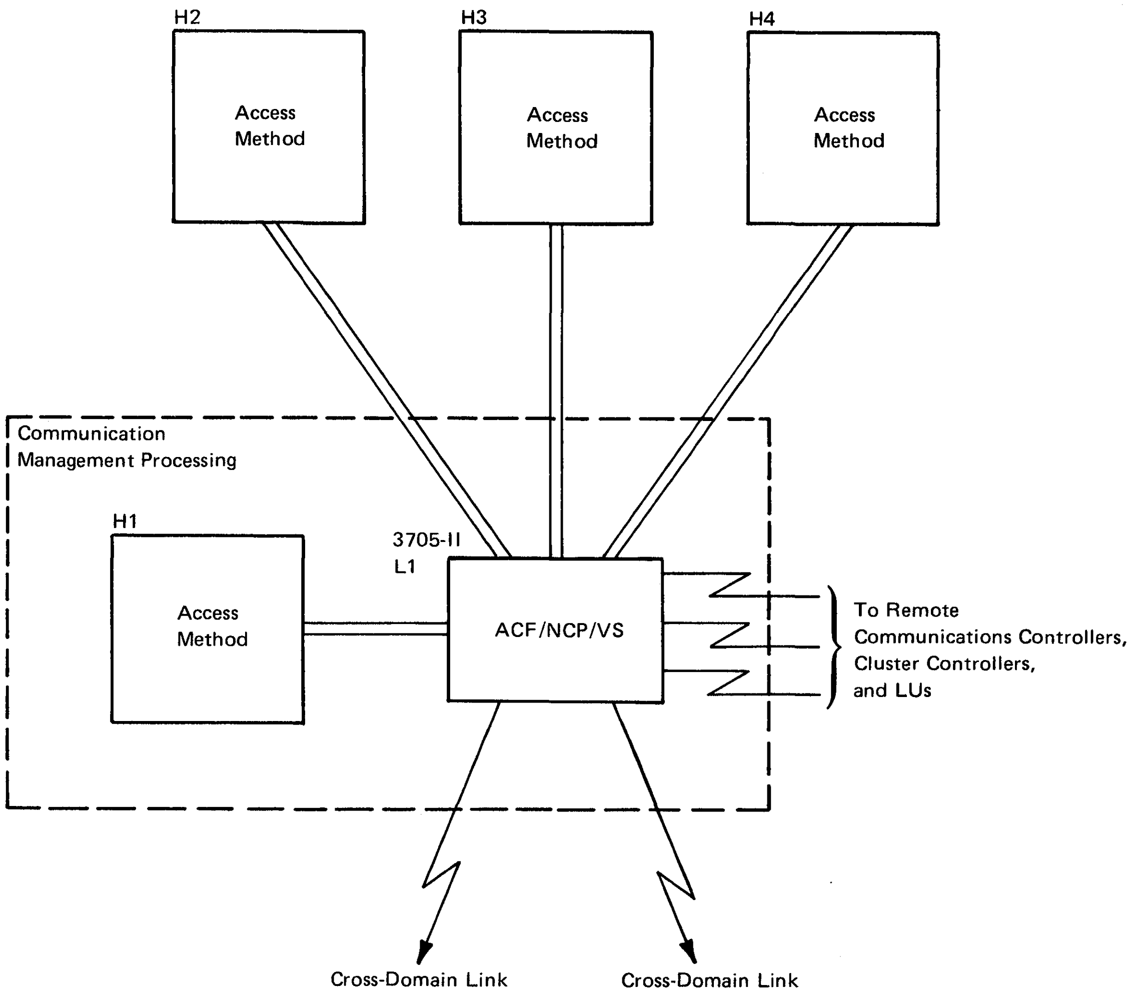

Figure 3-15. Possible Implementation of a Communication Management Configuration

Network Control

Complex networks impose significant problems in the area of network control. The solutions to these problems require careful analysis of network requirements, proper selection of user options available when defining the network structure, and development of operational procedures for controlling the network that will effect the desired results under both normal and recovery conditions.

ACF/VTAM and ACF/TCAM include operator commands and status displays for use in controlling each single-domain network. In addition, ACF/VTAM and ACF/TCAM provide status displays of all resources identified as cross-domain resources. The operator commands and status displays can be used with operational procedures to control all network elements in single and multiple domains under all network conditions with minimum disruption to established LU-LU sessions. Options that are available when the network is specified include:

- Selection and identification of same-domain and cross-domain resources
- Predefined backup configuration definitions

- Routing paths
- Symbolic names for same-domain and cross-domain resources.

ACF/VTAM allows an authorized application program (a *program operator)* to issue network operator commands and to receive appropriate responses. ACF/VTAM program operators can use this facility in cross-domain sessions to permit an operator in one ACF/VTAM domain to control resources in another ACF/VTAM domain. In addition, three complementary program products are available for use with single-domain or multiple-domain network configurations. One of these is the Network Operation Support Program (NOSP), which is available only for ACF/VTAM.

The other two program products are the Network Communications Control Facility (NCCF) and Network Problem Determination Application (NPDA). They may be used by either ACF/TCAM or ACF/VTAM users to aid in controlling network operation and handling problem determination. (NCCF is the successor to NOSP; NCCF includes the facilities provided by NOSP and offers additional facilities as well).

Detailed information on the benefits of these program products can be found in their respective General Information publications, which are listed in the Preface.

ACF/TCAM is similar to ACF/VTAM in that it permits any authorized station or application program to issue operator control commands and effect cross-domain operator control with other ACF/TCAM domains.

Availability and Recovery in a Multiple-Domain Network

The characteristics of the multiple-domain networking environment make it possible to maximize the availability of network resources. Continuity of LU-LU sessions is preserved whenever possible. When a network element participating in an online application becomes inactive because of a failure or deactivation, an alternative element, if available, can continue to support that application.

Releases 1 and 2 of ACF/NCP/VS provide the following backup capabilities for multiple-domain networks. Additional backup capabilities are available in Release 3 of ACF/NCP/VS, as described in Chapter 6.

Continuation of Intermediate Node Routing Function

Assume that host processor H2, in the configuration shown in Figure 3-16 *(a)* fails or is deactivated, and that communications controller L21 remains active. Upon loss of contact with host processor H2, controller L21 enters automatic network shutdown mode. During shutdown, cross-domain links remain active and cross-domain sessions using the intermediate node routing function of controller L21 between host processors H1 and H3 may continue. Cluster controllers and LUs on nonswitched lines in domain 2 (host processor H2) optionally can be allowed to continue cross-domain sessions with host processors H1 and H3.

Upon recovery, host processor H2 will activate communications controller L21. The recontacting procedure does not require reloading the ACF/NCP/VS. Cross-domain sessions between domains 1 and 3 through controller L21 are not affected by the restart.

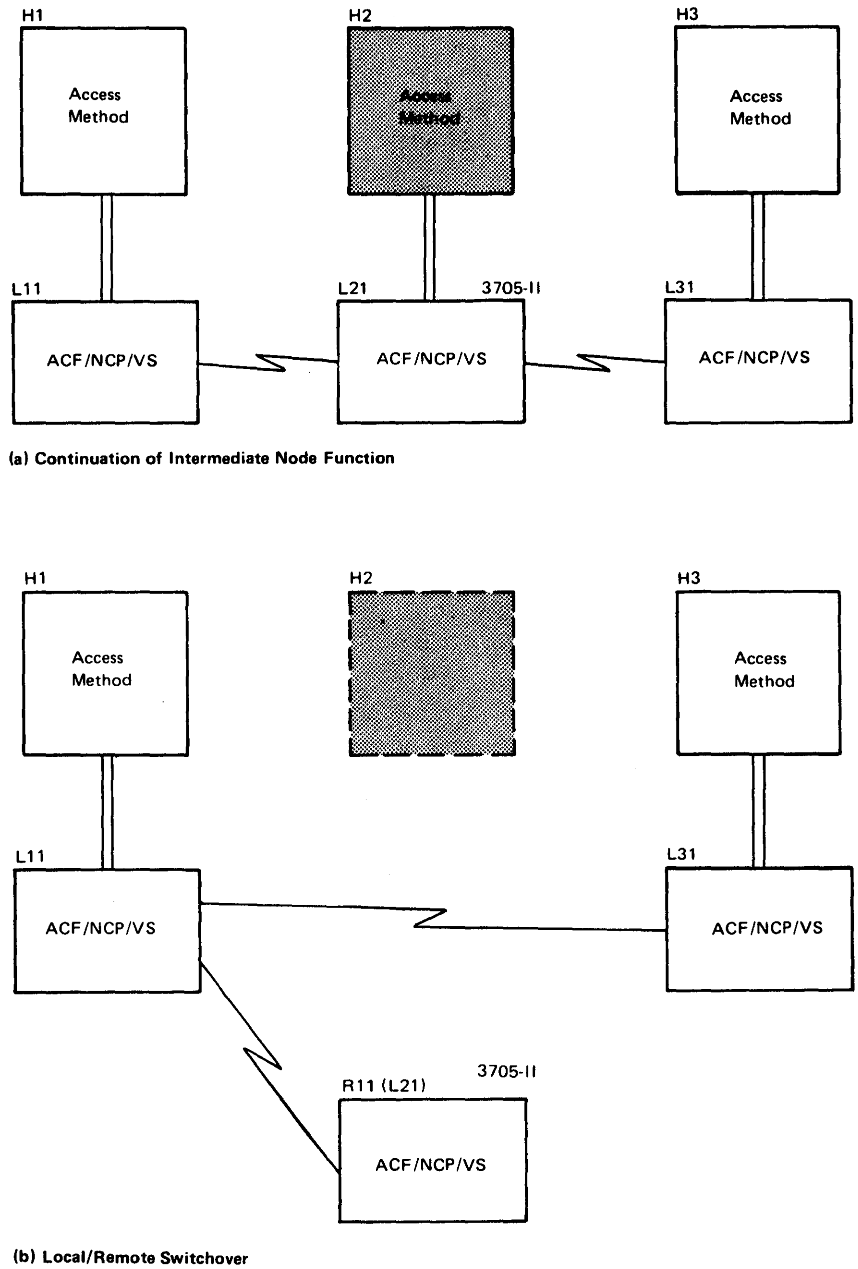

Figure 3-16. Recovery from Host Processor Failures

Host Processor Backup Capabilities

There are several options for establishing backup host processor relationships. These capabilities are especially useful in the event of a host processor failure but can be useful in other situations as well. These capabilities are:

- Taking over the resources of a communications controller by an adjacent-domain host processor
- Reconfiguring a local 3705-II communications controller to become a remote communications controller

- Connecting a remote communications controller to a local communications controller other than the one to which it was originally connected.

Resource Takeover by an Adjacent-Domain Host Processor: An operating local communications controller and its resources can be taken over by an access method in an adjacent domain. An adjacent domain is one that either shares a multiple-channel-attached communications controller, or is one cross-domain link away. Following acquisition of the communications controller by the access method in the new host processor, the resources of the communications controller may be activated incrementally. The takeover may be staged so that cross-domain sessions are not disrupted. The capability of ACF/NCP/VS to be shared by multiple SSCPs allows the takeover to be shared by multiple adjacent-domain host processors. Cross-domain links can also be taken over by a host processor performing the backup function.

Local-to-Remote Reconfiguration: The optional local/remote capability of the 3705-II communications controller allows a local communications controller to be taken over from another host processor as a remote communications controller in the new domain (see Figure 3-16 *b*.) This action changes the characteristic of the 3705-II from a local communications controller (L21) in domain 2 to a remote communications controller (R11) in domain 1. The changeover is initiated by an operator from the domain of the host processor taking over the communications controller and requires that the controller be reloaded. As a remote communications controller, the 3705-II does not support any cross-domain links. (Release 3 of ACF/NCP/VS removes this limitation; see Chapter 6 for details.) If the cross-domain link between communications controllers L11 and L31 is desired, the link can be manually dialed and both L11 and L31 reloaded with updated path routing information.

Connecting a Remote Communications Controller to a Different Local Controller: A remote communications controller can be connected to a different local communications controller in the same or a different domain, if the receiving communications controller has the proper network address defined in its ACF/NCP/VS. This action is accomplished by activating a data link to a new local communications controller and reloading the remote communications controller. If the changed remote communications controller is given a different subarea address, neither local communications controller needs to be reloaded.

Communications Controller Backup

If a local communications controller becomes inoperative, the inoperative condition is detected by all adjacent network elements. Only sessions requiring access to the inoperative communications controller are disrupted. If a remote communications controller is connected to the inoperative local controller, the remote controller enters automatic shutdown mode. Logical units may clean up processing for that session and wait for network recovery. The network operator may intervene to recover the capabilities of the inoperative controller by either reloading the communications controller or manually switching to a backup communications controller.

Cross-Domain Link Backup

The network operator can effect recovery from inoperative cross-domain links by activating a backup nonswitched link or by manually dialing a switched backup link. Network operator intervention is required in both domains to initiate activation of the appropriate cross-domain link. The communications controllers at each end of the failed link do not require reloading. The only disruption is to cross-domain sessions that use the failing link. The logical units for such sessions are notified of the failure if that failure has not also disrupted the SSCP-PU and SSCP-SSCP sessions used to notify SSCPs of the cross-domain session disruption.

Chapter 4. System Support Programs for ACF/NCP/VS

The IBM-supplied system support programs for ACF/NCP/VS are used for preparing, installing, and using ACF/NCP/VS. These programs are (1) a procedure for generating the ACF/NCP/VS, (2) an assembler, (3) a loader, and (4) a dump program, and (5) a dynamic dump program. The loader and the dump programs are utility programs.

The system support programs will run under DOS/VS, OS/VS1, or OS/VS2 (MVS) in a System/370.

Consult your IBM representative to determine the availability of the system support programs under each of the operating systems referred to above.

The generation procedure and the assembler are executed entirely in the System/370. (This need not be the host processor with which the ACF/NCP/VS will communicate.) Each of the utilities is divided into two portions, one of which runs in the host processor and the other in the communications controller.

ACF/NCP/VS Generation

The generation procedure for preparing a network control program includes (1) a generation language by which network configuration and program options are specified, and (2) a library of macro definitions from which the source statements are expanded. The procedure is similar for OS/VS and DOS/VS. Described below are the generation language, some coding conventions, and how the generation procedure works.

The Generation Language

The generation language provides a high-level means for generating the network control program for the 3705. The language is designed to minimize the programming effort for even the most complex configurations of lines and stations.

The generation language is made up of macro instructions that fall into four categories according to the type of parameters they define. The types of macros are (1) system macros, (2) configuration macros, (3) block handling (BH) macros, and (4) a generation delimiter macro.

System Macros

The system macros provide information pertaining to hardware features with which the communications controller is equipped, certain network control program options, and program generation information such as data set names. The system macros specify, for example:

- The amount of storage installed in the communications controller
- The size of buffers the network control program contains
- The type of channel adapter installed
- Optional facilities such as online terminal test
- Optional dynamic control functions to be included in the network control program
- The identifier of the controller

Configuration Macros

The configuration macros provide the information necessary to construct the tables needed by the network control program to control the flow of data between the 3705 and remote stations or other 3705s and between the 3705 and the host processor.

One group of these macros defines the characteristics of the elements in the network—line groups, communication lines, cluster controllers, terminals, components of terminals, and SNA physical and logical units. A macro is coded for each element in the network. The macros must be arranged in a specific order to associate a particular communication line with a particular line group, a particular station or SNA physical unit with a particular communication line, etc.

The rest of the configuration macros provide the following types of information:

- Information needed for data transfer between the communication access method and the network control program; for example, average block size and buffer-unit size in the host processor
- Information describing the communication scanner(s) installed in the 3705
- Definition of the remaining control tables necessary to control the network; for example, lists of valid identification sequences for binary synchronous stations that call in over the switched communication network

Block-Handling Macros

The block-handling (BH) macros apply only to messages transmitted over binary synchronous and start-stop communication lines. They describe optional processing that the network control program can perform on a block of data before transferring the block to a station or to the host processor.

Some BH macros define *block-handling routines* (BHRs) that perform specific processing functions. The routines specified by these macros can:

- Insert the date and/or time of day into blocks of data.
- Correct text incorrectly entered from a station. The macro defines the character to be recognized by the BHR as a backspace character. The BHR deletes these characters from the text and overlays the characters preceding the backspace characters with the text that follows.

  Example:
  From station: CHARACTRE bksp bksp ER
  After processing by BHR: CHARACTER

Using the controller assembler, additional block-handling routines can be written to process blocks in other ways. A block-handling macro allows inclusion of these routines in the network control program at the time the program is generated.

The remaining block-handling macros provide for the grouping of block-handling routines into block handlers and sets of block handlers. A *block-handler* consists of one or more block-handling routines defined by the individual block-handling macros. Many block handlers can be defined for a single network control program configuration. Special block-handling macros delimit the beginning and end of each block handler and provide a symbolic name for it. When multiple block handlers are defined, one must be completed before the next is defined.

Up to three block handlers are grouped into a *block-handler set,* defined by another block-handling macro. Each block handler in a set can be executed at one of three points in time, as follows:

1. After a command has been received from the host access method for a station but before the communication line is available
2. After a command has been received from the host access method for a station, but only after the communication line is available
3. When an input operation on a communication line ends

Each block-handler set may be associated with one or more stations by coding the name of the set as an operand of the configuration macro that defines the station.

Generation Delimiter Macro

The generation delimiter macro ends the program generation input stream.

Coding the Generation Language

The generation language is designed to make coding as easy as possible. All the operands of the individual macros are keywords, so the programmer need not be concerned with the sequence in which he codes the operands. However, the relative order of the macros in the input stream is to some extent fixed.

General Logic Flow of the Program Generation Procedure

When the programmer has coded the generation macros that describe the network to be controller by ACF/NCP/VS, he generates the control program using the program generation procedure. The generation procedure is a two-stage process, for OS/VS, or a three-stage process, for DOS/VS. The procedure is executed as a series of jobs in the host processor (or other System/370). Figure 4-1 illustrates the procedure for program generation under OS/VS.

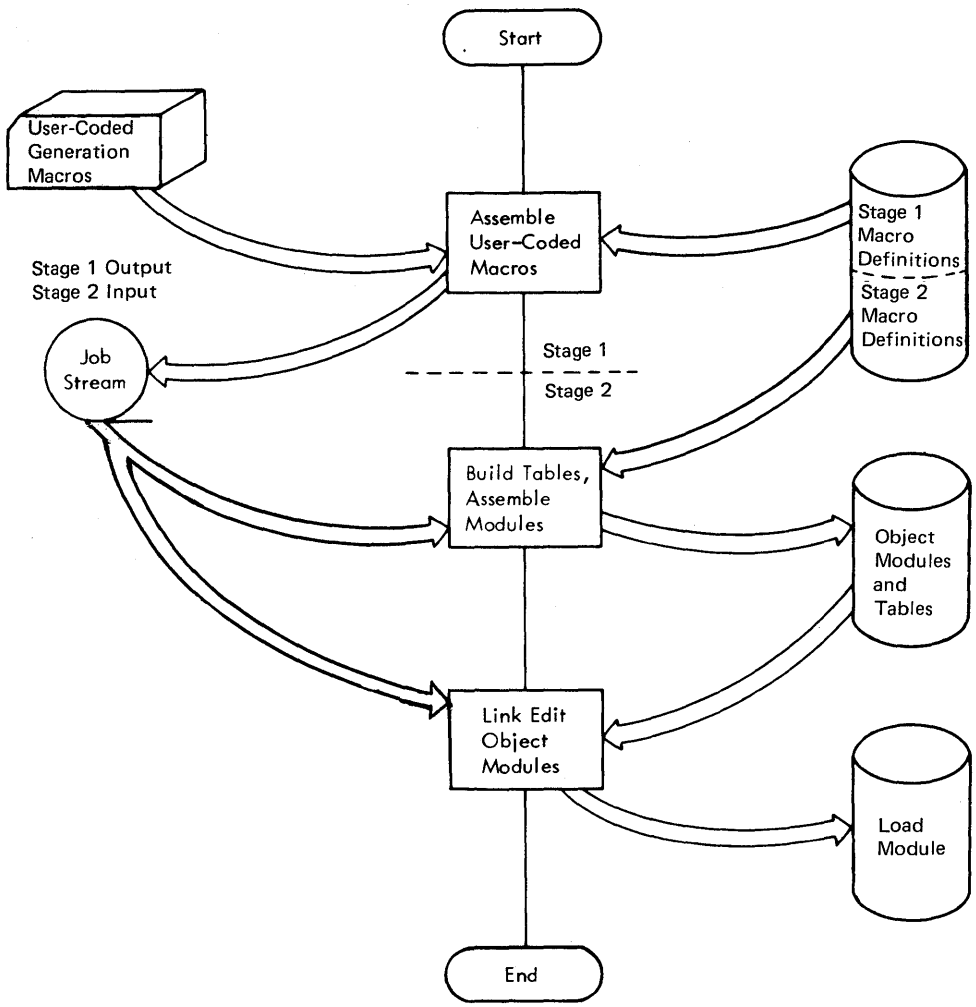

Figure 4-1. Example of the Program Generation Procedure under OS/VS

Stage 1

In the first stage of the generation procedure, the source statements (generation macros) are assembled by the controller assembler. (If generation is under OS/VS, an OS/VS assembler, instead of the controller assembler, may be used for stage 1 assembly. Generation under DOS/VS requires the controller assembler.) For OS/VS and DOS/VS, output from the assembly is a job stream containing the data and control statements necessary to create the desired network control program. The job stream is a sequential data set that can be directed to cards, tape, or a direct-access storage unit.

Intervention between the Stages of Program Generation

Intervention is required between stages of the generation procedure. If there are errors in the source statements entered as input to stage 1, a diagnostic message is issued for each statement that contains an error. For severe errors, the source statements must be corrected and resubmitted until no severe errors remain.

Stage 2

The second stage of the generation procedure, executed when stage 1 has been completed with no severe errors, creates the network control program.

For OS/VS, the job stream from stage 1 contains the data necessary to select the appropriate program modules and build the proper tables. Using the information coded in the generation macros, stage 2 first builds the tables and then selects and assembles (using the controller assembler) those modules that are dependent on the defined network. Then the OS/VS linkage editor is executed to combine the appropriate modules into the ACF/NCP/VS load module. Finally, the load module is stored on a direct-access storage unit.

For DOS/VS, stage 2 of the generation procedure selects and assembles the control tables and those program modules that are dependent on the defined network; it then creates job control statements and linkage editor statements for stage 3. Stage 3 catalogs the tables and modules assembled in stage 2 and link-edits them into an ACF/NCP/VS load module. The CSERV utility must then be executed to move the load module to a direct-access data set from which the loader utility may obtain it.

Generating Multiple Network Control Programs

As many network-control-program load modules as desired can be generated for a 3705. Each program requires a separate generation, and each must have a different symbolic name in order that the loader can identify the load module to be transferred into the 3705.

Multiple load modules are useful for installations that have several distinct applications for the network at different times. For example, if an installation uses only start-stop communication lines during the day and only binary synchronous lines at night, separate network control programs for the separate configurations could reduce the amount of storage required for the controller and make the program operation more efficient.

The Communications Controller Assembler

The communications controller assembler is available to assemble programs written in controller assembler language. In its external structure, it is very similar to the OS/VS and DOS/VS assemblers.

The assembler operates on three kinds of instructions: (1) machine instructions (written in controller assembler language notation); (2) macro instructions; and (3) assembler instructions. The assembler translates the machine instructions and the macro instructions into executable object code. The assembler instructions direct the assembler to perform certain operations during the assembly process, but the instructions are not converted into executable code. These three types of instructions are similar to the types of instructions processed by the OS/VS and DOS/VS assemblers.

The Instruction Set

The instruction set for the 3705 consists of 52 machine instructions. The instructions are represented to the assembler by mnemonic operation codes, usually followed by one or more operands. Most of the machine instructions are register-oriented. That is, they represent operations involving two registers, a register and immediate data, or a register and a storage area.

The assembler converts the machine instructions into 2 or 4 bytes of object code, depending on the length assigned to the particular instruction.

The publication *IBM 3704 and 3705 Principles of Operation* (GC30-3004) explains each of the machine instructions and gives the assembler language statement that corresponds to each.

Macro Capability of the Assembler

The macro language for the communications controllers is an extension of the controller assembler language. It provides a convenient method of generating a desired sequence of assembler language statements many times in one or more programs. Macro definitions can be coded in line or in assembler-language programs, or they can be stored in a host library and called in when needed by means of a macro instruction coded in the program.

The Assembler Instructions

Instructions to the controller assembler are written as assembler pseudo-operation codes, with or without operands. These instructions perform such functions as delimiting the beginning and end of section code, defining data areas, controlling the format of listings, and specifying base registers.

Uses of the Assembler

The uses of the controller assembler include: (1) assembling the ACF/NCP/VS generation macros and application-dependent modules during the program generation procedure and (2) preassembling user-written block-handling routines.

The assembler enables the programmer to augment the IBM-supplied ACF/NCP/VS modules with block handling routines that are unique to his installation's applications (for start-stop and binary synchronous communication lines). Using the controller assembler language, the programmer codes block-handling routines to process the data in message blocks going to or coming from a station. Then he uses the assembler to create object modules that are stored in the same library with the IBM-supplied ACF/NCP/VS object modules. During program generation, if the appropriate ACF/NCP/VS macros have been coded, the block-handling routines are link-edited together with IBM modules to form the ACF/NCP/VS load module.

The Utilities

The system support programs for ACF/NCP/VS include two utility programs: a loader and a dump program. Each is controlled by the appropriate job control statements and control cards.

The Loader

The loader has two loading functions: (1) it transfers a diagnostic routine, the *initial-test* routine, into the 3705; and (2) it transfers the network control program from host secondary storage into the 3705.

The initial-test routine is loaded before the ACF/NCP/VS load module is loaded into the 3705. This routine tests the hardware for conditions that could possibly result in failure of the 3705 after network operation begins. If the initial-test routine discovers any exceptional conditions, it causes a hard stop of the 3705 and cancels transfer of the ACF/NCP/VS load module across the channel. Indicators on the 3705 control panel are set to aid in isolating the problem.

Whenever the loader is invoked, the initial-test routine is executed automatically unless it is suppressed by a utility control card entered as input to the loader.

The loader is invoked to load the network control program during start-up of the 3705 and when the 3705 fails because of some error condition. In either case, the operator starts the loader with job control statements in the job stream.

Both ACF/TCAM and ACF/VTAM have optional loader routines that may be used instead of the loader supplied with the system support programs to load the network control program into the 3705 during start-up or after 3705 failure.

How the Loader Operates

The part of the loader that resides in the host processor handles all external input and output requirements of the loading process. This portion reads the ACF/NCP/VS load modules from secondary storage and issues a Write command for each block of code to be transferred across the channel into the 3705.

The portion of the loader that resides in the 3705 initializes the 3705 to prepare it for the data written from the host processor. It then communicates with the host portion, accepting blocks of code from the channel and positioning then appropriately in 3705 storage.

The Dump Program

The dump program dumps the contents of 3705 storage to help a programmer to isolate and correct error conditions that arise. These options are available when requesting a dump:

- The limits of the storage area to be printed out from the dump can be specified. Otherwise, the program produces a printout of the complete dump (except for a small area at the beginning of storage).
- A formatted dump of the network control program can be requested. In this case, the dump program isolates and labels certain control blocks, printing them at the beginning of the dump.

Both formatted and unformatted dumps contain a hexadecimal representation of 3705 storage. In addition, all dumps include the contents of the general registers and the EBCDIC representation of all letters and numbers in the dump.

After any dump (whether it be a complete storage dump or a dump of only a portion of storage), the network control program must be reloaded into the 3705 before network operation can be resumed.

How the Dump Program Operates

The dump program for OS/VS has two job steps. The first step, which requires code in both the host processor and the 3705, dumps the entire contents of 3705 storage and the contents of the general registers to a data set on host disk storage.

The first step then automatically invokes the second step, which runs entirely in the host processor. Step 2 first analyzes the utility control cards, on which the programmer has specified the options desired for this dump (storage limits, formatting, etc.). Printouts of as many different areas of the dump as desired can be requested. Then step 2 formats the dump when appropriate; reads the requested contents from the disk data set; translates them into printable hexadecimal characters; and, finally, writes the requested portion(s) of the dump to an output data set to be printed out. (Step 2 can also be executed independently as many times as desired to produce extra dump printouts or dumps with different formats.)

The communication access method can also perform the functions of step 1 of the dump program. In this case, step 2 is not invoked automatically; it must be initiated as an independent job in order to print out the dump.

Unlike the OS/VS dump program, the DOS/VS dump program consists of only one job step and uses no intermediate disk data set. It is, however, functionally the same as the OS/VS dump program.

After a dump, the 3705 is idle and must be reloaded before it can operate again.

The Dynamic Dump Utility

The dynamic dump utility is an optional utility program that allows the contents of controller storage to be transferred from the controller to the host processor without interrupting operation of the network control program. A full storage dump or a dump of the trace tables for lines in emulation mode can be obtained. In addition, portions of storage can be displayed on the operator's console at the host processor.

The host processor module of the dynamic dump program is executable in any System/370 that will accommodate OS/VS1 or OS/VS2.

The OS/VS1 dynamic dump utility operates in a minimum virtual partition. The OS/VS2 dynamic dump utility operates in a minimum virtual region.

The ACF/TAP Service Aid

The Advanced Communications Function/Trace Analysis Program (ACF/TAP) is an IBM service aid that assists in analyzing trace data produced by ACF/VTAM, ACF/TCAM, and ACF/NCP/VS. ACF/TAP provides a common trace analysis facility for different types of trace input data. It produces output reports showing SNA and SDLC trace data formatted into merged, easy-to-use, and easy-to-read formats. Unusual conditions occurring in the trace data that may indicate error situations are highlighted.

Chapter 5. ACF/NCP/VS Release 2

Release 2 of ACF/NCP/VS is an IBM licensed program product that provides additional capabilities for SNA networks beyond those available in Release 1. Release 2 can thus assist in optimizing the management, control, and flexibility of the network. (Further capabilities that offer still more assistance in optimizing network operation are available in Release 3; see Chapter 6.)

Note: The Release 2 capabilities require that ACF/NCP/VS Release 2 be communicating with Release 2 or Release 3 of ACF/VTAM or Version 2 (Release 1, 2, or 3) of ACF/TCAM in the host processor.

The new capabilities of ACF/NCP/VS Release 2 may be summarized in two categories—communications systems management and enhanced data communications capabilities—as follows.

Communication System Management Capabilities

These are the communication system management capabilities:

- *Enhanced SDLC data link test* permits improved testing of SDLC links in two ways. When initiated by the network operator, the link test can involve only one station on the link, thus leaving the other stations available for normal operations. When initiated by the operator of a station, the link test allows that operator to determine that his station is capable of communicating with the host processor.
- *Intensive mode error recording* provides, through operator control, the capability of informing ACF/NCP/VS (Release 2) to record detailed information about temporary SDLC data link errors. Such detailed information could preclude the need for specific testing of the link to re-create the error condition.
- *Dynamic reconfiguration* allows selective alteration of the physical network definition for SNA-SDLC physical unit types 1 and 2 (that is, SNA-SDLC stations excluding host processors and communications controllers) without stopping network operation. The advantage is greater flexibility of network configurations for accommodating additional or relocated SDLC stations. This capability allows temporary configuration changes (as for testing purposes) without disrupting the network.
- *Dynamic display of network control program storage* allows the network operator, through ACF/TCAM Version 2 or ACF/VTAM Release 2, to display at a console any contiguous 256 bytes of ACF/NCP/VS storage without disrupting normal ACF/NCP/VS operation.
- *Dynamic dump of network control program storage* allows the network operator, through ACF/VTAM Release 2, to dump the entire ACF/NCP/VS storage, in 256-byte increments, without disrupting normal ACF/NCP/VS operation.
- *Support of the Network Communications Control Facility* program product and related program products. These program products are described in *Network Communications Control Facility General Information* (GC27-0429).

Enhanced Data Communication Capabilities

These are the enhanced data communication capabilities:

- *Negotiable session-initialization parameters* allow network components to have a more direct and timely control over the session-initialization process. Actual session-initialization parameters can be selected when the session is initiated; the receiver of the bind image may alter the image and return it to the issuer, which may accept or reject it.

- *Support for the Network Terminal Option program product and related capabilities* allows the inclusion in the network control program of IBM-provided code that extends the record application program interface of the communications access method to certain non-SNA stations and also allows the inclusion of user-written routines and network addressable units.
- *Support of 3705-II enhancements,* including a cycle utilization counter. The cycle utilization counter accumulates statistical data on cycle utilization within the 3705 for examination by the user. (This data includes cycles taken for instruction execution, cycle-steal operations, and maintenance cycles.)

All of these capabilities, in both categories, are of value to SNA networks having either single host-processor sites (with one or more host processors) or multiple sites (each with one or more host processors). The dynamic reconfiguration and Network Communications Control Facility capabilities increase in importance as the network expands in size; the value of the remaining capabilities is relatively independent of the size of the network.

Functional Descriptions

Described below are the individual functions that Release 2 of the ACF/NCP/VS makes available. Also discussed are the usefulness of the function and the factors to be considered when planning for migration from previous versions of the NCP to ACF/NCP/VS Release 2 or from previous versions of TCAM or VTAM to ACF/TCAM Version 2 or ACF/VTAM Release 2.

Enhanced SDLC Data Link Testing

To facilitate pinpointing the cause of errors occurring on an SDLC link, ACF/NCP/VS provides link tests. A link test is an "echo test" at the data link control level. As directed by the access method, the primary SDLC station on the link sends a test command to a secondary station on the link. The secondary station returns the command unchanged if no error occurred. In Release 1 of ACF/NCP/VS, no stations on the link being tested can carry on normal message traffic; the entire link is dedicated to the test. This disruption of message traffic means that a link testing is often deferred until minimal impact on normal operations can be expected; as, for example, at the end of the business day. In Release 2 of ACF/NCP/VS, the link test does not require disruption of traffic involving any station but the one participating in the test. Further, more than one link test can be run at the same time, each involving a different station on the link. Once each link test is completed, the result of the test is sent to the requester of the test.

Link testing is important whenever degradation or failure of a communication line is suspected. The nondisruptive nature of the enhanced capability means that a link test will more likely be run as soon as the problem is suspected rather than be deferred until normal operations are at a minimum.

This enhanced link testing capability is of particular use to SNA networks having nonswitched multipoint SDLC links.

Intensive Mode Error Recording

Intensive mode error recording is the act of recording detailed information about temporary errors on an SDLC link. The error-recovery procedures in ACF/NCP/VS can record the initial error status that started the error-recovery procedure and the final error status, in the case of a permanent (unrecoverable) error, that caused the permanent error record to be generated. In Release 1 of the ACF/NCP/VS, the information about the temporary errors that may have occurred between the initial and the final error is not sent to the host processor. The enhanced capability—intensive mode error recording—available in Release 2 can provide data about all temporary errors as they occur. When the access method has initiated intensive mode error recording, the error statistics are sent to the host processor after the retry sequence for each error is completed.

Intensive mode error recording is useful for all SNA networks having SDLC links. The expanded statistics that this function makes available will often preclude the need to run specific link tests and thus provide more efficient, timely problem determination.

Dynamic Reconfiguration

The dynamic reconfiguration capability of ACF/NCP/VS Release 2 allows the user to temporarily alter the physical network definition without stopping the network operation. That is, physical units (types 1 and 2) and logical units can be added or deleted dynamically without the need to immediately perform a new ACF/NCP/VS generation. Dynamic reconfiguration is not a substitute for the program generation procedure. It simply allows the user to make temporary changes in the network configuration at will and to defer until a more convenient time the program generation needed to incorporate the changes permanently. Until the program is regenerated, the previous program generation listings will not reflect the configuration changes made via dynamic reconfiguration.

Dynamic reconfiguration is a significant new capability in that it permits a network configuration to be changed without disruption of ongoing operations, and it can reduce the machine time that repeated regenerations would require.

SNA physical units and logical units can be dynamically added to or deleted from nonswitched SDLC links or manual-dial backup links. The dynamic reconfiguration capability does not, however, apply to host processors, network control programs, or non-SNA stations.

To prepare for dynamic reconfiguration, the system programmer specifies, in the ACF/NCP/VS generation deck, the resources (dummy PUs and LUs) to be allocated to the dynamic reconfiguration function. The actual PUs and LUs to be added or deleted are not defined in the generation deck; rather, they are defined later as needed. The ACF/NCP/VS uses the dummy resources to store the definition information for physical and logical units as they are added. In addition to specifying dummy resources, the programmer can specify which generation-defined PUs and LUs may later be deleted dynamically. For example, the programmer can specify that a master console cannot be dynamically deleted.

Actual PUs and LUs may be added to SDLC links any time the ACF/NCP/VS is active. They can be deleted only when they are in the inactive state. Both generation-defined resources and dynamically added resources may be deleted. When they are deleted, the storage allocated to them is released for use by later dynamic resource additions. Thus, many dynamic reconfigurations can be made with the same dummy PUs and LUs. Once the pool of dummy PUs and LUs is entirely occupied by actual PUs and LUs, a new ACF/NCP/VS generation is necessary. Certain physical limitations also make a new generation necessary; for example, the need to add a new SDLC link, and certain parameters specified in the generation deck that limit the number of resources that may be associated with a device.

Among the reasons the dynamic reconfiguration capability is of value to ACF/NCP/VS networks are the following:

- A central site can respond more quickly to user demands for additional or relocated stations because the network need not be shut down to make the change and because no immediate regeneration of the ACF/NCP/VS is necessary.
- A central site can correct certain errors in the network definition without the need to regenerate the ACF/NCP/VS.
- A central site can accommodate a temporary or test configuration without disrupting normal operations of the network.
- After an SDLC link failure, a new physical connection can be defined for affected LUs without the need to change the LU names and without the need for added access-method or ACF/NCP/VS storage.

In general, dynamic reconfiguration is especially important to SNA networks characterized by large numbers of SNA stations, rapid addition of new stations, frequent configuration changes requiring fast reaction time, or 24-hour-a-day operation—or any combination of these characteristics.

Dynamic Display of Network Control Program Storage

The dynamic display facility allows a network operator to display up to 256 contiguous bytes of ACF/NCP/VS storage. The operator requests the display by means of an operator control command to the access method. The command specifies the address of the ACF/NCP/VS and the starting address of the network control program storage to be displayed. No interruption of network operation is necessary as this action by the operator does not require that the network control program be stopped, the 3705 storage dumped, and then the program reloaded.

The dynamic display function is of value when operation or performance of the network is incorrect or marginal and the ACF/NCP/VS is suspected of being the cause. For example, an incorrect ACF/NCP/VS generation can cause problems of this kind. This function is especially useful for examining data areas that are static (for example, areas that are associated with an inoperative station) rather than areas that are dynamically changing (as, for example, buffer pools or areas that are associated with an operative station). Use of this function is appropriate for users who have considerable knowledge of internal ACF/NCP/VS operations and need to examine ACF/NCP/VS storage when performing problem determination.

The network operator, through ACF/VTAM, can use the dynamic display facility to dump the entire ACF/NCP/VS storage in 256-byte increments without disrupting ACF/NCP/VS operation.

Network Communications Control Facility

The Network Communications Control Facility is a program product designed for use with ACF/TCAM and ACF/VTAM in medium-to-complex SNA networks. This facility provides a base upon which the user may employ IBM-supplied and/or user-written application programs to provide functions collectively called communication systems management. These programs, such as the IBM-supplied Network Problem Determination Application, are referred to as communication systems management processors. The purpose of the Network Communications

Control Facility is to provide these processors with (1) an environment that insulates them from operating system and access method dependencies, and with (2) application functions common to the various command processors. The facility provides those unique tasks that must be performed in order to monitor, control, change, and improve the overall operation of an SNA network.

ACF/NCP/VS Release 2 provides enhanced error recording functions for remote stations controlled by the network control program in support of the Network Communications Control Facility.

Further information about the Network Communications Control Facility program product may be found in *Network Communications Control Facility General Information* (GC27-0429).

Negotiable Session-Initialization Parameters

To establish a session with a secondary logical unit (LU), a primary logical unit sends a Bind command to the secondary LU. The primary LU includes with the Bind command certain session-initialization parameters that represent the way in which the primary LU wishes to conduct the session. An example of a session-initialization parameter is the maximum request/response unit (RU) size that will be accepted from the secondary LU.

In Release 1 of ACF/NCP/VS, the secondary LU may either accept the initialization parameters, in which case the session is established, or it may reject them, in which case the session is not established. If the secondary LU rejects the session, the primary LU must try again to establish a session by sending a Bind command with a different set of parameters. In order to be sure of establishing a session with any given logical unit, the primary LU must maintain a list of session-initialization parameters that are acceptable to that secondary LU. The term *non-negotiable* is applied to the situation in which the secondary LU either accepts or rejects outright the parameters accompanying the Bind command.

In Release 2 of ACF/NCP/VS, some session-initialization parameters can be negotiable. That is, upon receiving the suggested parameters from the primary LU, when the session begins, the secondary LU can substitute session-initialization parameters that the secondary LU wishes to have govern the session. The secondary LU includes its selected parameters in its response to the primary LU. If the primary LU accepts the substituted parameters, the session proceeds using those parameters. If the primary LU does not accept the substituted parameters, it must send an Unbind command to end the session.

The ability for a secondary LU to suggest its own choice of session-initialization parameters means that primary LUs are relieved of the need to maintain a list of parameters suitable to all of the secondary LUs with which they may establish sessions. Besides reducing the storage needed to store these parameter lists, this capability allows more flexibility in modifying the session-initialization parameters to suit the needs of the session. That is, the access method need not be regenerated and reinitialized simply to include new or changed parameter lists each time a secondary LU in the network requires new session-initialization parameters.

Support for the Network Terminal Option Program Product and Related Capabilities

ACF/NCP/VS Release 2 and Release 3 allow the inclusion of IBM-provided code that extends the record application program interface of the communications access method to IBM 2740 Model 1, IBM 2741, Western Union TWX Models

33/35, and World Trade teletypewriter terminals and that permits these terminals to participate in a multisystem as well as a single-system environment.

The support provided by ACF/NCP/VS Release 2 and Release 3 allows the inclusion of user routines to be executed during ACF/NCP/VS initialization, and the inclusion of user-written network addressable units that execute as part of ACF/NCP/VS. These routines may be used to add programmed SNA resources, to support non-SNA devices through user-written code, or to add line control support not provided by ACF/NCP/VS but compatible with type 2 or type 3 communication scanners in the IBM 3705.

The Network Terminal Option is a program product available for use with ACF/VTAM Release 2 and Release 3. A description of this program product may be found in *Network Terminal Option General Information* (GC38-0297).

3705-II Hardware Additions

ACF/NCP/VS Release 2 and Release 3 support the following hardware additions to the IBM 3705-II Communications Controller:

- 900 nanosecond storage cycle speed
- Main storage expansion to 512K bytes
- CCITT V.35 local attachment interface for 14.4/57.6 K bps links between network control programs.
- Cycle utilization counter

The cycle utilization counter accumulates statistical data on 3705 cycle utilization for access by the user. The data includes cycles taken for instruction execution, cycle steal operations, and maintenance. From this data ACF/NCP/VS Release 2 or Release 3 can provide information on the percentage of available cycles utilized. In a similar manner, ACF/NCP/VS Release 2 or Release 3 can provide information on the percentage of buffers available.

The 3705-II Models J, K, and L include this additional hardware, details of which may be found in *Introduction to the IBM 3704 and 3705 Communications Controllers,* GA27-3051.

In addition to the foregoing items, Release 2 and Release 3 of ACF/NCP/VS support other 3705-II hardware additions that require no new ACF/NCP/VS programming changes.

Chapter 6. ACF/NCP/VS Release 3

Release 3 of ACF/NCP/VS is an IBM licensed program product that provides major new capabilities for SNA-based networks beyond those available in Releases 1 and 2. Release 3 of ACF/NCP/VS, together with Release 3 of ACF/VTAM and Version 2 Release 2 of ACF/TCAM, contributes to the usability, reliability, and availability of SNA-based networks.

The new capabilities of ACF/NCP/VS Release 3 may be summarized in two categories: enhanced configuration flexibility and expanded communication system management capabilities.

Enhanced Configuration Flexibility

These are the capabilities that permit the user to arrange the components of an SNA-based network in a more flexible way.

Extended NCP interconnection allows an ACF/NCP/VS Release 3 to perform identical network control functions regardless of whether it was loaded into the communications controller over a channel attachment to a host processor or over an SDLC link from another communications controller. For single-domain networks, this capability means that more than two communications controllers can be connected in tandem, as Figure 6-1 shows. All controllers connected in tandem must execute Release 3 of ACF/NCP/VS.

For multiple-domain networks, extended interconnection means that (1) cross-domain SDLC links can be established between network control programs in remote controllers, as shown in Figure 6-2, and (2) a single ACF/NCP/VS Release 3 can communicate via cross-domain links with network control programs in multiple 3705s, as shown in Figure 6-3. Either kind of attachment permits ownership of a network control program in a remote controller to be shared by two or more access methods.

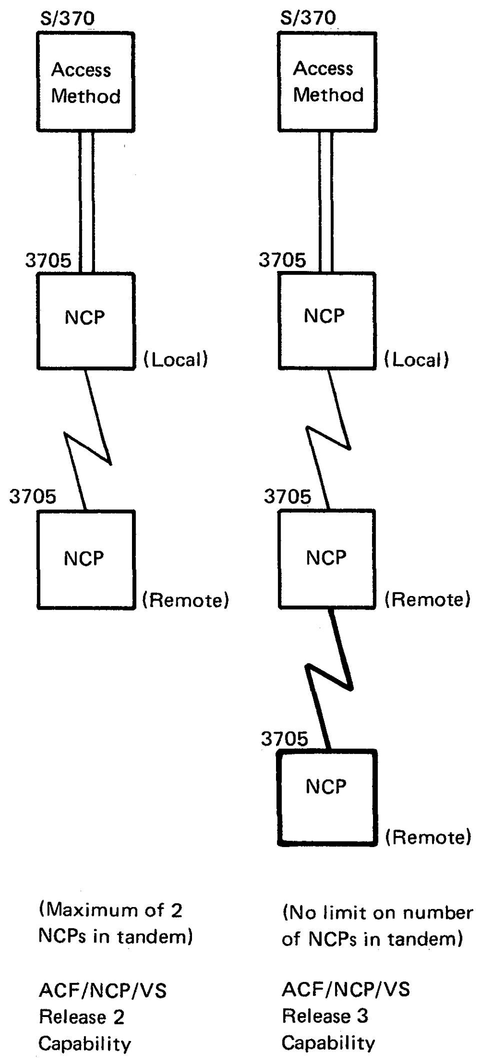

Figure 6-1. Multiple NCPs in Tandem

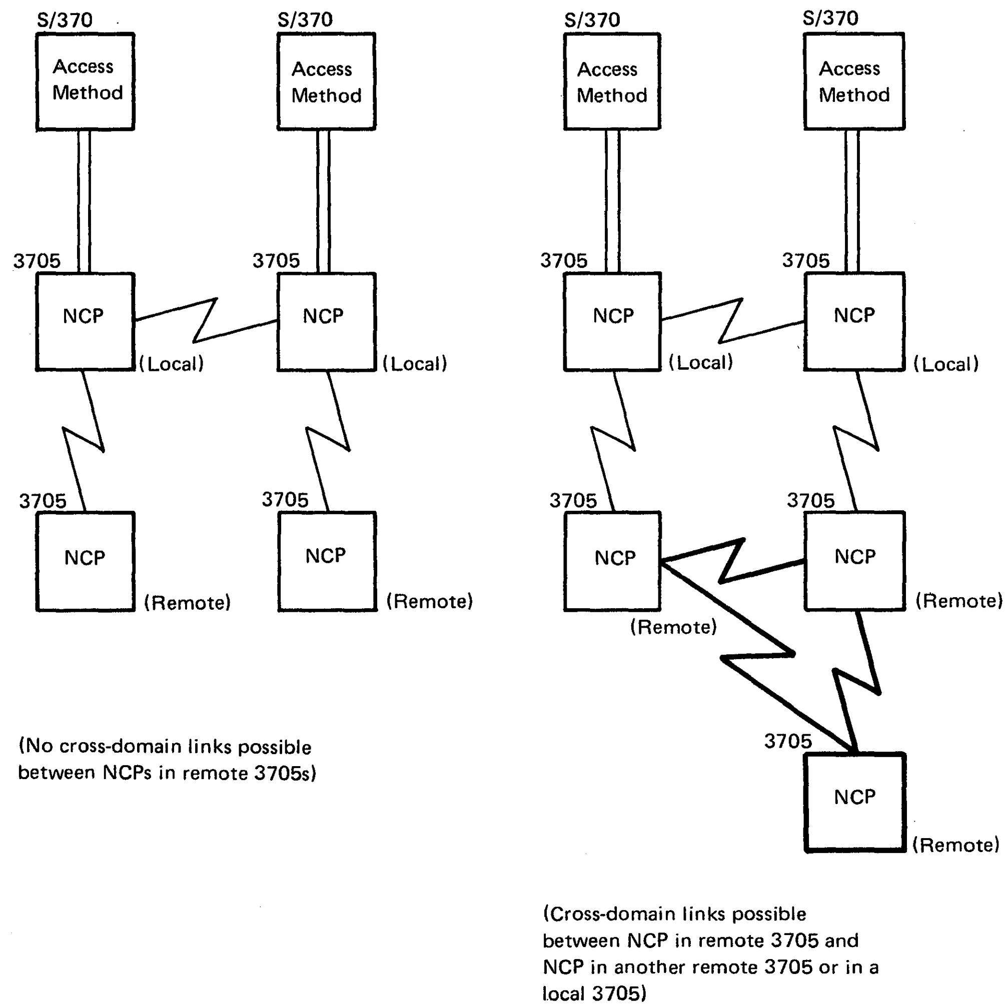

Figure 6-2. Cross-Domain Links to NCPs in Remote 3705s

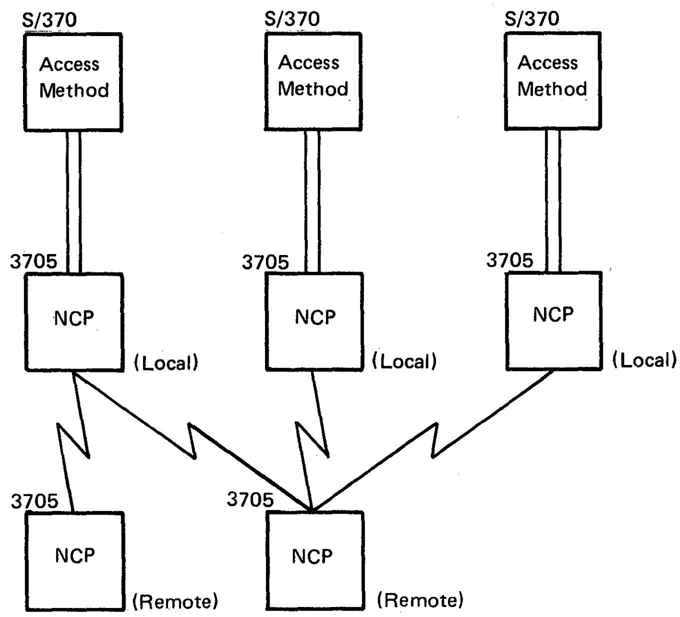

Figure 6-3. NCP in Remote 3705 Communicating with Multiple Access Methods via Cross-Domain Links to NCPs in Local 3705s

For either single- or multiple-domain networks, extended interconnection can have two major benefits. First, because the functions of a network control program are the same regardless of whether it was loaded over a channel or a link, controllers may be dispersed geographically without concern with whether the controllers are local or remote. Second, ownership of any network resource may be shared concurrently by multiple access methods, as Figure 6-4 shows. Increased shareability of resources, together with the availability of alternate connections between subareas (described later in this chapter under "Multiple Routing"), can simplify the ability of an access method to take over control of a resource whose owner (or a link to that owner) has failed or been otherwise deactivated. The takeover can occur without disruption of any cross-domain sessions in which the resource is engaged. The availability of the resource is thus improved.

Figure 6-5 shows an example of an SNA-based network in which there are three local-local and five local-remote links, as supported by ACF/NCP/VS Release 2. By replacing ACF/NCP/VS Release 2 with ACF/NCP/VS Release 3 in each controller, the network configuration can be extended and data paths within the network can be made more flexible, as shown in Figure 6-6. One additional remote controller (connected in tandem to an existing remote controller), three remote-remote links, and three new local-remote links have been added. (One of the added links, in combination with link 8 in the figure, replaces a former local-remote link, link 6). The advantage of having alternate paths (routes) in a network is discussed further under "Multiple Routing," later in this chapter.

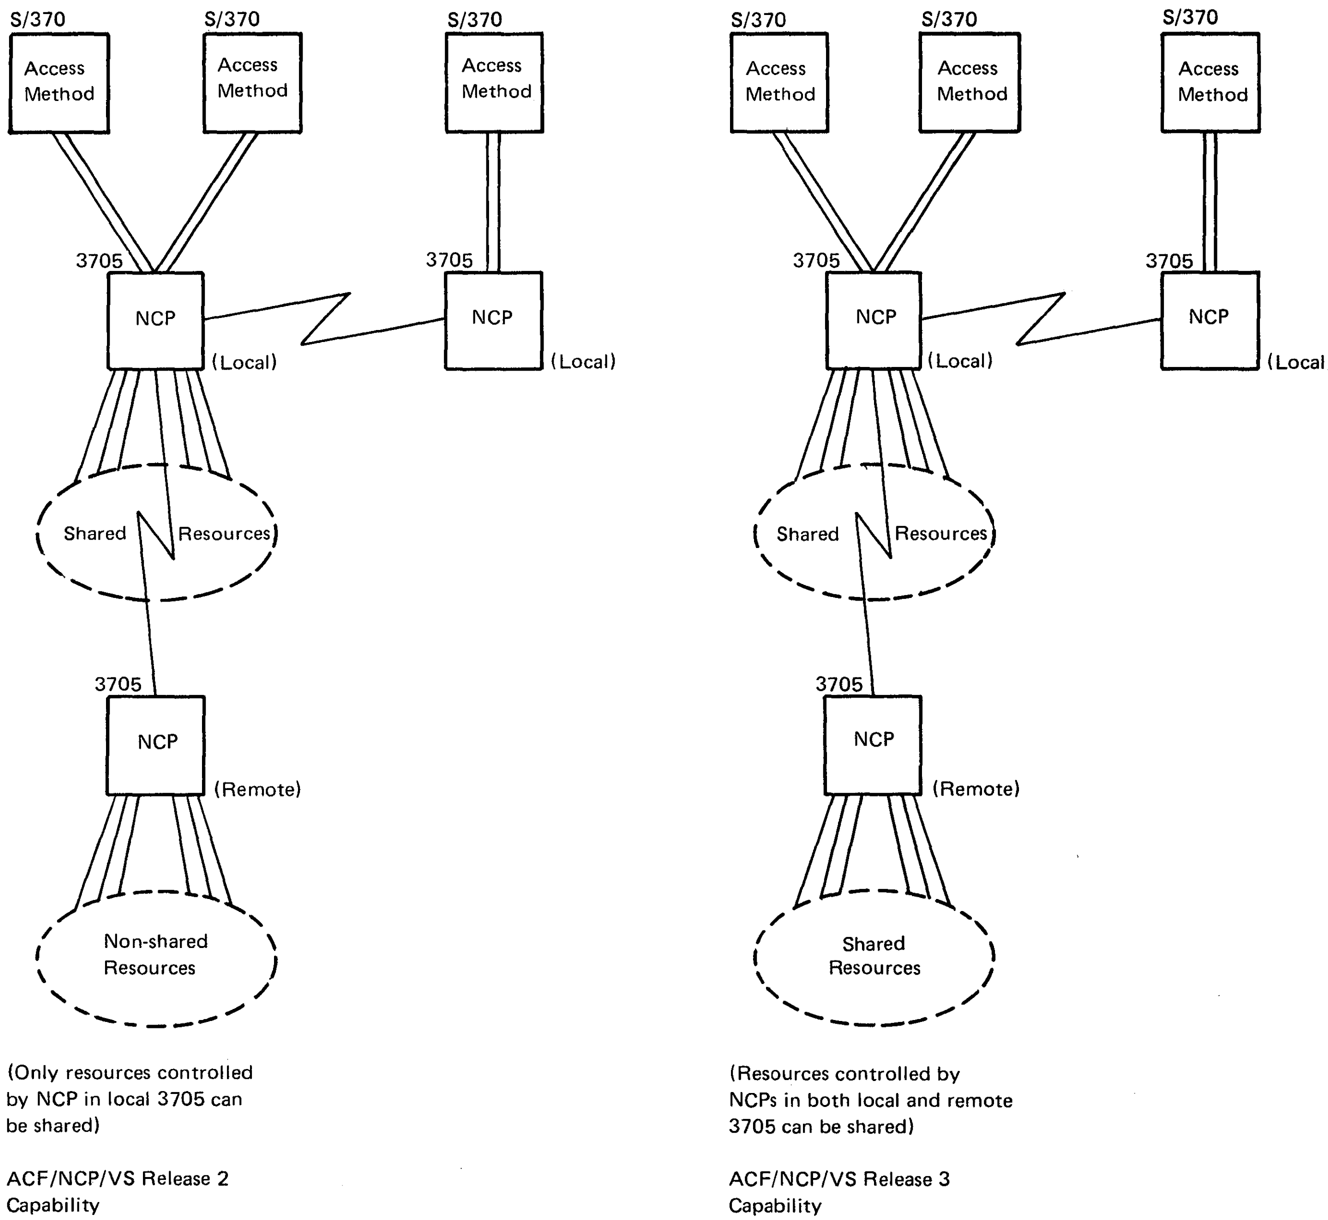

Figure 6-4. Sharing of Resources Controlled by NCP in Remote 3705

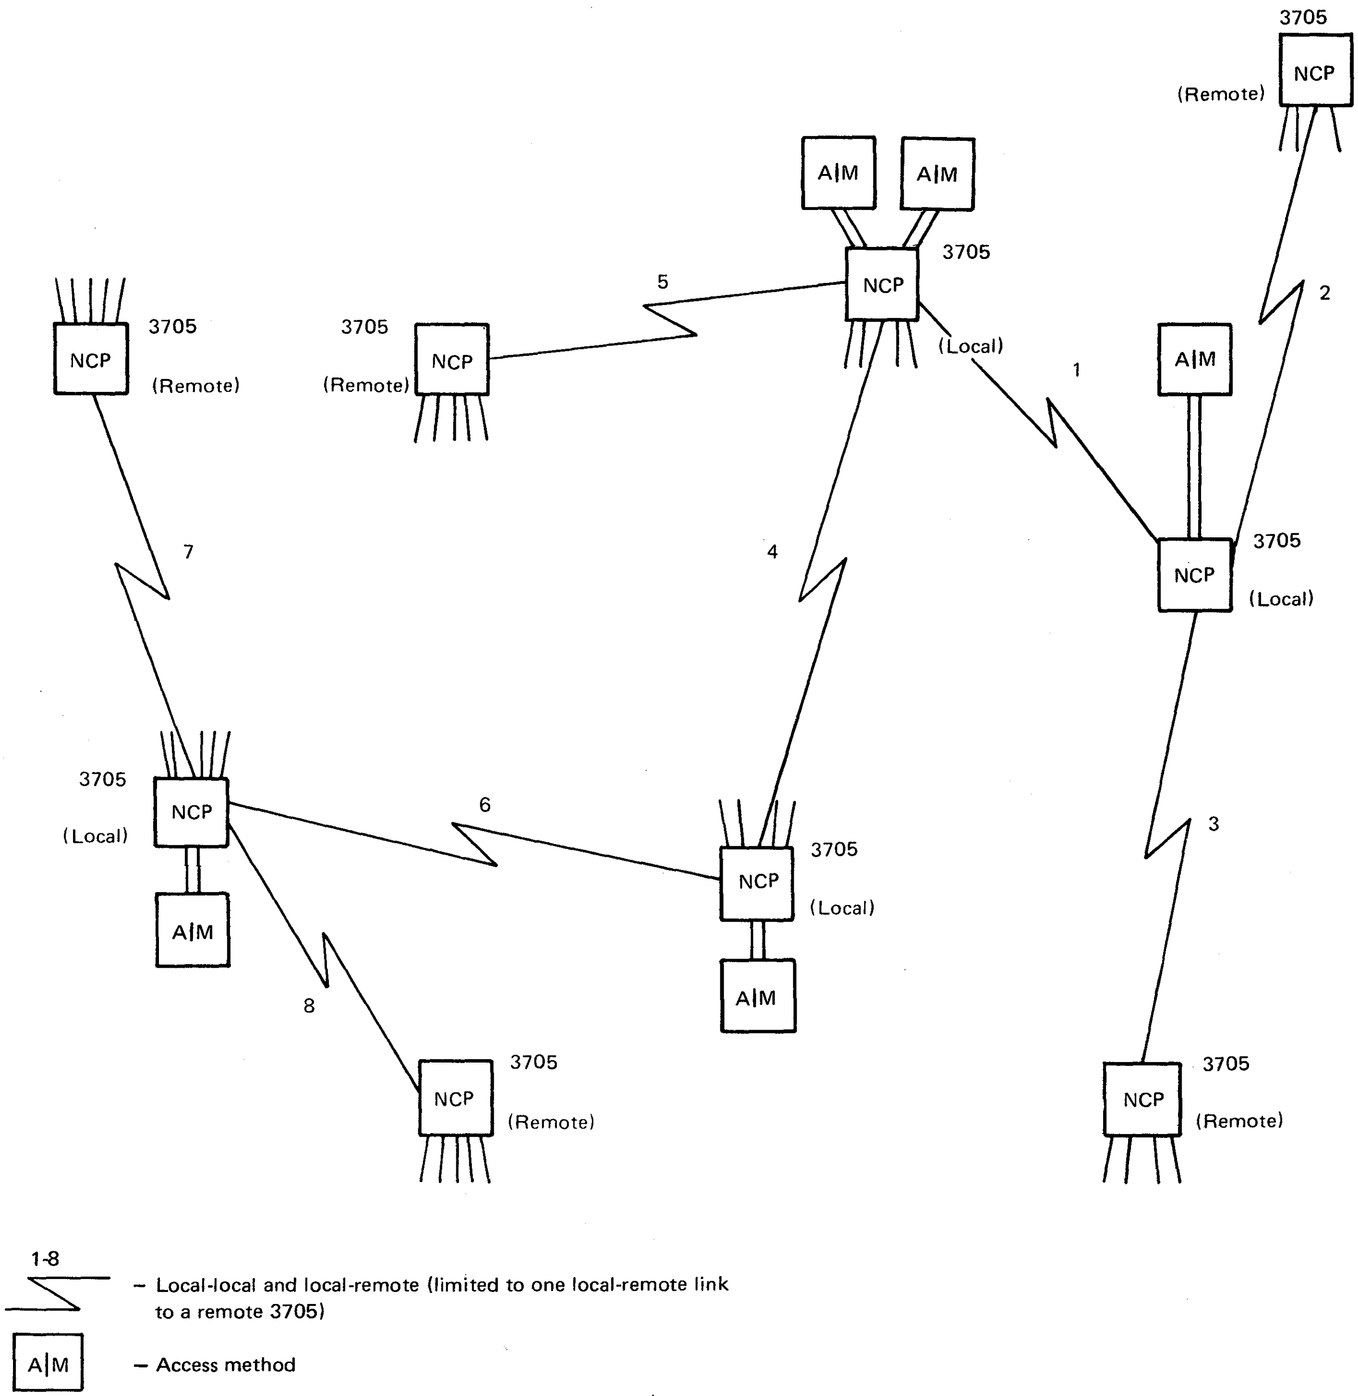

Figure 6-5. Example of a Network Controlled by ACF/NCP/VS Release 2

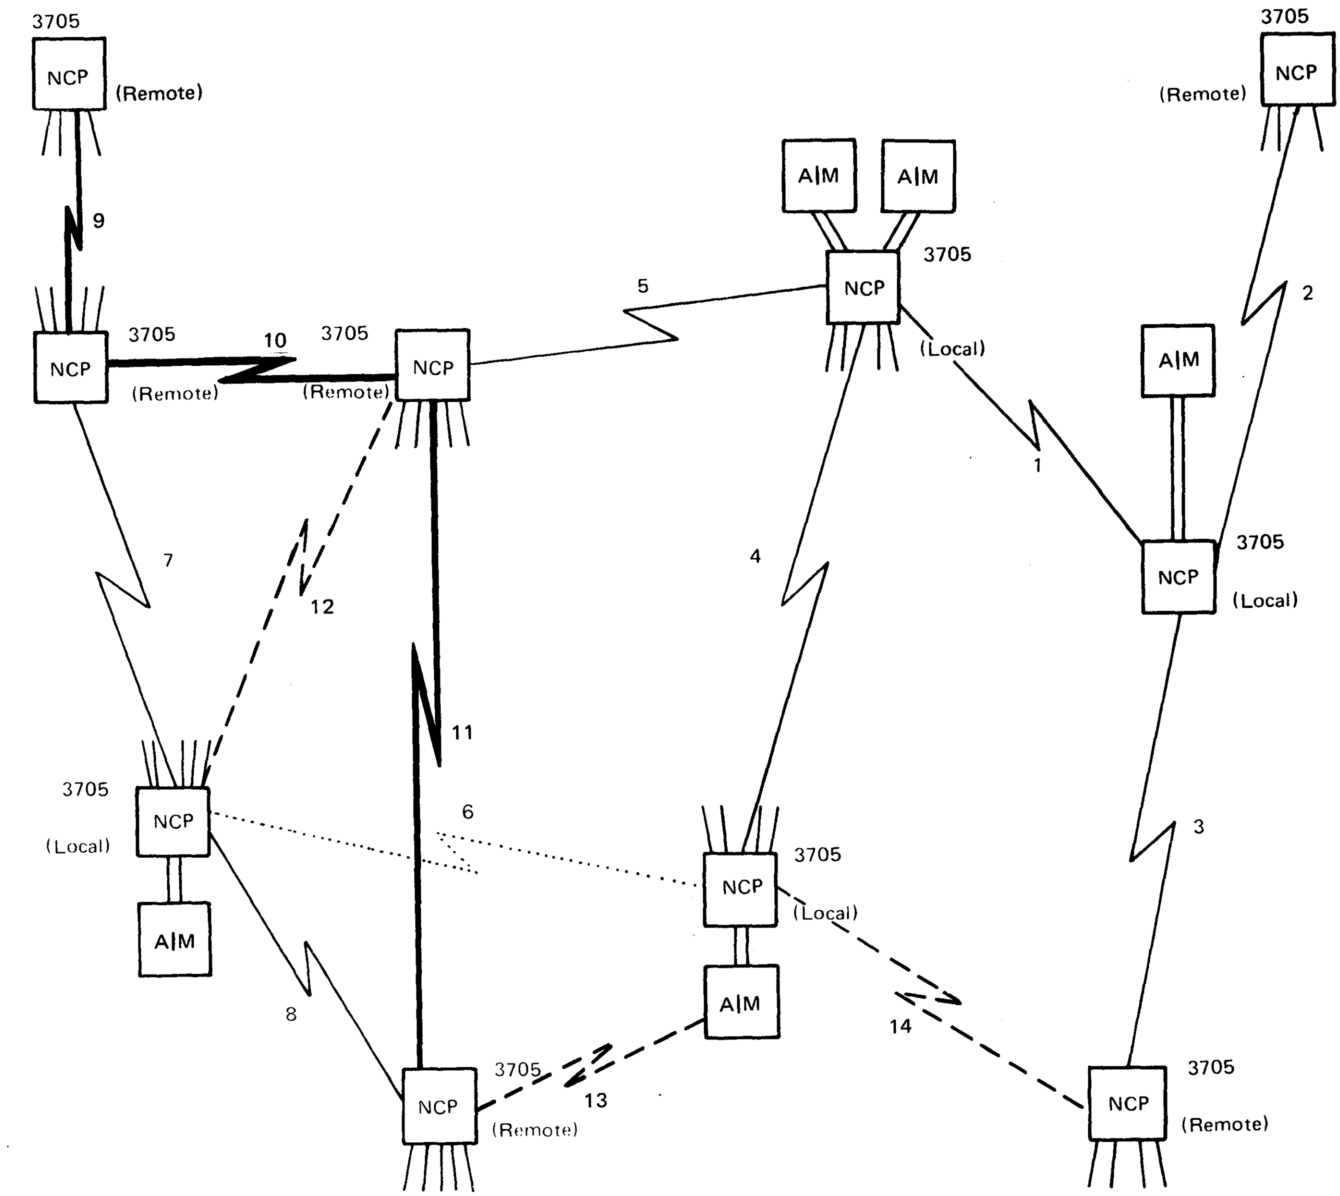

1-8 - Local-local and local-remote links in network of Figure 6-5

9-11 - Remote-remote links that can be added to network because network control programs are ACF/NCP/VS Release 3

12-14 - Further local-remote links that can be added to network because network control programs are ACF/NCP/VS Release 3

6 - Former link replaced by combination of existing link 8 and new link 13.

A/M - Access method

Figure 6-6. Example of a Network Controlled by ACF/NCP/VS Release 3

Support of parallel SDLC links between adjacent NCPs allows data traffic between any two adjacent network control programs (provided that both are ACF/NCP/VS Release 3) to flow simultaneously over two or more SDLC links. The number of parallel links that can be operated simultaneously depends only on the line attachment capabilities of the communications controllers in which the network control programs are executed. By distributing the total message flow among multiple links, the user can minimize the disruption to sessions caused by failure of a single link or by the need to deactivate a link for reasons such as maintenance. Thus, the reliability and availability of the paths between network control programs are improved.

The *extended NCP ownership* capability of ACF/NCP/VS Release 3 allows ownership of any link between network control programs to be exercised by any or all SSCPs in the network, regardless of how many domains distant the link may be from an SSCP. Thus several SSCPs can activate a link between network control programs. Information maintained in the programs reflect the current number of SSCPs that own (have activated) the link. Each SSCP may "deactivate" the link when it has no further use for it, without actually causing the link to be deactivated. Only when a Deactivate Link command is received from the last remaining owner does the network control program actually deactivate the link.

Expanded Communication System Management Capabilities

Release 3 of ACF/NCP/VS provides the following capabilities that permit an SNA-based network to be more readily managed in ways that are less dependent on the physical configuration of the network than is the case with Release 2.

Multiple routing allows several routes to be carrying data concurrently between a pair of subareas in the network. In networks using prior releases of ACF/NCP/VS, only one route (path) can carry data between a pair of subareas; all data traffic sent during sessions established between the subareas must be transmitted over that path. If a physical network element in that path fails, all sessions are interrupted and all data transmission stops until the element is repaired or another element, such as a backup link, is substituted. Multiple routing allows up to eight physical routes to be established between a pair of subareas; ACF/NCP/VS Release 3 distributes the sessions between the subareas among the active routes. Any given session and its data flow over one route. That is, the session data is not distributed among routes. Several sessions, however, may be assigned to the same route.

A major advantage of multiple routing is that failure of one route of several between two subareas does not halt all traffic between the subareas. Only the sessions using the inoperative route are disrupted, and new sessions to replace the disrupted ones can be established over the remaining routes. Thus, reestablishment of communication between the subareas does not depend on repair of the inoperative route. The multiple routing capability can therefore improve the reliability and availability of SNA-based networks in which this capability is used.

Multiple priority level support allows the user to establish up to three levels of priority for transmission of data within a session. A specific route can be assigned to each priority level so that high-priority traffic can be sent on one route while low-priority traffic flows over a different route. For example, highly time-dependent interactive transactions can be sent with higher priority than are batches of collected data that only requires processing at the end of the day.

The *transmission group capability* of ACF/NCP/VS Release 3 allows multiple SDLC links between adjacent network control programs to be grouped logically to appear to the network as a single connection. ACF/NCP/VS Release 3 dynamically distributes the data traffic among the links in the transmission group. Failure or deactivation of any of the active links in the group (so long as it is not the sole remaining link) results in automatic redistribution of data traffic among the remaining links. Sessions currently in progress over those links are not disrupted. Application programs in the host processor and user sessions are entirely unaware of and are not concerned with any redistribution of traffic within a transmission group resulting from a link becoming inoperative. Figure 6-7 shows five links between adjacent network control programs divided into two transmission groups.

Two network control programs can be joined up to eight transmission groups; any single active link can be assigned to only one transmission group at a time.

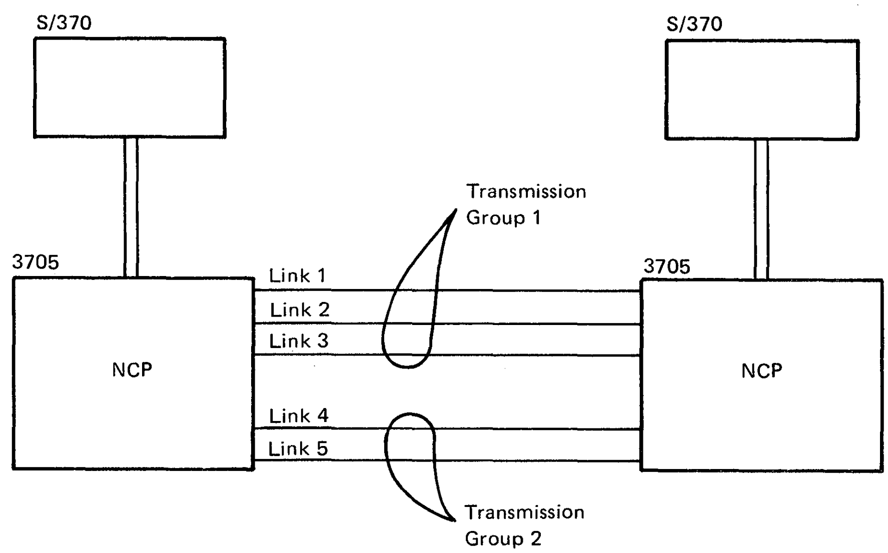

Figure 6-7. Two Transmission Groups between NCPs

Route verification permits a network operator to (1) verify that a given route is usable for data transmission before assigning it to message traffic, (2) take corrective action for a route that has become inoperative, and (3) verify that a route that has been inoperative is once again usable for data transmission after being reactivated. The route verification function identifies the physical element in the route that failed; through the access methods, it notifies the network operator who requested the verification and the network operators for all host subareas that share ownership of the subarea in which the inoperative element is located.

Flow-control enhancements allow the network control program and the access methods to regulate the flow of data in the network. Through continuous monitoring of data flow and application of various control mechanisms, the entry of data into the network from logical units is limited to an amount that can be routed through the network and delivered to its destination without delay and without overloading of network resources such as network control program buffers.

Enhanced non-disruptive session recovery lessens the disruptive effect on LU-LU sessions of reestablishing related SSCP-PU or SSCP-LU sessions. In prior releases of ACF/NCP/VS, any reestablishment of an SSCP-PU or SSCP-LU session following failure of a network element in the route used by such sessions disrupts any active LU-LU sessions involving corresponding LUs. With Release 3 of ACF/NCP/VS, the LU-LU sessions are not disrupted, provided that the associated physical units are capable of being reactivated without themselves resetting their logical units. Non-disruptive restart is available only for SNA stations on nonswitched links; this function does not apply to switched SNA stations or to any non-SNA stations.

Chapter 7. Planning for Use of ACF/NCP/VS

This chapter describes some factors that must be considered when an installation plans to replace earlier versions of the network control program with ACF/NCP/VS Release 2 or Release 3.

Communications Controller Hardware Requirements

ACF/NCP/VS Release 2 and Release 3 can be executed in the following communications controllers, whether channel-attached to a host processor or remotely connected to a host processor by an SDLC link to a channel-attached (local) communications controller:

ACF/NCP/VS Release 2:

- IBM 3705-I Communications Controller (with a minimum of 80K bytes of storage)
- IBM 3705-II Communications Controller (with a minimum of 96K bytes of storage)

ACF/NCP/VS Release 3:

- IBM 3705-I Communications Controller (with a minimum of 144K bytes of storage)
- IBM 3705-II Communications Controller (with a minimum of 128K bytes of storage)

Prerequisite NCP/VS Requirement

Use of ACF/NCP/VS Release 2 or Release 3 requires the concurrent installation of prerequisite NCP/VS system control programming (SCP).

Levels of ACF/NCP/VS Required in a Network Having Multiple Network Control Programs

Every network control program that is to communicate with another network control program over a cross-domain or a single-domain link within the same network must be an ACF/NCP/VS. That is, all network control programs that communicate with each other over cross-domain and/or same-domain links must be ACF/NCP/VS Release 1, Release 2, or Release 3.

Compatibility of ACF/NCP/VS Release 2 with ACF/TCAM, ACF/VTAM, TCAM, and VTAM

ACF/NCP/VS Release 2 is compatible with ACF/TCAM Version 2 and with ACF/VTAM Release 2. ACF/NCP/VS Release 2 is also compatible with ACF/TCAM Version 1, ACF/VTAM Release 1, TCAM Level 10, and VTAM Level 2, provided that these access methods have the appropriate PTFs, if required. However, many of the enhanced capabilities offered by ACF/NCP/VS Release 2 cannot be used with Version 1 of ACF/TCAM or Release 1 of ACF/VTAM, or with TCAM or VTAM. See "Migration Paths to ACF/NCP/VS Release 3" for further information on compatibility.

Compatibility of ACF/NCP/VS Release 3 with ACF/TCAM and ACF/VTAM

ACF/NCP/VS Release 3 is compatible with ACF/TCAM Version 2 Release 3 and with ACF/VTAM Release 3. ACF/NCP/VS Release 3 will also operate with ACF/TCAM Version 1, ACF/TCAM Version 2 Releases 1 and 2, ACF/VTAM Release 1, and ACF/VTAM Release 2, provided that these access methods have the appropriate PTFs, if required. However, support is at the functional level of the access method; the full functional capability of

ACF/NCP/VS Release 3 depends upon coexistence in the network with ACF/TCAM Version 2 Release 3 and/or ACF/VTAM Release 3.

Customer Responsibilities

To install and use ACF/NCP/VS Release 2 or Release 3, the customer must:

- Design the network
- Meet the minimum 3705 hardware requirements, such as minimum storage size, to support ACF/NCP/VS Release 2 or Release 3 in the planned environment
- Order and install all required communication equipment
- Order and install ACF/NCP/VS Release 2 or Release 3 (Program Number 5735-XX1) and System Support Programs for ACF/NCP/VS (Program Number 5735-XX3) and prerequisite system control programming
- Generate the ACF/NCP/VS Release 2 or Release 3 program, or programs, appropriate to the network configuration
- Load the generated ACF/NCP/VS load module into the 3705 communications controller, using the loader program furnished by the ACF/NCP/VS system support programs or the loader facility of the access method with which the ACF/NCP/VS will communicate
- Meet the requirements of the host communications access method that is to communicate with the generated and loaded ACF/NCP/VS

Migration Paths to ACF/NCP/VS Release 3

ACF/NCP/VS Release 2 and Release 3 can be used with levels of ACF/TCAM and ACF/VTAM other than the most current level. Similarly, Version 2 of ACF/TCAM and Releases 2 and 3 of ACF/VTAM can be used with versions (releases) of the NCP/VS and ACF/NCP/VS other than the most current version (release). This capability of these programs permits a network to be upgraded a step at a time, thus reducing the complexity of the upgrading task.

Permissible combinations of network control programs and access methods are as follows. (The figures referred to in this section show graphically the levels of access methods with which each release of NCP/VS and ACF/NCP/VS will operate.)

NCP/VS Version 5 (see Figure 7-1) will operate with the following access methods (provided that the access methods have the appropriate PTFs, if required):

- TCAM Level 10
- ACF/TCAM Version 1
- ACF/TCAM Version 2 Release 1
- ACF/TCAM Version 2 Release 2
- ACF/TCAM Version 2 Release 3
- VTAM Level 2
- ACF/VTAM Release 1

ACF/NCP/VS Release 1 (see Figure 7-2) will operate with the following access methods (provided that the access methods have the appropriate PTFs, if required):

- TCAM Level 10
- ACF/TCAM Version 1
- ACF/TCAM Version 2 Release 1
- ACF/TCAM Version 2 Release 2
- ACF/TCAM Version 2 Release 3

- VTAM Level 2
- ACF/VTAM Release 1
- ACF/VTAM Release 2
- ACF/VTAM Release 3

ACF/NCP/VS Release 2 (see Figure 7-3) will operate with the following access methods (provided that the access methods have the appropriate PTFs, if required):

- TCAM Level 10
- ACF/TCAM Version 1
- ACF/TCAM Version 2 Release 1
- ACF/TCAM Version 2 Release 2
- ACF/TCAM Version 2 Release 3
- VTAM Level 2 (OS/VS1 and OS/VS2 MVS only)
- ACF/VTAM Release 1
- ACF/VTAM Release 2
- ACF/VTAM Release 3

ACF/NCP/VS Release 3 (see Figure 7-4) will operate with the following access methods (provided that the access methods have the appropriate PTFs, if required):

- ACF/TCAM Version 1
- ACF/TCAM Version 2 Release 1
- ACF/TCAM Version 2 Release 2
- ACF/TCAM Version 2 Release 3
- ACF/VTAM Release 1
- ACF/VTAM Release 2
- ACF/VTAM Release 3

Functions supported by NCP/VS Version 5 or ACF/NCP/VS Release 1, with the appropriate level of TCAM, ACF/TCAM, VTAM, or ACF/VTAM are also supported by ACF/NCP/VS Release 2 with the appropriate level of access method.

Functions supported by ACF/NCP/VS Release 1 or Release 2 with the appropriate level of ACF/TCAM or ACF/VTAM are also supported by ACF/NCP/VS Release 3 with the appropriate level of access method. *Exception:* ACF/NCP/VS Release 3 does not support the SDLC/BSC path function.

ACF/NCP/VS Program Generation Considerations

ACF/NCP/VS Release 2: Stage 1 of the program generation procedure of ACF/SSP allows the user to generate an ACF/NCP/VS Release 2 using the ACF/NCP/VS Release 1 source statements. The generated program will function as an ACF/NCP/VS Release 1 but will include some control blocks and tables that are new with Release 2. The generated program will be able to communicate with ACF/TCAM (Version 1 or 2), ACF/VTAM (Release 1 or 2), and other ACF/NCP/VS (Release 1 or 2) programs. Some of the enhanced capabilities of ACF/NCP/VS Release 2 will, however, be realized only when the generation source statements are augmented with the parameters required to specify those capabilities and only when the program is communicating with Version 2 of ACF/TCAM or Release 2 of ACF/VTAM.

Once an ACF/NCP/VS has been upgraded, by regeneration, from Release 1 to Release 2, it will not be necessary to regenerate it when some other network control program in the network is similarly upgraded.

ACF/NCP/VS Release 3: Stage 1 of the program generation procedure of ACF/SSP allows the user to generate an ACF/NCP/VS Release 3 using the ACF/NCP/VS Release 2 (or Release 1) source statements. The generated program will function as an ACF/NCP/VS Release 2 (or Release 1) but will include some control blocks and tables that are new with Release 3. The generated program will be able to communicate with ACF/TCAM (Version 1 or Version 2 [Release 1, 2, or 3]), ACF/VTAM (Release 1, Release 2, or Release 3), and other ACF/NCP/VS (Release 1, Release 2, or Release 3) programs. Some of the enhanced capabilities of ACF/NCP/VS Release 3 will, however, be realized only when the generation source statements are augmented with the parameters required to specify those capabilities and only when the program is communicating with Version 2 Release 3 of ACF/TCAM or with Release 3 of ACF/VTAM.

Once an ACF/NCP/VS has been upgraded, by regeneration, from Release 1 or Release 2 to Release 3, it will not be necessary to regenerate it when some other network control program in the network is similarly upgraded.

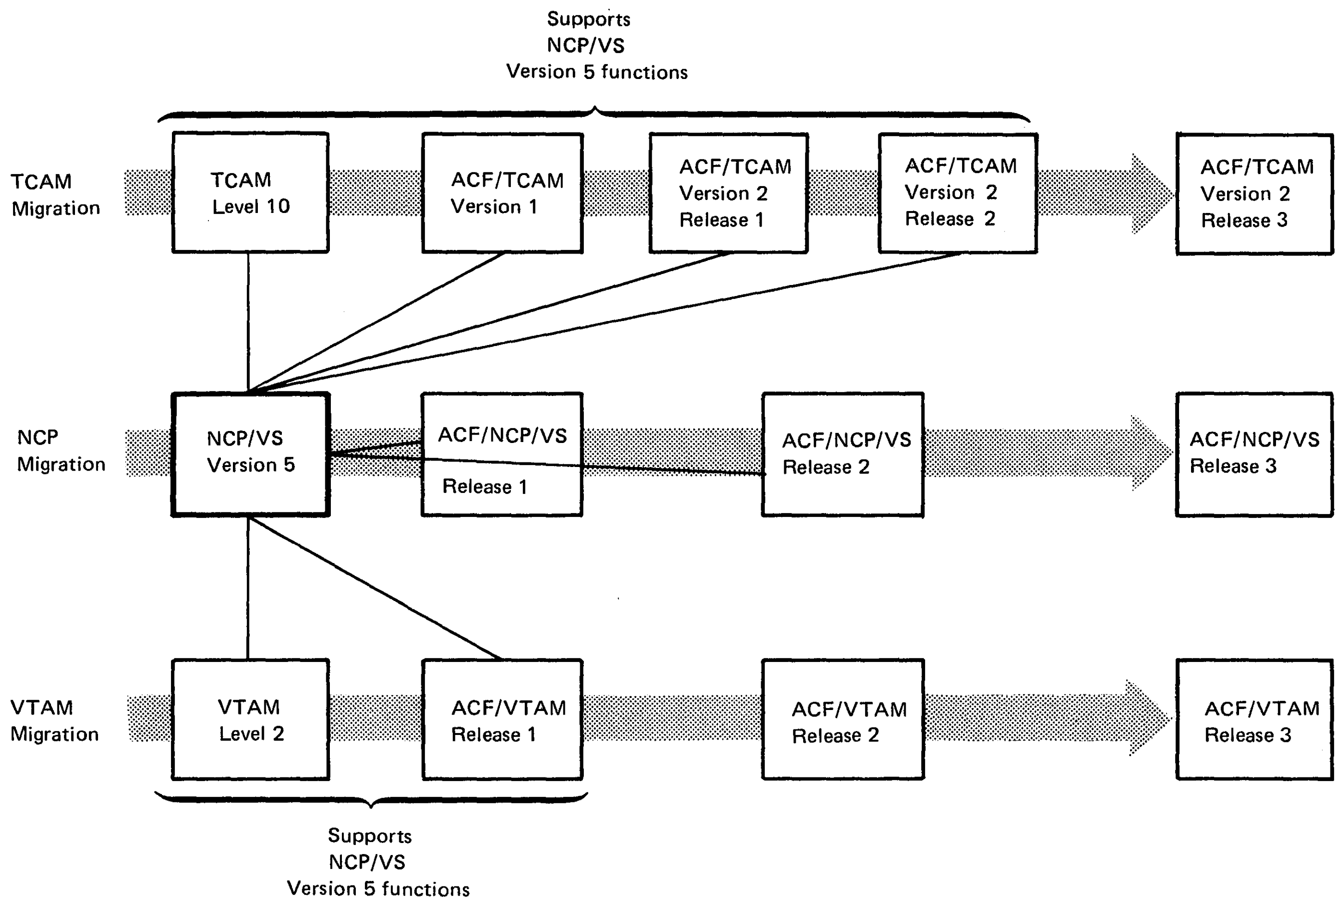

Figure 7-1. Levels of Access Method and NCP with which NCP/VS Version 5 Can Communicate

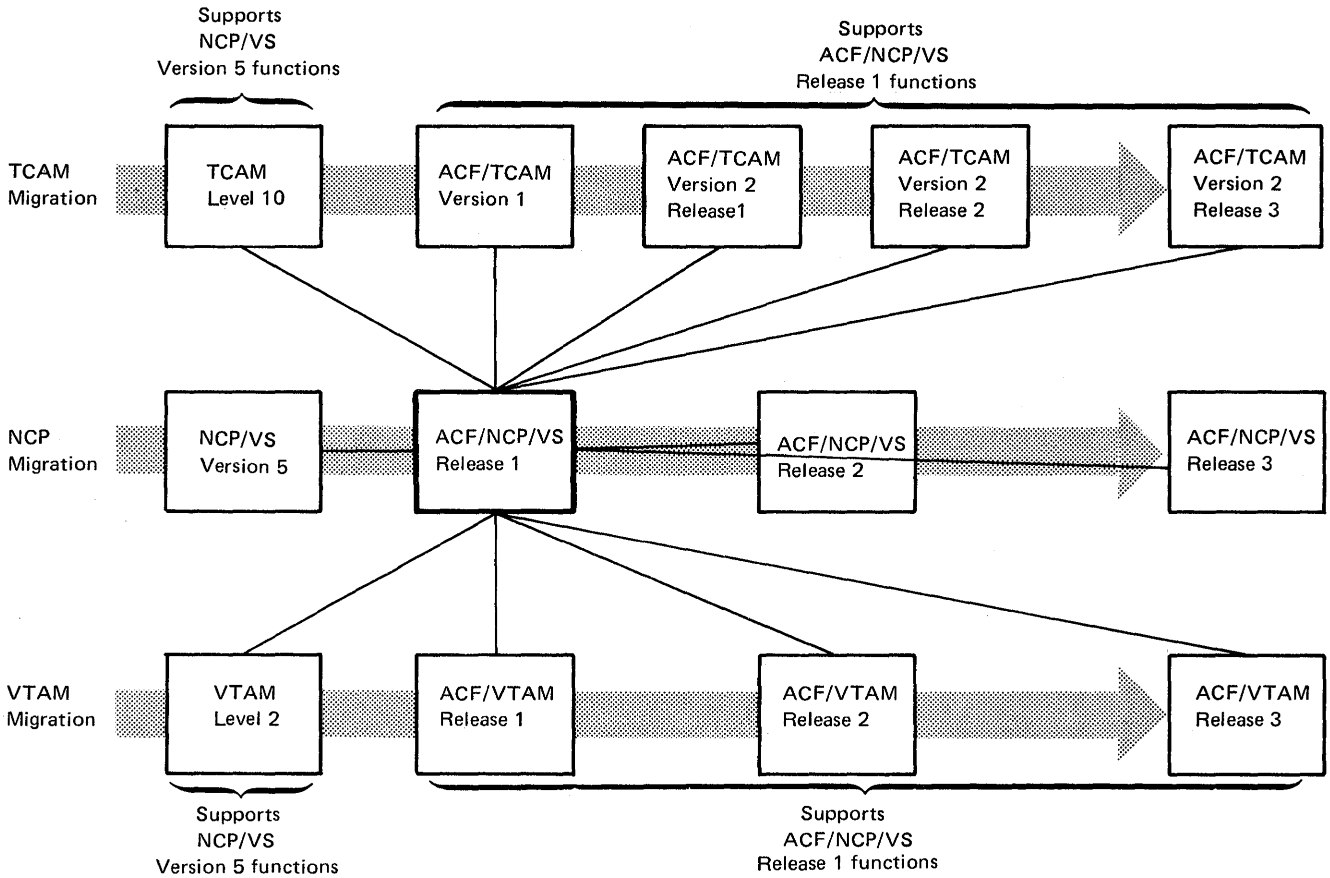

Figure 7-2. Levels of Access Method and NCP with which ACF/NCP/VS Release 1 Can Communicate

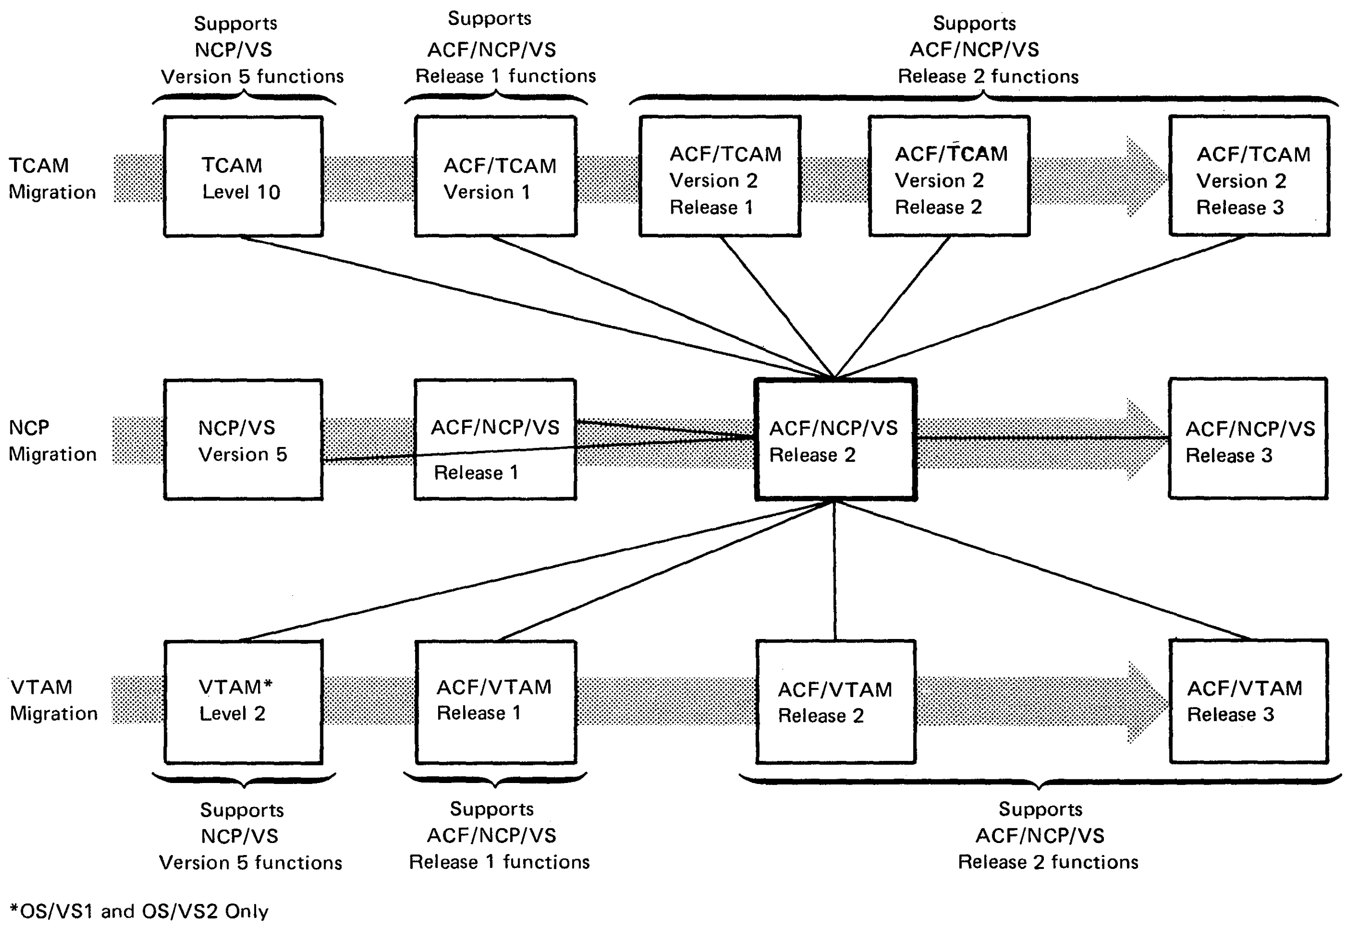

Figure 7-3. Levels of Access Method and NCP with which ACF/NCP/VS Release 2 Can Communicate

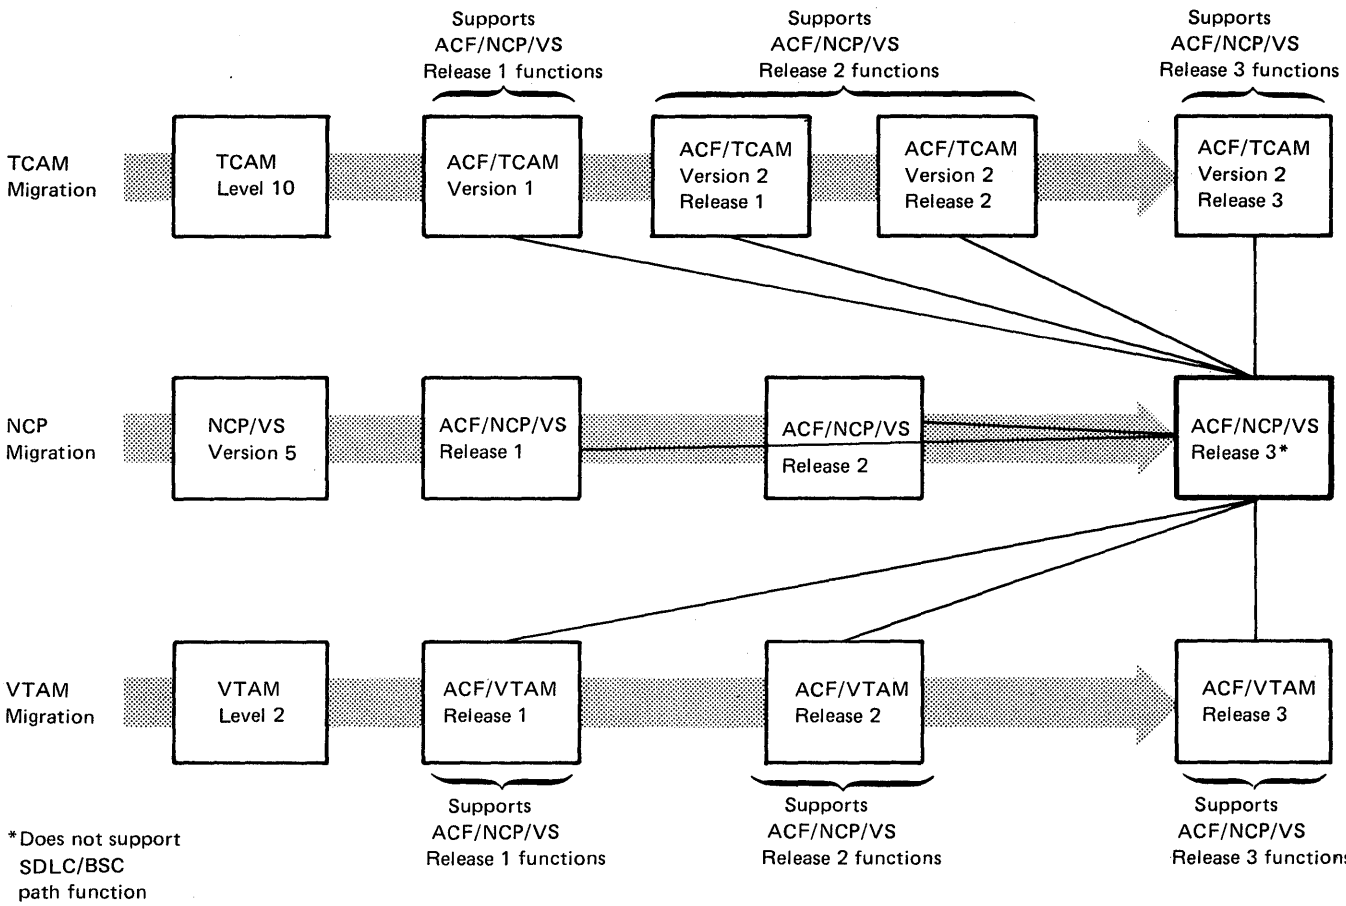

Figure 7-4. Levels of Access Method and NCP with which ACF/NCP/VS Release 3 Can Communicate

Figure 7-4. Levels of Access Method with Which ACF/NCP/VS Release 3 Can Communicate

Chapter 8. Preinstallation Planning

A great many factors should be considered when an installation is planning to convert a data communications network from one operational level to another through changes in equipment or programs. The need for careful preinstallation planning increases as the network configuration becomes more complex and as the number of different applications it supports grows. While each installation is unique and no specific preinstallation planning procedure can be prescribed here, some general areas of planning can be identified. These are mostly posed as questions that the installation manager, or other person responsible for the conversion effort, must be able to answer.

Objectives

Before objectives can be set for the upgraded network, the characteristics of the present network must be considered.

What is the present network configuration? What applications presently use the network? Does the network meet the needs of the applications? If not, in what ways does it not? (Identifying these lacks can help in setting objectives for the upgraded network.)

How does the present network perform? For example, is the average response time experienced by users of the network satisfactory—during periods of average traffic flow in the network? during periods of peak flow? What backup plans are in effect for various kinds of failures?

It is important to document these and any other pertinent characteristics of the present network as a basis for comparison when planning the improved network.

What are the objectives of the improved network—adding new applications? upgrading existing applications, as by improving average response time? adding new types of stations? increasing flexibility of network operation, as by allowing more types of stations to communicate with the same application programs? arranging for further distribution of function into the network to relieve processing loads at the central site(s)? consolidating two or more separate networks to reduce communication line costs?

Must all of the identified objectives be met at once? If not, consider how the network can be upgraded in steps. Also consider what effect each step will have on existing network operations. Will present users of the network notice any changes, as in procedures for access to the network, or in response times? What can be done to minimize the effects on and possible inconvenience to present users?

These and similar questions should be answered early in the preinstallation planning process.

Resources

While objectives are being set for the new network, the resources needed to fulfill those objectives must be identified. Some resources are permanent, such as stations to be added to the network. Others are temporary, such as the programmer time needed to program new applications and the extra computer time that may be needed for testing the upgraded network before it goes into normal operation.

What equipment is needed: at the host processor site(s)? at remote sites? Will timely maintenance service be available for the new equipment? (This question is especially important for remote sites, which may be distant from centers of data processing expertise.)

What programming will be needed to meet the objectives of the upgraded network? Most of the enhanced functions of ACF/NCP/VS Release 2 and Release 3 require the complementary functions that ACF/TCAM Version 2 (Release 1, 2, or 3) and ACF/VTAM Release 2 and Release 3 provide. Should the access method(s) be upgraded first? Or should the network control program(s) be upgraded first? Are the program resources compatible? For example, will the functions provided by the network control program be supported by the access method? by the operating system? by programming subsystems such as CICS/VS?

What service aids and programs (such as the Network Problem Determination Application) are required or desirable?

What communication services are required to implement the upgraded network? Will they be available in all the locations to be connected? Will they be available at the appropriate time?

What personnel resources will be required? What new people hired and trained? What existing people retrained? These questions should be asked and answered for both the network development effort and the continuing operations of the network after it is placed in service.

Network Control

A most important aspect of a complex network is the manner in which it will be controlled. Multiple-domain networks provide the opportunity for distribution of the control function between two or more network control operators.

Will the network be controlled from one site or more than one? At any moment, will the entire network be controlled from one site, or from more than one? If more than one, how will the control responsibility be divided? In case of failure of one point, will the other point be able to assume control of the entire network? How will the multiple control sites coordinate their activities if failure of one part of the network interrupts communication between them?

Backup and recovery procedures are especially important in a complex network. Each possible kind of failure should be anticipated.

What will be the effect on the network of failure of a host processor or access method? of a communications controller or network control program? of a communication line, a modem, a cluster controller, or a terminal? of a cross-domain link? In each case: what problem determination procedures will be used? how can the effect on the network be minimized? What backup resources (such as alternate communication lines) should be available?

What procedures are required by the network control operator(s) and users of the network to operate in backup mode? to return to normal operation? to notify network users of changing to backup mode and reverting to normal operation?

These procedures should be documented and made available to the individuals who will be affected.

Test Plans

Preinstallation planning for upgrading of a network should include appropriate testing to make sure not only that new applications and network functions work correctly, but also that existing applications and network functions are not impaired. A test plan should be prepared and agreed to by all who may be affected.

What new functions or applications will the upgraded network perform? How should each be tested? What existing functions could be impaired through conflict with the new ones? Should old and new functions be tested at the same time?

Exceptional conditions, as well as normal network operations, should be tested. Consideration should be given to testing the network under peak traffic loads as well as during periods of normal traffic. Backup and recovery functions should be tested by forcing error conditions, as, for example, turning off power on various network components such as controllers and modems and verifying that the network recovers properly from that action.

Training

Some training will be necessary for users and personnel responsible for network operations: terminal operators, computer operators, network control operators, among others. System programmers and others involved in the development effort may need to learn about new versions of access methods, network control programs, and other programs to be used in the upgraded network. If network control programs are to be modified to meet special needs, the programmers responsible for the activity will need to become intimately familiar with the internal logic of the network control program and knowledgeable about Systems Network Architecture.

As mentioned earlier, network operating procedures, including those for backup and recovery from failure conditions of various kinds, should be documented. Some training of operators and others may be required for them to understand these procedures.

Scheduling

During consideration of all of the factors outlined above, and any others that may apply to a given installation, the question of timing will necessarily arise. Once all of the network requirements and the resources to meet them are identified, a detailed schedule of events should be developed in cooperation with all responsible individuals affected.

The schedule should reflect realistic estimates of the availability of the needed resources, with contingency factors included. Some resources will take longer to put in place than others. Lead times for new equipment and installation of communication services, especially, may be considerable. Also important is a realistic assessment of the personnel resources involved in developing the network improvements. For example, will the needed system programmers be available and have been appropriately trained at points in the schedule when they are required?

When the detailed schedule of events is developed from all of the preceding factors, it should be agreed to by all who will be providing resources and by those responsible persons whose operations will be affected by the network improvement effort. The proper review and approval of the detailed schedule should minimize the problems in the development process that could arise due to misunderstandings as to priorities and responsibilities involved in the process.

Appendix A. Communication Stations and Systems Supported by ACF/NCP/VS

This Appendix lists the communication products supported by ACF/NCP/VS Releases 1, 2, and 3. ACF/NCP/VS supports these products when they are attached to the 3705-I or 3705-II communications controller by either nonswitched or switched communication lines, except as noted.

SNA-SDLC Stations and Systems

IBM 3271 Control Unit (Models 11, 12)[1]
IBM 3274 Control Unit (Model 1C)[1]
IBM 3275 Display Station (Models 11, 12)[1]
IBM 3276 Control Unit Display Station (Models 11, 12, 13, 14)
IBM 3601 Finance Communication System
IBM 3602 Finance Communication System
IBM 3614 Consumer Transaction Facility[1]
IBM 3624 Consumer Transaction Facility[1]
IBM 3651 Store Controller (Models A50, B50)
IBM 3651 Store Controller (Models A60, B60)[2]
IBM 3661 Store Controller (U.S.A. only)[2]
IBM 3767 Communication Terminal (Models 1,2,3)
IBM 3770 Data Communication System:
  IBM 3771 Communication Terminal (Models 1, 2, 3)
  IBM 3773 Communication Terminal (Models 1, 2, 3, P1, P2, P3)
  IBM 3774 Communication Terminal (Models 1, 2, P1, P2)
  IBM 3775 Communication Terminal (Models 1, P1)
  IBM 3776 Communication Terminal (Models 1, 2, 3, 4)
  IBM 3777 Communication Terminal (Models 1, 3)
IBM 3791 Communications Controller
IBM Series/1

Binary Synchronous (BSC) Stations and Systems

IBM 1131 Central Processing Unit
IBM 1826 Data Adapter Unit
IBM 2715 Transmission Control (Model 2)
IBM 2772 Multipurpose Control Unit
IBM 2780 Data Transmission Terminal
IBM 2972 Station Control Unit (Models 8, 11) (U.S.A. only)[1]
IBM 3271 Control Unit (Models 1, 2)[1]
IBM 3275 Display Station (Models 1, 2)[1]
IBM 3735 Programmable Buffered Terminal
IBM 3741 Data Station (Models 2, 4)
IBM 3747 Data Converter
IBM 3750 Switching System (World Trade only)[1]
IBM 5110 Portable Computer
IBM System/3
IBM System/360 Model 20
IBM System/360 Models 25, 30, 40, 50, 65, 65MP, 67 (in 65 mode),
  75, 85, 91, 195
IBM System/370 Models 115—168MP

---

[1]nonswitched lines only
[2]switched lines only

Start-Stop Stations

IBM 1051 Control Unit
IBM 2740 Communication Terminal (Model 1)
IBM 2740 Communication Terminal (Model 2)[1]
IBM 2741 Communication Terminal
AT & T 83B3 line control type[1] (U.S.A. only)
Western Union 115A line control type (U.S.A. only)
CPT-TWX (Models 33, 35) line control type[2] (U.S.A. only)
World Trade teletypewriter (teleprinter) terminals[1]

Compatible Stations

The following stations are supported as compatible with the types of stations indicated in parentheses:

SNA-SDLC Stations

IBM 3274 Control Unit (Model 1C) (as a 3791)
IBM 3276 Control Unit Display Station (Models 1, 2, 3, 4, 11, 12, 13, 14) (as a 3791)
IBM 3631 Plant Communication Controller (as a 3601 or 3602)
IBM 3632 Plant Communication Controller (as a 3601 or 3602)
IBM 3771 Communication Terminal (Models 1, 2, 3) (as a 3767)
IBM 3773 Communication Terminal (Models 1, 2, 3) (as a 3767)
IBM 3774 Communication Terminal (Models 1, 2) (as a 3767)
IBM 3775 Communication Terminal (Models 1, 2) (as a 3767)
IBM 5937 Industrial Terminal (as a 3270)
IBM 8130 Processor (as a 3276 Model 11, 12, 13, or 14 or a 3791)
IBM 8140 Processor (as a 3276 Model 11, 12, 13, or 14 or a 3791)
IBM System/32 (as a 3770)
IBM System/34 (as a 3767, 3770, or 3791)
IBM System/38 (as a 3770)

Binary Synchronous (BSC) Stations

IBM 3274 Control Unit (Model 1C) (as a 3271 Model 1 or 2)
IBM 3276 Control Unit Display Station (Models 1, 2, 3, 4) (as a 3271 Model 1 or 2)
IBM 3770 Data Communication System terminals (as a 2772)
IBM 3780 Data Communication Terminal (as a 2772)
IBM 5275 Direct Numerical Control Station (as a 3275 Model 1 or 2)
IBM 5937 Industrial Terminal (as a 3270 Model 1 or 2)
IBM 8130 Processor (as a 3271 Model 1 or 2 or a 2772)
IBM 8140 Processor (as a 3271 Model 1 or 2 or a 2772)
IBM System/7 Processor Station (as a System/3)
IBM System/32 Batch Work Station (as a System/3)
IBM System/34 (as a 3767, 3770, or 3791)
IBM System/38 (as a 3770)

---

[1]nonswitched lines only
[2]switched lines only

Start-Stop Stations

IBM 3767 Communication Terminal (Models 1 and 2) (as a 2740 Model 1 or 2 or a 2741)
IBM 3767 Communication Terminal (Model 3) (as a 2740 Model 2)
IBM 5100 Portable Computer (as a 2741)
IBM 5110 Portable Computer (as a 2741)
IBM Communicating Magnetic Card Selectric— Typewriter (as a 2741)[2]
IBM System/7 Processor Station (as a 2741)

Equivalent Stations

Stations that are functionally equivalent to those listed above may also operate satisfactorily with ACF/NCP/VS; the customer is responsible for establishing equivalency.

---

[1]nonswitched lines only
[2]switched lines only

Appendix B. ACF/NCP/VS in a Remote Communications Controller

A local communications controller is a controller directly attached to a host processor data channel. A remote communications controller is one located at a distance from the host processor, and communicates with it via a local communications controller over a nonswitched SDLC link. If provided with the required hardware and program options, the two controllers can alternatively communicate via the switched telephone network in case the regular nonswitched link fails.

Multiple remote communications controllers may be connected to the same local controller, each over a separate SDLC link.

The same communications controller may be channel-attached to a host processor and remotely connected to a different host processor via a local controller.

ACF/NCP/VS in a remote communications controller supports the same types of SNA and non-SNA stations as it does in a local communications controller.

ACF/NCP/VS Definition and Generation

An ACF/NCP/VS for a remote communications controller is defined in the same way as for a local controller. Generation parameters that reflect channel and line connections may differ; for example, the absence of a channel adapter in the remote controller would be reflected in the generation source statements. Further, when executing in a remote controller, ACF/NCP/VS cannot operate lines attached to the controller in emulation mode. That is, ACF/NCP/VS cannot be generated with the partitioned emulation programming (PEP) extension. (See Appendix C for information about the PEP extension.)

A program generation parameter specifies whether the ACF/NCP/VS is to be executed in a local or a remote communications controller. Beyond this, there is no difference in the way the program is generated for the two kinds of controllers, and no difference in the network control functions the programs perform.

Appendix C: Partitioned Emulation Programming Extension

When executing in a local communications controller, ACF/NCP/VS is capable of operating communication lines in emulation mode, as well as in network control mode, through the partitioned emulation programming (PEP) extension to ACF/NCP/VS.

The PEP.extension allows the communications controller to emulate the operation of an IBM 2701 Data Adapter Unit or a 2702 or 2703 Transmission Control (or any combination of the three) for certain communication lines, while performing network control functions for other lines.

When defining ACF/NCP/VS, the system programmer specifies each communication line in the network as being operable in network control mode, or in emulation mode, or both. If specified in either network control or emulation mode, the line always operates in that mode. If specified in both modes, the line can be alternated between modes by command from the network operator. Program generation parameters specify the mode in which the line is to be placed when the ACF/NCP/VS is initiated.

The partitioned emulation programming (PEP) extension allows many programs written for support of the IBM 2701, 2702, and 2703 to operate with the communications controller without modification. These programs include (1) IBM Type I access methods that support the 2701, 2702, and 2703 and (2) IBM Type II and III programs and user-written programs that support these units in a manner equivalent to Type I access method programs. Programs that involve timing dependencies and support of certain special and custom features may have to be changed.

Advantage of the PEP Extension

The principal advantage of the PEP extension is that it allows concurrent operation of (1) existing application or access method programs, designed to communicate with a 2701, 2702, or 2703; and (2) new (or converted) application programs designed to communicate with ACF/NCP/VS.

With concurrent operation, not all existing application or access method programs need be converted before the benefits of operating communication lines in network control mode, rather than emulation mode, can be realized. Programs may be converted gradually, over any desired period of time. As conversion of each program is completed, the communication lines associated with that program may be changed from emulation mode to network control mode. Spreading the conversion process over an extended period allows each program to be more thoroughly tested than if all programs had to be converted at once. However, the sooner all programs are converted, the sooner the installation will realize the full benefits of operating all lines in network control mode.

Conversion of any existing application programs may be deferred until new application programs, designed for use with ACF/NCP/VS, are completed. This may be advantageous when installation of ACF/NCP/VS coincides with the development of new applications for the network.

Further information on operations in emulation mode, including channel-attachment requirements for the communications controller, may be found in *Introduction to the IBM 3704 and 3705 Communications Controllers,* GA27-3051.

Appendix D. Glossary of Terms and Abbreviations

This glossary contains definitions reproduced from the *American National Dictionary for Information Processing,* copyright 1977 by the Computer and Business Equipment Manufacturers Association, copies of which may be purchased from the American National Standards Institute at 1430 Broadway, New York, New York 10018.

access method: A technique for moving data between main storage and input/output devices.

ACF: Advanced Communications Function

ACF/NCP/VS: Advanced Communications Function/Network Control Program/Virtual Storage *See also network control program.*

ACF/TCAM: *Advanced Communications Function for the Telecommunications Access Method.*

ACF/VTAM: *Advanced Communications Function for the Virtual Telecommunications Access Method*

adjacent domains: Domains sharing a common network control program (ACF/NCP/VS) or two domains connected by a cross-domain link. See also *domain.*

adjacent subareas: Two subareas connected by (1) a subchannel, (2) a local-local link, or a local-remote link. See also *subarea.*

Advanced Communications Function (ACF): A group of IBM program products (principally ACF/TCAM, ACF/VTAM, and ACF/NCP/VS), that use concepts of Systems Network Architecture (SNA), including distribution of function and resource sharing. The Multisystem Networking Facility of ACF/TCAM, ACF/VTAM, and ACF/NCP/VS allows the interconnection of two or more domains into one consolidated and coordinated multiple-domain network.

application layer: In SNA, the functional layer of each individual session in which the end user's application is executed. See also *function management layer* and *transmission subsystem layer.*

application program: (1) A program written for or by a user that applies to a particular application. (2) A program used to connect and communicate with terminals in a network, enabling users to perform application-oriented activities.

asynchronous: Without regular time relationship; unexpected or unpredictable with respect to the execution of a program's instructions.

AT&T: American Telephone & Telegraph Company.

basic information unit (BIU): In SNA, a unit of data and control information that is passed between connection point managers. It consists of a request-response header (RH), followed by a request-response unit (RU).

basic link unit (BLU): In SNA, a unit of data and control information transmitted over a data link by data link control.

basic transmission unit (BTU): In SNA, the unit of data and control information passed between path control components. The BTU can consist of one or more path information units (PIUs), depending on whether blocking of PIUs is done by the path control component that builds the BTU.

binary synchronous communication (BSC): Communication using binary synchronous line discipline. See *binary synchronous transmission.*

binary synchronous transmission: Data transmission in which synchronization of characters is controlled by timing signals generated at the sending and receiving stations. Contrast with *start-stop transmission* and *synchronous data link control.*

bind image: A string of parameters in a Bind Session request that specifies the protocols for an LU-LU session.

BIU: Basic information unit.

BIU segment: In SNA, a portion of a basic information unit. A large BIU can be divided into segments for piecemeal transmission to accommodate the receiver's buffer-size constraints. See also *segment.*

blocking of PIUs: In SNA, combining multiple path information units (PIUs) into a basic transmission unit.

BLU: Basic link unit.

bps: Bits per second.

BSC: Binary synchronous communication.

BTU: Basic transmission unit.

buffer: (1) *A routine or storage used to compensate for a difference in rate of flow, or time of occurrence of events, when transmitting data from one device to another. (2) An area of storage that is temporarily reserved for use in performing an input/output operation, into which data is read or from which data is written.

chain: For logical units, one or more related units of data; each unit is identified as the first, middle, last, or only element of a chain. Each physical transmission is a chain.

channel: A device that connects the host processor and main storage with the I/O control units.

channel adapter: A communications-controller hardware unit used to attach the controller to a System/370 channel.

communication common carrier: In the USA, a government-regulated private company that furnishes the general public with telecommunication service facilities; for example, a telephone or telegraph company.

communication line: A physical connection, such as a wire or a telephone circuit, between communications controllers or between a communications controller and a station. Synonym for line. See also *data link.*

communications controller: A type of communication control unit whose operations are controlled by a program stored and executed in the unit. It manages the details of line control and routing of data through a network. It can route data to the host processor, or it can route data to or from a cluster controller or terminal. Examples are the IBM 3704 and 3705 Communications Controllers.

configuration services: In SNA, one of the types of network services in the system services control point; these services activate, deactivate, and maintain the status of physical units, logical units, and data links. They also start up, shut down, and restart network elements. See also *maintenance services, session services,* and *system services control point.*

connection point manager: In SNA, a component of the transmission subsystem layer that provides a common mechanism by which session control, network control, and network addressable units communicate with their corresponding elements through the network. The unit of information handled by the connection point manager and interpreted by the receiving connection point manager is the request/response header (RH).

cross-domain communication: In a multiple-domain network, communication between domains; synonymous with *networking.*

cross-domain LU: A logical unit (LU) located in another *domain.*

cross-domain session: A session between logical units (LUs) in different domains; a cross-domain LU-LU session. Contrast with *same-domain session.*

data link: A generic term for any physical facility used to carry data from one resource to another; examples are channels, communication lines, and communication loops.

data link control: (1) The noninformation exchanges that set up, control, check, and terminate the information exchanges between two stations on a data link. (2) In SNA, a part of the transmission subsystem layer. It initiates, controls, checks, and terminates the data transfer over a data link between two nodes. See also *path control.*

distributed function: The use of programmable stations to perform operations that were previously done by the host processor, such as network control, processing, and error-recovery operations.

domain: The collection of network resources controlled by one system services control point (SSCP). Synonymous with *single-domain network.*

DOS/VS: Disk Operating System/Virtual Storage.

emulation mode: The function of a network control program that enables it to emulate a transmission control unit. Contrast with *network control mode.*

emulation program (EP): A control program that allows a local IBM 3704 or 3705 Communications Controller to emulate the function of an IBM 2701 Data Adapter Unit, an IBM 2702 Transmission Control, and/or an IBM 2703 Transmission Control. Contrast with *network control program.*

end user: In SNA, the ultimate source or destination of information flowing through the network. It may be an application program, an operator, or a physical device medium (such as cards of tapes).

EP: Emulation program.

function management layer: In SNA, the layer of functional capability between the application layer and the transmission subsystem layer.

host access method: The access method, either ACF/TCAM or ACF/VTAM, that controls data communication with a domain.

host LU: A logical unit (LU) in a host processor.

host node: Synonym for *host processor.*

host processor: The System/370 with its operating system, access methods, and application programs; the host processor oversees the entire domain. The system services control point (SSCP) is located in the host processor. Synonymous with *host* and *host node.*

I/O: Input/output.

K: Thousand (1 024, when referring to bytes of storage).

link: See *communication line* and *data link.*

link header: Control information that is added to the beginning of a basic link unit.

link trailer: Control information that is added to the end of a basic link unit.

load module: A program in a format suitable for loading into storage for execution.

logical unit (LU): One of three types of network addressable units; an end user's means of accessing the facilities and services of a network in order to communicate with another end user. See also *host LU, outboard LU, primary LU,* and *secondary LU.*

logical unit services: The portion of the function management layer that (1) supports SSCP-LU sessions and (2) aids in initiating and terminating LU-LU sessions.

LU: Logical unit.

LU-LU session: In SNA, a session between two logical units in the network. It provides communication between two end users, each associated with one of the logical units. See also *cross-domain session* and *same-domain session.*

maintenance services: In SNA, one of the types of network services in the system services control point (SSCP). These services provide facilities for testing the links and nodes for collecting and recording error information. See also *configuration services, session services,* and *system services control point.*

MCP: Message control program.

message: A combination of characters and symbols transmitted from one point to another in a network.

message header: The leading part of a message that contains information such as the source or destination code of the message, the message priority, and the type of message.

multiple-domain network: A data communication network with more than one system services control point (SSCP). Contrast with *single-domain network.*

multipoint line: A line or circuit interconnecting several stations. Contrast with *point-to-point line.*

network: (1) The assembly of equipment through which connections are made between installations. (2) A configuration in which two or more station installations are connected. See also *domain, multiple-domain network,* and *single-domain network.*

network addressable unit: In SNA, a logical unit, a physical unit, or a system services control point. It is the origin or the destination of information transmitted in the transmission subsystem layer. Each network addressable unit has a network address that represents it to the transmission subsystem layer.

network control mode: The functions of a network control program that enable it to direct a communications controller to perform activities such as polling, device addressing, dialing, and answering. Contrast with *emulation mode.*

network control program: A control program for the IBM 3704 and 3705 Communications Controllers generated by the user from a library of IBM-supplied modules. ACF/NCP/VS operates only in the IBM 3705 Communications Controller.

networking: The concept of communication between domains in a multiple-domain network.

node: A junction point in a network represented by a physical unit; an addressable point in a network. The host processor, communications controllers, cluster controllers, and some SNA stations are nodes.

nonswitched line: A connection between a remote station and a communications controller that does not have to be established by dialing. See also *point-to-point line* and *multipoint line.*

OS/VS: Operating System/Virtual Storage.

outboard LU: An SNA logical unit located outside a host processor in an SNA cluster controller or terminal Contrast with *host LU.*

partitioned emulation programming (PEP) extension: A function of a network control program that enables a communications controller to operate some communication lines in network control mode while simultaneously operating others in emulation mode.

path control: In SNA, one of the components of the transmission subsystem layer. It is responsible for managing the sharing of data link resources of the network and for routing basic information units (BIUs) through it. See also *data link control.*

path information unit (PIU): In SNA, the unit of transmission consisting of a transmission header (TH) and either a basic information unit (BIU) or a BIU segment.

PEP: Partitioned emulation programming.

physical transmission: (1) For non-SNA devices, the amount of data entered on a line during an entire transmission sequence, from the first byte of data to the end-of-transmission character. (2) For logical units, a physical transmission is a chain.

physical unit (PU): In SNA, one of three types of network addressable units; a PU is associated with each mode whose existence has been defined to the system services control point (SSCP). A physical unit controls the resources local to its associated node. The SSCP establishes a session with the physical unit as part of the bring-up process.

PIU: Path information unit.

point-to-point line: A line that connects a remote station to a host processor; it may be either switched or nonswitched. Contrast with *multipoint line.*

polling: A technique by which a device is periodically interrogated to determine whether it needs servicing.

primary LU: A logical unit (LU) that issues a Bind Session request to establish an LU-LU session; the logical unit that controls an LU-LU session. Contrast with *secondary LU.*

request: In SNA, synonym for *request/response unit.*

request/response header (RH): In SNA, a control field, attached to a request/response unit (RU), that specifies the type of RU being transmitted—request or response—and contains control information associated with that RU. It is used by sending and receiving connection point managers to coordinate data traffic between network addressable units. See also *request/response unit* and *connection point manager.*

request/response unit (RU): In SNA, the basic unit of information entering and exiting the transmission subsystem layer. It may contain data, acknowledgment of data, commands that control the flow of data through the network, or responses to commands.

response: In SNA, synonym for *request/response unit.*

RH: Request/response header.

RU: Request/response unit.

same-domain session: An LU-LU session between logical units (LUs) in the same domain. Contrast with *cross-domain session.*

SCP: Systems control program.

SDLC: Synchronous data link control.

secondary LU: In an LU-LU session, the partner that receives the Bind Session request. Contrast with *primary logical unit.*

segment: A portion of a message that can be contained in a buffer. See also *BIU segment.*

session services: In SNA, one type of network services provided by the system services control point (SSCP). These services provide facilities for a logical unit or network operator to request that the SSCP establish, terminate, or alter sessions between logical units. See also *system services control point, maintenance services,* and *configuration services.*

single-domain network: A data communication network with one system services control point (SSCP). Contrast with *multiple-domain network.*

SNA: Systems network architecture.

SSCP: System services control point.

SSCP backup: The changing of domain boundaries so that network resources move from one domain to another.

SSCP-LU session: An SNA session between the system services control point (SSCP) and a logical unit (LU). It is used to support logical-unit-related control and use of the communication system. Each logical unit in the network must participate in a session with the SSCP that provides services for that logical unit.

SSCP-PU session: An SNA session between the system services control point (SSCP) and a physical unit (PU) that is used to control the physical configuration of a network and to control an individual node.

start-stop transmission: Asynchronous transmission whereby a group of bits is preceded by a start bit that prepares the receiving mechanism for the reception and registration of a character, and is followed by at least one step bit that enables the receiving mechanism to come to an idle condition pending the reception of the next character. Contrast with *binary synchronous transmisison* and *synchronous data link control.*

station: A point in a data communication network at which data can enter or leave.

subarea: In SNA: (1) A subfield in the network address. (2) The group of network addressable units sharing a common subarea address. The network address space is partitioned into subareas and each subarea is further divided into elements. Subareas are the basic units of routing in SNA.

switched line: A communication line in which the connection between the communications controller and a remote station is established by dialing.

synchronous data link control (SDLC): A discipline for managing synchronous, transparent, serial-by-bit information transfer over a communication channel. Transmission exchanges may be duplex or half-duplex over switched or nonswitched data links. The communication channel configuration may be point-to-point, multipoint, or loop. Contrast with *binary synchronous transmission* and *start-stop transmission.*

system services control point (SSCP): In SNA, a network addressable unit that provides configuration, maintenance, and session services via a set of command processors—network services—supporting physical units and logical unit. The SSCP must be in session with each logical unit and each physical unit for which it provides these services.

systems network architecture (SNA): The total description of the logical structure, formats, protocols, and operational sequences for transmitting information units through a communication system.

TH: Transmission header.

transmission header (TH): In SNA, a control field attached to a basic information unit (BIU) or to a BIU segment, and used by path control. It is created by the sending path control component and interpreted by the receiving path control component. See also *path information unit.*

transmission subsystem layer: In SNA, the innermost layer of the communication system. It provides the control in each session to route and move data units between network addressable units (NAUs), and to manage the NAUs and their interconnecting paths. See also *application layer* and *function management layer.*

TWX: Teletypewriter Exchange.

Index

ACF/NCP/VS
Network Control Program
System Support Programs
General Information

Order No. GC30-3058-1

READER'S COMMENT FORM

This form may be used to communicate your views about this publication. They will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

Note: *Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.*

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you would like a reply, complete the following (Please Print):

Your Name ______________________ Date __________

Company Name ______________________

Department ______________________

Street Address ______________________

City __________ State __________ Zip Code __________

Thank you for your cooperation. No postage stamp is 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 This Line

ADDITIONAL COMMENTS:

Please Do Not Staple

Fold and Tape ... Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Dept. E02
P.O. Box 12195
Research Triangle Park
North Carolina 27709

Fold and Tape ... Fold and Tape

IBM®

International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

Cut or Fold Along Line

ACF/NCP/VS Network Control Program System Support Programs General Information Printed in U.S.A. GC30-3058-1

International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

GC30-3058-1