



Washington
Systems
Center

Technical
Bulletin

OSI And SNA: A Perspective

J. H. Rutledge
Washington Systems Center

DAPS Code 0895
GG22-9225-00
April 1981

Washington Systems Center
Gaithersburg, Maryland
Technical Bulletin

OSI and SNA: A Perspective

James H. Rutledge

DAPS Code 0895
GG22-9225
April 10, 1981

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "as is" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

In this document, any references made to an IBM program product are not intended to state or imply that only IBM's program product may be used; any functionally equivalent program may be used instead.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

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.

A form for reader's comments is provided at the back of this publication. If the form has been removed, comments may be addressed to: Washington Systems Center, 18100 Frederick Pike, Gaithersburg, MD 20760.

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.

ACKNOWLEDGMENTS

I wish to extend my thanks to the following individuals: J. R. Aschenbrenner, F. P. Corr, W. F. Emmons, M. L. Hess, and G. D. Schultz. With little advance notice, they were called upon to review either the document or the associated presentation foils. Their efforts provided me with significantly improved insights into the relationship of OSI and SNA. I have endeavored to pass these on to you, the audience for this paper.

PREFACE

A presentation based on this paper was given by the author in session C110 of SHARE 56 in Houston, Texas. The date of the session was March 11, 1981.

CONTENTS

Introduction	1
OSI and SNA	2
Open Systems Interconnection	4
Systems Network Architecture	13
Comparison of SNA and OSI	27
Conclusions	34
Appendix - Architecture	35
References and Notes	42

LIST OF ILLUSTRATIONS

Figure 1.	Foil 12 - Intermediate Nodes and Peer-to-Peer Communications	5
Figure 2.	Foil 18 - Five Component Networks of SNA.	14
Figure 3.	Foil 19 - SNA Layers Derived from Nested Networks.	16
Figure 4.	Foil 21 - Virtual Route Implementation with Layers	22
Figure 5.	Foil 22 - Implementation of SNA CPU with Possible Mapping to OSI	24
Figure 6.	Foil 23 - Implementation of SNA Communication Controller	26
Figure 7.	Foil 27 - Adjacent Node Sense for X.25 Interpretation of OSI	30
Figure 8.	Foil 28 - SNA and Other	33
Figure 9.	Foil 6 - Party A and the Five Levels	36
Figure 10.	Foil 7 - Peer-to-Peer Communications	38
Figure 11.	Foil 8 - Data Formats for Movement from B to A	40

INTRODUCTION

This paper will discuss similarities and differences between Systems Network Architecture (SNA) of IBM and the ISO Reference Model for Open Systems Interconnection (OSI). The definitive document for SNA is an IBM manual entitled Systems Network Architecture Format and Protocol Reference Manual: Architectural Logic (15). The definitive document for OSI is International Standards Organization (ISO) "Draft Proposal 7498" (4). In working form, this document is also identified as ISO/TC97/SC16/537 Rev. It is dated December 3, 1980.

The discussions in this paper will be high level. They are not intended to exhaustively cover each relevant detail of both architectures. Rather they will present a high level discussion of the two architectures with emphasis on common objectives and the means of attaining these objectives. SNA is much further along in evolution than OSI. As a result, much more detail about SNA is available. Awareness of the stages of evolution is essential in making judgments about the two architectures.

The Appendix presents a general discussion of architecture with specific points on communications architectures. Layering, environmental, and data flow concepts common to both OSI and SNA are introduced. This should be read by those unfamiliar with layered architectures.

The first section is a brief discussion of the objectives and benefits associated with OSI and SNA.

Next will be specific discussions of OSI and then SNA. This is where the more advanced evolution of SNA will become apparent.

Finally will be discussions comparing the two structures followed by some suggested conclusions.

An effort has been made to reduce the technical detail in this paper. The paper has concepts and insights for the uninitiated in data communications. However, those that will gain the most from it are those with some awareness of current issues associated with data communications. A conceptual understanding of standards such as CCITT X.21 and CCITT X.25 or EIA RS-232 is necessary to appreciate some of the discussion on OSI. The reader is also presumed to have a familiarity with basic SNA concepts as defined by IBM and with the IBM product line that supports SNA. It is not the intent of this paper to provide a tutorial in either standards, SNA concepts, or IBM products.

OSI AND SNA

The following quotations from the draft (4) establish the environment to be addressed by OSI.

The purpose is "... to provide a common basis for the coordination of standards developments for the purpose of systems interconnection"

"OSI is concerned with the exchange of information between systems (and not the internal functioning of each individual system)."

"... A system is a set of one or more computers, the associated software, peripherals, terminals, human operators, physical processes, information transfer means, etc., that forms an autonomous whole capable of performing information processing.

"'Openness' ... refers to the mutual recognition and support of the applicable standards."

OSI addresses standardization of protocols that are required to allow communication among discrete data processing entities. There is a sense that the discrete entities have the capability to meaningfully process data when operating stand-alone. Enhanced data processing capabilities may result from interconnection with other data processing entities. There is also a sense that the data processing entities can be dissimilar in structure.

SNA provides a foundation for a unified teleprocessing strategy from IBM. The products supporting and guided by SNA constitute the implementation of the strategy. The primary benefits accruing from the strategy are fourfold:

1. protection of customer application investment,
2. flexibility of product development in order to take advantage of new technologies,
3. uniformity of function to be supported by communications products, and the
4. creation of an application foundation that facilitates addressing new opportunities.

Since SNA is the basis for a coherent product offering rather than the basis for interconnecting dissimilar systems, it goes beyond OSI in scope. SNA includes a control structure that provides for management of the network resources. Additionally, there are increasing systems management offerings to assist in administration of the problems and change associated

with any complex system. These offerings can grow to be quite complete since the entire communications environment functions under one set of rules.

As can be seen, OSI and SNA are considerably different in their origins and objectives. However, if a discussion of SNA is restricted to inter-system communication with no consideration for resource control and management, much similarity between OSI and SNA can be established. Although this similarity will be shown in the next sections, the reader must recognize that systems conforming to SNA will enjoy a higher degree of homogeneity than systems developed with no conformity to any architecture.

The next two sections individually address OSI and SNA. The following section offers some comparisons of the two.

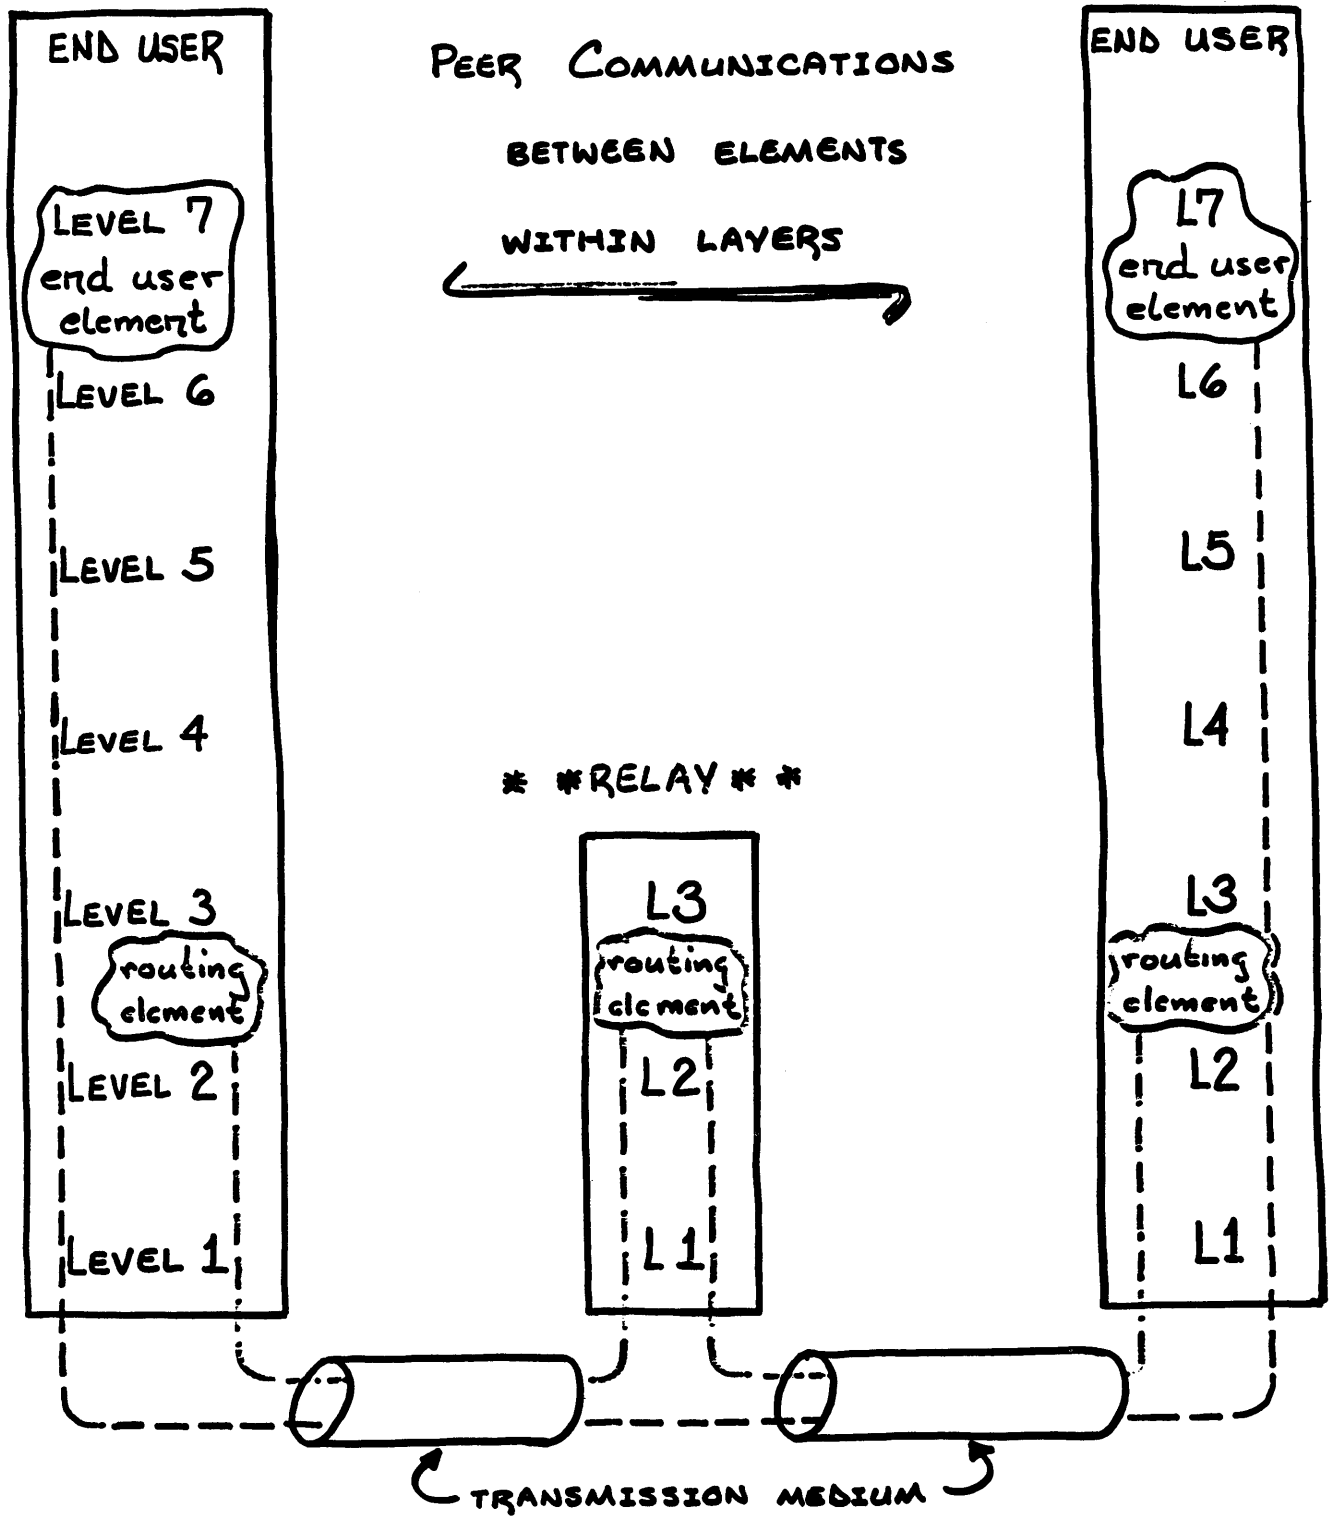
OPEN SYSTEMS INTERCONNECTION

OSI is a provisional model proposed by the International Standards Organization (ISO). It is to provide a framework for standards development supporting interconnection of systems using peer-to-peer communication between equivalent layers. OSI has seven layers. These seven layers describe the seven component processes that define the total data communications process for one party of a two-party conversation.

In the architecture, each layer (component process) must act in concert with its peer. The peer represents the other system in the conversation. Therefore, it takes at least two implementations of all seven layers to permit communication between two end-users. In addition, there may be intermediate implementations of layers concerned only with data transport. This is shown in Figure 1 on page 5.

The current OSI draft proposal (4) describes the relationship among the seven layers. Data formats and inter-layer protocols are in varying stages of development.

'Layer' is synonymous with 'level' in this OSI nomenclature. When referring to a level/layer by its number, level will be used, e.g., level 4. When referring to a level/layer by its name, layer will be used, e.g., the Transport layer.



12
CHR
8MAR88

Figure 1

The seven layers are discussed below.

- Level 7 - Application Layer

The OSI-defined communications process is represented to the user by this layer. Based on requests from the network user, this layer selects appropriate services to be supplied from lower layer functions. Lower layer functions are comprised of functions acting on behalf of a local party to the conversation as well as peer functions acting on behalf of the remote party.

As specified in the OSI draft (4), the following services are within the scope of the application layer acting on behalf of a particular network user:

- identification of intended communications partners and their availability and authenticity,
- establishment of authority to communicate,
- agreement on required privacy mechanisms,
- determination of cost allocation methodology,
- determination of resource adequacy to provide an acceptable quality of service,
- synchronization of cooperating applications,
- selection of dialog discipline,
- establishment of error recovery responsibility,
- agreement on data validity commitment,
- identification of data syntax constraints, and
- information transfer.

- Level 6 - Presentation Layer

The following definition is quoted from the OSI draft (4):

The purpose of the Presentation layer is to represent information to communicating application-entities in a way that preserves meaning while resolving syntax differences.

Toward this objective, this layer can provide the following functions:

- data transformation,
- data formatting, and

- syntax selection.

- Level 5 - Session Layer

The following definition is from the OSI draft (4):

The purpose of the Session layer is to provide the means necessary for cooperating presentation-entities to organize and synchronize their dialogue and manage their data exchange. To do this, the Session layer provides services to establish a session-connection between two presentation-entities, and to support their orderly exchange interactions.

In support of these objectives, the Session layer provides the following services to the Presentation layer:

- session-connection establishment,
- session-connection release,
- normal data exchange,
- expedited data exchange,
- quarantine service,
- interaction management,
- exception reporting, and
- mechanism for session-connection synchronization.

- Level 4 - Transport Layer

This layer exists to provide "... transparent transfer of data between session-entities." Transport protocols will have end-to-end significance. Transport layer users will be "... provided with the means to establish, maintain and release transport connections which represent a two-way simultaneous data path between a pair of transport-addresses."

The OSI proposal (4) defines three phases of operation within the Transport layer. They are listed below with associated services.

1. Establishment phase.

The objective is to establish connections between peer transport functions on behalf of service requests from

higher levels. The service quality of the connection can be negotiated during this phase. Services provided include:

- selection of network service as a function of parameters such as throughput, transit delay, set-up delay, and error characteristics,
- management of transport connections to lower level connections,
- establishment of appropriate data unit size,
- selection of usable functions for data transfer, and
- transport of data from higher levels.

2. Data transfer phase.

These services have the objective of providing data transfer in accordance with agreed-upon service quality from the establishment phase. Services provided include:

- blocking,
- concatenation,
- segmenting,
- multiplexing of connections provided by lower levels,
- flow control in a session-oriented, end-to-end sense,
- maintenance of the identity of data units received from the Session layer,
- maintenance of connection identity between the two transport functions acting on behalf of the parties of the conversation,
- error detection for lost, damaged, duplicated, misordered, or misdelivered data units,
- error recovery to address problems detected in this layer or signalled from lower levels, and
- transport of expedited data which flows outside normal flow control mechanisms.

3. Termination phase.

These services allow either end of the session to terminate the connection with notification to the other party. Services include:

- notification of termination reason,
- identification of connection terminated, and
- optional information as required.

The following condition from the OSI proposal (4) is noteworthy:

Only connection-oriented services are defined: transaction-oriented services and broadcast-oriented services are anticipated as future extensions to the basic definition.

• Level 3 - Network Layer

The basic function of this layer "... is to provide the transparent transfer of all data submitted by the Transport layer" while allowing "... the structure and detailed content of submitted data to be determined exclusively by levels above the Network layer."

The purpose is to allow the higher levels to have "... independence from routing and switching considerations associated with the establishment and operation of a ..." connection. The establishment, maintenance, and termination of connections on behalf of using parties is included in the service provided by this layer.

These functions and services are listed below.

- Network addressing and end point identification.
- Multiplexing network connections onto data link connections provided by the next lower level.
- Segmenting and/or blocking to facilitate data transfer.
- Service selection when different services are available.
- Selection of service quality based on parameters such as: residual errors, service availability, reliability, throughput, transit delay, and connection-establishment delay.

- Error detection and recovery to support desired quality of service.
- Error notification to higher levels when service quality cannot be maintained.
- Sequenced delivery of data if available from a particular implementation.
- Flow control. This is support of flow control indicators provided at one end by the Transport layer. This does not address global flow control issues. Internal network flow control requirements may exist that can be reflected to the user at the network interface.
- Expedited data transfer as an optional service.
- Connection reset with loss of enroute data and notification to using parties.
- Termination services when requested by a using party.

Current thinking for Network layer definitions centers around Level 3 of the CCITT X.25 specification (13). This recommendation addresses protocols and formats for communication between an end-user node and a network-access node. Accordingly, limited end-to-end protocols are defined. This is consistent with the OSI layer definitions.

The CCITT X.25 recommendation does not address protocols, formats, or services between network-access nodes. The owner of a data movement service with X.25 interfaces may address internal data link and physical layer functions in any manner that he chooses. This is an opportunity for implementation choices to serve the convenience of the transport service creator.

• Level 2 - Data Link Layer

The OSI proposal (4) states: "The Data Link layer provides functional and procedural means to establish, maintain and release data-link-connections among network-entities." The objectives are to provide data transmission services to the Network layer and "... to detect and possibly correct errors which may occur in the Physical layer."

Significant functional characteristics of this layer are listed below.

- Data-link-connection activation and deactivation. These functions include the use of physical multipoint facilities to support connections between peer Network layer functions.

- Mapping data units provided from the Network layer into data-link-protocol units for transmission
- Multiplexing of one data-link-connection onto several physical connections.
- Delimiting of data-link-protocol units for transmission.
- Error detection, recovery, and notification as appropriate.
- Identification and parameter exchange with peer data link parties.

Two current examples of Data Link layer implementations are:

1. High-level Data Link Control (HDLC). This is an ISO-developed link protocol that has enjoyed a favorable reception in the international data communications community.
2. Link Access Protocol B (LAP-B). This is a subset of HDLC asynchronous balanced mode (ABM) protocols. It is identified as an option for Level 2 of the CCITT X.25 specification (13).

The reader should understand that different interpretations of OSI could support different protocols in this layer, e.g., SDLC, ADCCP, etc. Support of different protocols here could also be driven by implementation or market factors. For example, the access protocols associated with well-known packet networks include BSC and asynchronous disciplines among others.

- Level 1 - Physical Control Layer

The OSI proposal (4) states: "The Physical layer provides mechanical, electrical, functional and procedural characteristics to activate, maintain and deactivate physical connections for ... transmission of transparent bit streams between ..." using parties.

The current focus for new networks is on the CCITT X.21 specification (6). Existing equipment and associated market requirements have caused the evolution of two modes of X.21 bis;

1. a V.24/RS-232 mode, and
2. a V.35 mode.

This completes the discussion of OSI. The next section will describe SNA. Some differences between OSI and SNA will be identified under appropriate topics. The following section will highlight the significant differences between OSI and SNA.

SYSTEMS NETWORK ARCHITECTURE

When SNA was first announced in 1974, we described three layers:

- the transmission subsystem layer,
- the function management layer, and
- the application layer.

The following definitions of the SNA layers are given (12):

... The transmission management layer controls movement of user data through the network, independently of the contents of the user data unit.... A transmission management layer exists in every intermediate node through which data units flow, and transmission management may utilize a variety of physical connections and protocols between the nodes of an SNA network.

The function management layer controls the presentation format of information sent from and received by NAU services manager layer.

Function management also manages the protocols supporting the exchange of user information.

The services of the function management layer are invoked by requests from the application layer. In the computer, the application layer consists of the application programs from which the terminal user requests information-processing services. At the terminal, the application layer is represented by the terminal operator or an application in a programmable control unit. SNA refers to these sources destinations as 'end users'

As SNA evolved, particularly into interconnected CPU environments, refinements of the functions associated with these layers led to the definition of five nested networks. These are shown in Figure 2 on page 14 (15).

The correct interpretation of this picture follows, working from the outside toward the center.

Pairs of end-users desiring to communicate use NTWK.SNA to do so. The users are provided access to the network through an associated NAU Services Manager (16).

THE FIVE NETWORKS OF SNA

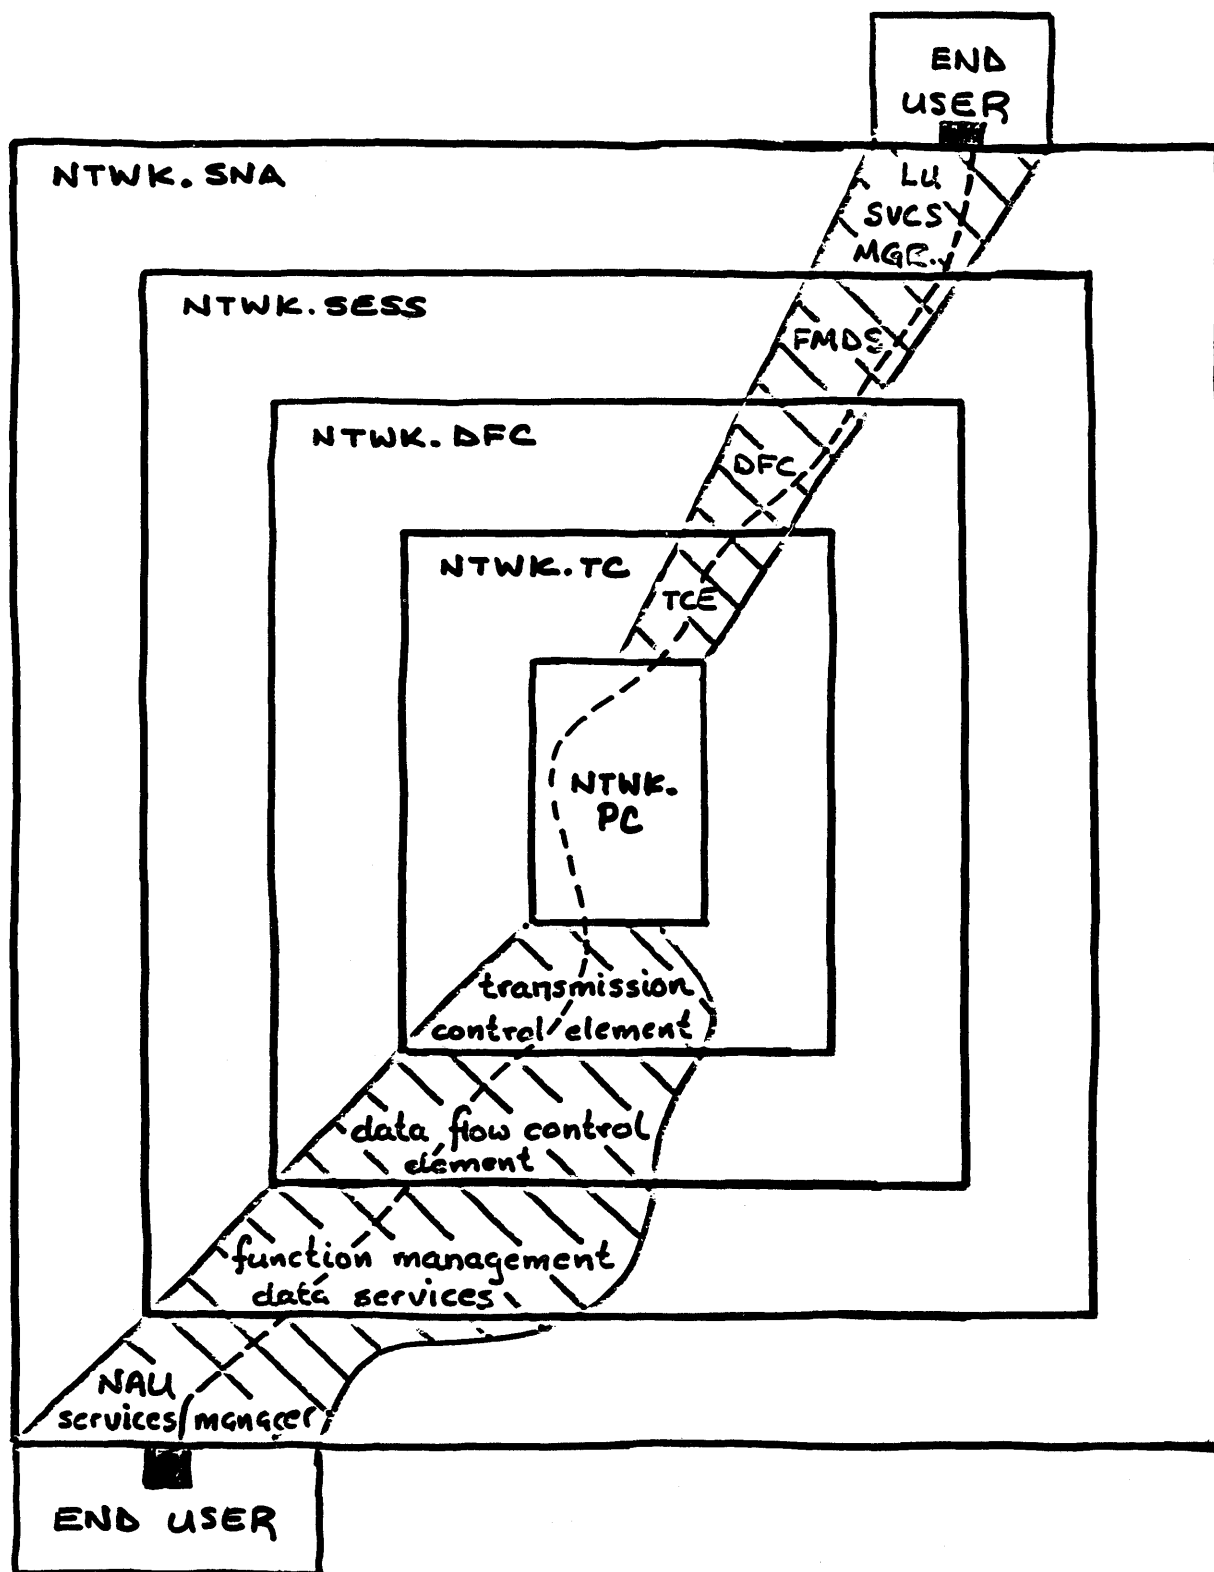


Figure 2

NAU Services Managers use NTKW.SESS for peer communications on behalf of the end users. The component process associated with the session on behalf of end users is the Function Management Data Services (FMDS).

Function Management Data Services will use NTKW.DFC. The DFC element pair acting on behalf of the of the end users will enforce data flow protocols chosen for the session.

The DFC element pairs will use NTKW.TC as a communication vehicle. The DFC element pairs are associated with transmission control element pairs for the provision and management of a two-way data flow for the end user session.

The transmission control element pairs will use NTKW.PC to "... route PIU's based on the destination and origin addresses"

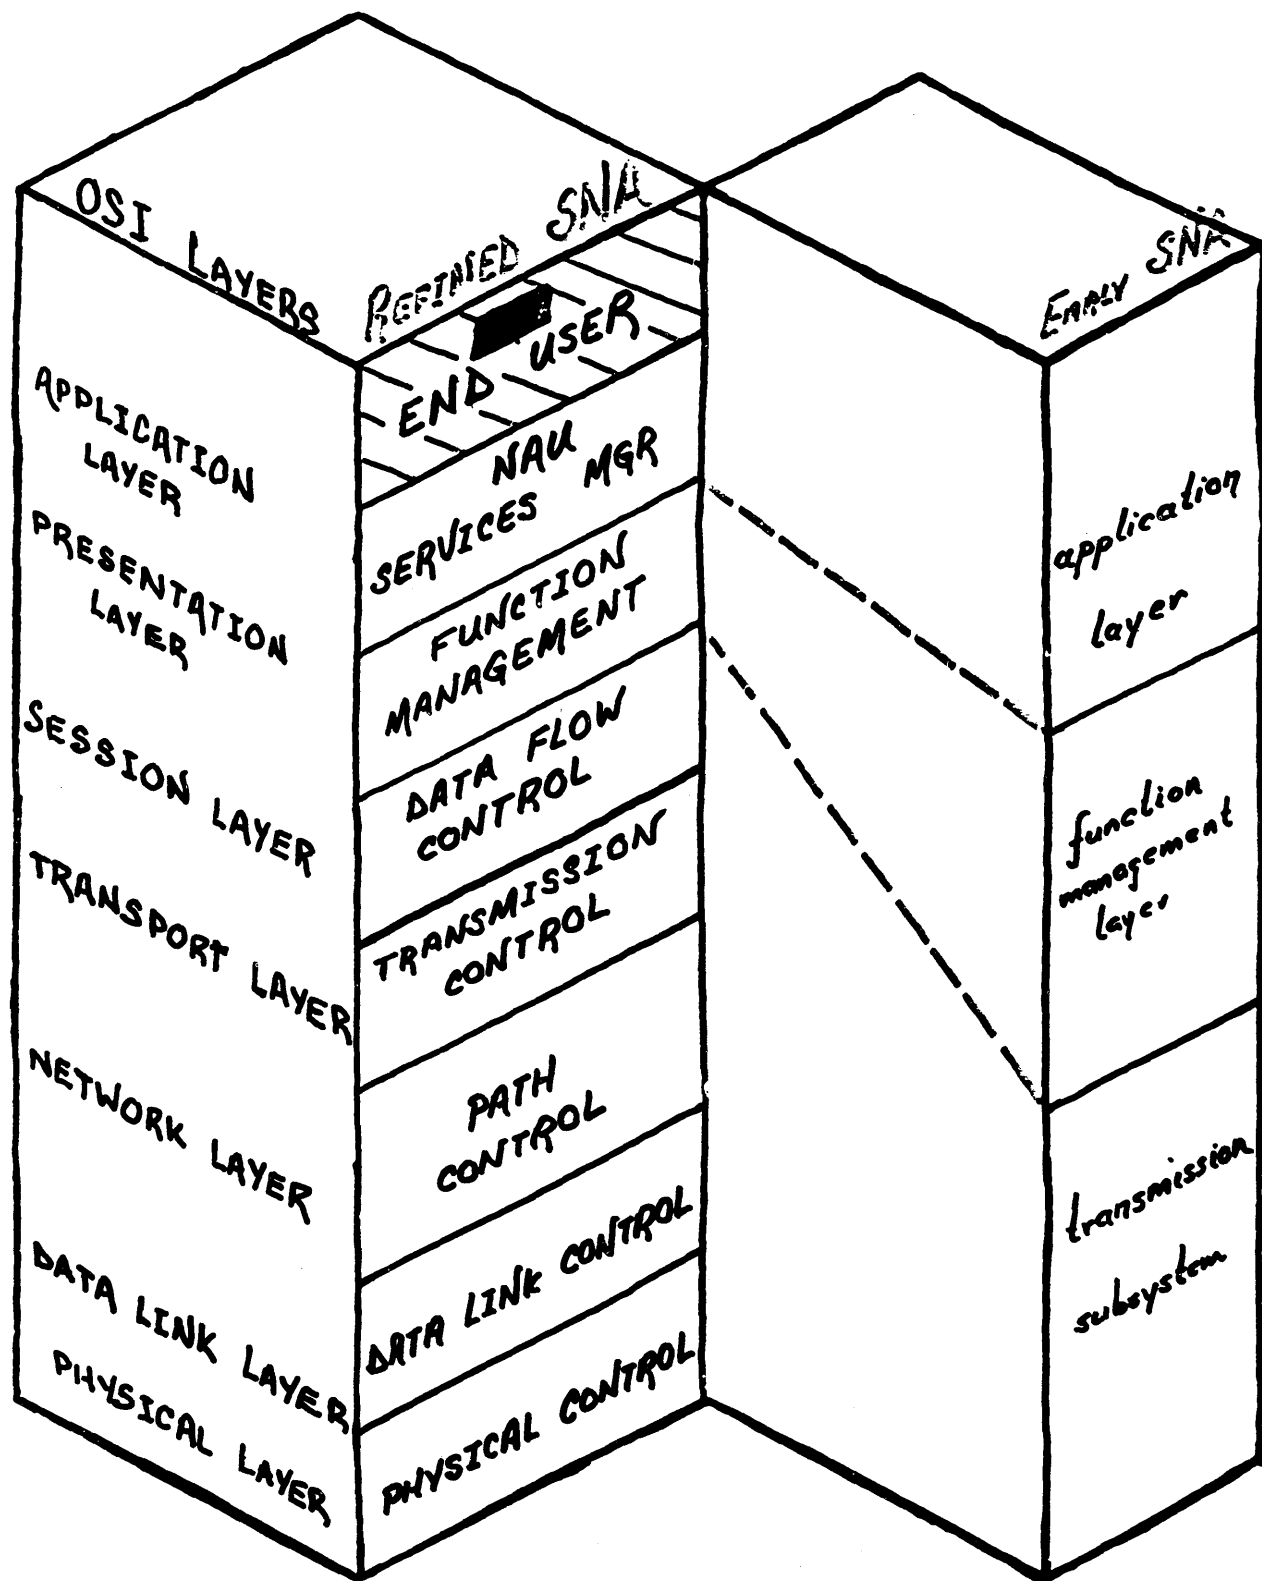
By taking three liberties with Figure 2 on page 14, we can derive a layered picture depicting an end user as shown in Figure 3 on page 16. The specific liberties are listed below:

1. taking a slice of the combined networks based on end user Party A shown in Figure 2 on page 14,
2. adding the Data Link Control layer under Path Control, and
3. illustrating the Physical Control function beneath Data Link Control.

On the left side of Figure 3 on page 16, the OSI layers are shown. This is not to suggest any one-to-one mapping between the OSI layers and the SNA layers. The suggestion is that the complete function addressed by SNA implementations for inter-system communication has some equivalence to the complete function modeled by OSI. Complete function includes all tasks that must be executed between the entry into the network of a service request from the end user and the transmission of that service request by the Physical layer onto the transmission medium.

SNA implementations also include internal network control facilities and network management features that are not addressed by the OSI model.

The seven SNA layers from Figure 3 on page 16 are described individually below. The basis for these descriptions is the SNA Format and Protocols Manual (15).



SNA LAYERS DERIVED FROM THE FIVE NETWORKS

19
JHR
9MARE

- **NAU Services Manager Layer**

There are three service manager functions included in this layer: systems services control point (SSCP) service manager, physical unit service manager, and logical unit service manager. The first two categories and companion portions of the logical unit services manager are associated with internal network control of its resources (17).

The remainder of the functions of the logical unit services manager are similar to the services suggested for OSI level 7, the Application layer.

An end user (party to a conversation) using the services of this layer attaches to a port into the SNA network. This port represents an application-entity that is referred to as a logical unit (LU) (16). This entity, or port, will communicate with its peer representing the other party of the desired two-party communication. It can communicate with network control functions such as the SSCP or an associated physical unit (PU). Accessing this port provides management services that initialize parameters for several layers.

Some specific examples of functions executed on behalf of the end user through the port provided by this layer are:

- resolution of network addresses from network names,
- checking of end-user access authority,
- selection and matching of session parameters,
- management and maintenance services,
- presentation services support, and
- sync point services.

Measurement services and network operator services are topics set aside for future architectural definition. Currently, services available in these areas are defined by the product developer(s).

- **Function Management Layer**

The services provided by this layer are similar to those suggested for OSI level 6, the Presentation layer. In SNA implementations, the services are provided by a function management data services (FMDS) element acting on behalf of an end user.

The FMDS element pair "... provides a connection for passing message units between pairs of ..." LU services managers representing end users.

This layer can provide format translation between services managers in accordance with format or presentation services available for the session. Specific examples are:

- device selection and control for multi-device workstations,
- data compression and compaction, and
- formatting data for diskettes when transmitting load modules for programmable devices.

Additionally, the FMDS element pair representing the end users will check and maintain current states for controlling and synchronizing some network services associated with session control.

- Data Flow Control Layer.

NOTE: Up until now, a one-for-one mapping between SNA layers and OSI layers may have been implied. While this has some validity for the top two and bottom two layers, it breaks down for the intermediate layers. The correct sense should be that the composite functions defined for OSI levels 5, 4, and 3, (Session, Transport, and Network layers respectively) are roughly addressed by the composite functions provided by the SNA layers of Data Flow Control, Transmission Control, and Path Control. Within these large segments, the boundaries are not distinct and do not map together.

The function of the Data Flow Control (DFC) layer is to control the flow of data between the FMDS element pairs of a session. Network and session control data is not managed by this layer.

Specific functions are listed below (15):

- enforcement of correct data formats,
- enforcement and checking of chaining,
- correlation of requests and responses including assignment of sequence numbers,
- enforcement of different response modes, e.g., immediate or delayed,

- coordination of sending and receiving according to session parameters, e.g., full-duplex, half-duplex contention, half-duplex flip-flop,
 - enforcement of bracket protocols,
 - enforcement of data flow suspension when requested, and
 - regulation of response queuing.
- Transmission Control Layer

There are three components of Transmission Control.

There are session control functions which include "... session-specific support for starting, clearing, and resynchronizing session-related data flows."

There are the connection point manager functions which include sequence number checking, enciphering/deciphering, separation of normal and expedited data flows, pacing enforcement, and routing of data to Data Flow Control and Path Control.

Finally, there is the boundary function. This is implemented within the network to support peripheral nodes. Peripheral nodes are isolated from global network considerations. The functions implemented on behalf of peripheral nodes include header transformation, address translation, routing to proper link station, optional segmenting of message units, session pacing, and coordination of local flow control with global flow control.

- Path Control Layer.

Path Control provides a full duplex path, independent of the physical configuration. Path Control performs the path-selection functions, ensuring that the correct transmission group or route extension is selected. Path Control assures the the message unit format is appropriate for the transmission medium. Path Control routes data over available links and through intermediate nodes enabling many end users to share common network resources.

There is a path control element in each SNA node.

There is a path control function associated with boundary function support of peripheral nodes.

Specific examples of Path Control functions are:

- segmenting and reassembly of message units,
- routing among SNA subareas,
- transmission group control,
- virtual route control including pacing,
- explicit route control,
- route extension control in support of boundary functions,

- Data Link Control Layer

This layer is defined below (15):

Data link control (DLC) supports protocols for (1) executing and coordinating the transfer of message units across a link connection between a single primary DLC user and a set of secondary DLC users, and for (2) performing link-level flow management and error recovery procedures.

Excepting transmission group management in Path Control, this layer is a functional equivalent to OSI level 2.

Current SNA implementation employ the following data link protocols:

1. S/370 channel protocols, and
2. SDLC.

Regarding the relationship of SDLC to HDLC, the following point is noteworthy (11).

It is IBM's technical judgement that SDLC, as implemented in IBM telecommunication products, conforms with a defined operational subset of ISO HDLC: the Unbalanced Normal Class of Procedure. An important point in understanding this technical judgement is that SDLC, as implemented in IBM telecommunication products, is more precise in certain aspects than the HDLC standards--both as approved and as currently proposed.

Support of channel data link protocols is required for the movement of data between the computer and the communications controller.

TCAM supports a wide range of link protocols which can be brought into an SNA environment. VTAM supports BSC as a special case for some models of the 3270 display family.

Software features of the ACF/NCP/VS such as NTO provide still another technique for supporting non-SDLC link protocols within an SNA network. Disciplines supported via NTO include the asynchronous disciplines associated with 2740/41 terminals, TWX terminals, and World Trade Telegraph terminals.

Devices attached to an SNA network by non-SDLC disciplines often do not enjoy the full range of functional support accorded those devices that were developed in conformity with SDLC and SNA.

- Physical Control Layer

Implementations of this layer manage the physical interface with the attached transmission medium. This includes data presentation, interface control, error detection, recovery, and notification.

This layer is the functional equivalent of the OSI level 1.


IBM SNA products support the following interfaces:

1. V.24 (RS-232-C),
2. V.35,
3. X.21,
4. X.21 bis, and the
5. S/370 channel interface.

This completes the discussion of SNA as an architecture. The remainder of this section will address three examples of SNA implementation.

The first example is Figure 4 on page 22. The important idea is that data transiting intermediate nodes will be operated upon by those layers that implement the transmission subsystem. The higher layers, more involved with end user issues, are only involved at the end points of the session or connection.

VIRTUAL ROUTE IMPLEMENTATION with LAYERS



21
JHR
9 MAR 81


Figure 4

This illustration suggests a mapping of IBM products implementing virtual routes to OSI layers. For this picture to be accurate, it is a requirement that the CPU's be executing a current ACF Release 3 access method (VTAM or TCAM) and the 3705's be executing ACF/NCP/VS Version 1 Release 3.

'Virtual route' is a term associated with ACF (Advanced Communications Function) Release 3, the most recent release of SNA implementation enhancements. "A virtual route identifies a full-duplex connection between two subarea nodes and only indirectly refers to physical connections" (7). For any two subareas that have a virtual route identified between them, end users in those subareas that desire to communicate may use any virtual route that is available. The selection is made at the time the session is set up.

The second implementation example is shown in Figure 5 on page 24. This figure shows how existing IBM and user programs in a mainframe might map against both the OSI layers and the SNA layers. There are 6 salient points to this picture as listed below.

1. The significance of user code is not properly emphasized. As shown, it sits above the Application or NAU Services Manager layers. Some user code can be in these layers. One should never lose sight of the tremendous investment that data processing users have in this code. I believe it is a marketplace reality that a successful implementation of lower layers will be one that minimizes impact on the layers above it. At the same time, lower layer implementations must address the issues of proliferation and flexibility loss that led to their creation.
2. Subsystems play a relatively small role. They can be thought of as providing: presentation services, a supporting environment for user application code, specialized support such as data bases, and supporting services such as checkpoint/restart. This is not to minimize the role of subsystems in relieving users of code development and maintenance as well as providing a stable framework for future enhancements and development.
3. The number of levels encompassed by the access method is an indicator of the function that is provided. The levels provide a system description and the implementation defines the parameters that must be passed down from network users.



SNA MAIN COMPUTER
no control structure shown


22
JHR
9MAR8

4. In the implementation depicted, the interfaces between the levels are not accessible and may not exist. If an architecture specification goes beyond defining and grouping functions to rigidly define interfaces between groups of functions, the result can be unnecessary constraints on implementation. This can adversely impact cost and production of associated products. The OSI model does not standardize the interfaces between levels.
5. User code as described above layer can be totally provided by a vendor. Examples of this within the IBM product line are NJE, TSO, and VM.
6. This SNA-based interpretation of OSI includes support of a S/370 channel as a data link.

The third implementation is shown in Figure 6 on page 26. This picture suggests a mapping of a 3705 with NCP against OSI and SNA layers. It is very similar to the above picture and many of the same comments apply. Some new ideas elicited by this picture are given below.

1. There is no user code shown because these products are concerned only with data transport and providing the route extension to terminals which support end users.
2. Boundary function is brought into play for support of peripheral nodes that are not capable of providing all the functions necessary to participate in the global network. The programmable resources of the 3705 are used to provide Transmission Control services on behalf of the peripheral node. These are described in the section on Transmission Control. The strengths of this approach are reduced cost for peripheral nodes, increased flexibility, and simplified systems definition.

This concludes the implementation examples. The next section will address different functional interpretations of OSI and SNA.



SNA Communications Controller

23
JHR
9 MAR 81

COMPARISON OF SNA AND OSI

This section will address interpretations of OSI that are implied by current SNA implementations. SNA implementations represent one example for OSI implementations.

As stated in "OSI and SNA" on page 2, the origins and objectives of SNA and OSI are considerably different. The purpose of OSI is to provide a vehicle for ongoing development of standards for system interconnection. The purpose of SNA is to provide a basis for a unified communications product line. This, in turn, will offer the users of that product line improved protection of their investments in end-user systems. To the degree that SNA must address communication between systems--systems in the OSI sense--SNA can be compared with OSI with a high degree of similarity. This paper addresses this similarity.

In focusing on the system-to-system communication aspects of SNA, two important features of the set of SNA products are ignored.

The first is the control structure that is used for management and control of sessions and network resources, e.g., controllers, lines, terminals, application programs, routes, etc. The primary vehicles for this are the architectural entities of systems services control point (SSCP), physical units (PU's) and associated services, logical units (LU's) and associated services, session control, and network control. (17) A similar, underlying control structure does not exist for OSI. This means that many questions about how it really works and who has responsibility for what function cannot be answered except on a case-by-case basis.

The second feature is the increasing set of products devoted to communications network management. Examples of these are the architecture for error reporting, and supporting products such as Threshold Analysis and Remote Access (TARA) for the 3600 Finance Communication System, Network Problem Determination Application (NPDA), and the 386X micro-processor based modems. The existence of an architecture facilitates development of consistent, meaningful network management tools. OSI does not address network management issues associated with system interconnection.

Any complete comparison of OSI and SNA must include these aspects as well as the inter-system similarities discussed in this paper. However, these topics are beyond the scope of this paper. Before continuing with a comparison of the system interconnection similarities of OSI and SNA, the evolutionary progress of each architecture must be understood.

SNA is an architecture driven by a single vendor. It has been visible since 1974. Since then, the implementations of SNA have been expanded and enhanced considerably. As a result of ongoing development and experience gained through implementation, the architecture itself is maturing.

On the other hand, OSI is a young architecture driven by an international standards organization. OSI has only been visible since 1978. The definitions of lower levels of OSI are supportable by, but not limited to, already existing standards work, e.g., CCITT X.25 or ISO HDLC. Architecture for higher levels is not yet advanced enough to permit coherent implementation by the vendor community. There are specific exceptions to this statement, e.g., CCITT X.3, X.28, and X.29 interpretations of OSI level 6, the Presentation layer. Active architectural development is underway in support of other layers. Architecture must be firm and seen to provide significant benefits before vendors can begin to commit resources to product development or enhancement.

Products that purport to conform to an architecture or specification may not be able to communicate with each other. This can result from different interpretations of the specification or mutually exclusive choices made by the product creators. History has shown that time and experience can solve these types of problems when the motivation exists.

For the functions listed earlier for OSI levels 7 and 6, Application and Presentation respectively, current SNA implementations address most in some manner. Exceptions for the Application layer include determination of cost allocation methodology and resource adequacy to provide the desired quality of service. These are OSI objectives to be provided in this layer. No specification of how this should be done is provided. For many of the other topics, the implementation may not be user-specifiable in any way; it may be fixed by the SNA architects or product creators.

The functions described for OSI level 5, the Session layer, are addressed in some manner by current SNA implementations. SNA chaining protocols can accomplish a form of the OSI quarantine service. The services of session establishment and release are controlled by logical unit services of the SNA system services control point (SSCP) acting through session control of the common session control (CSC) manager. The remaining services associated with an established session are controlled by the session control manager of the transmission control element (17).

SNA Transmission Control and Path Control implement most of the functions described in "Open Systems Interconnection" on page 4 for levels 4 and 3, the Transport and Network layers respectively. ACF Release 3 has the ability to select different classes of service. The selection is done from the application

layer at session set-up and is fixed for the duration of the session.

The correlation between an SNA class of service and parameters such as transit time, throughput, availability, etc. is dependent on the resources made available when the network was designed and built by the network owner.

Both SNA and OSI have defined a transport subgroup. In SNA, it is known as the transmission subsystem and includes the functions of Transmission Control, Path Control, Data Link Control, and Physical Interface Control. For OSI, it is known as the transport service and includes OSI levels 4 through 1, the Transport, Network, Data Link, and Physical layers respectively. The SNA Path Control network roughly embraces the total set of functions described for OSI Levels 1, 2, 3, and part of Level 4. With SNA, all end-to-end communication between NAU's is routed by Path Control. In the OSI model, end-to-end functions are only implemented in the Transport Layer, level 4.

The OSI proposal (4) makes the following statement with regard to the Network layer:

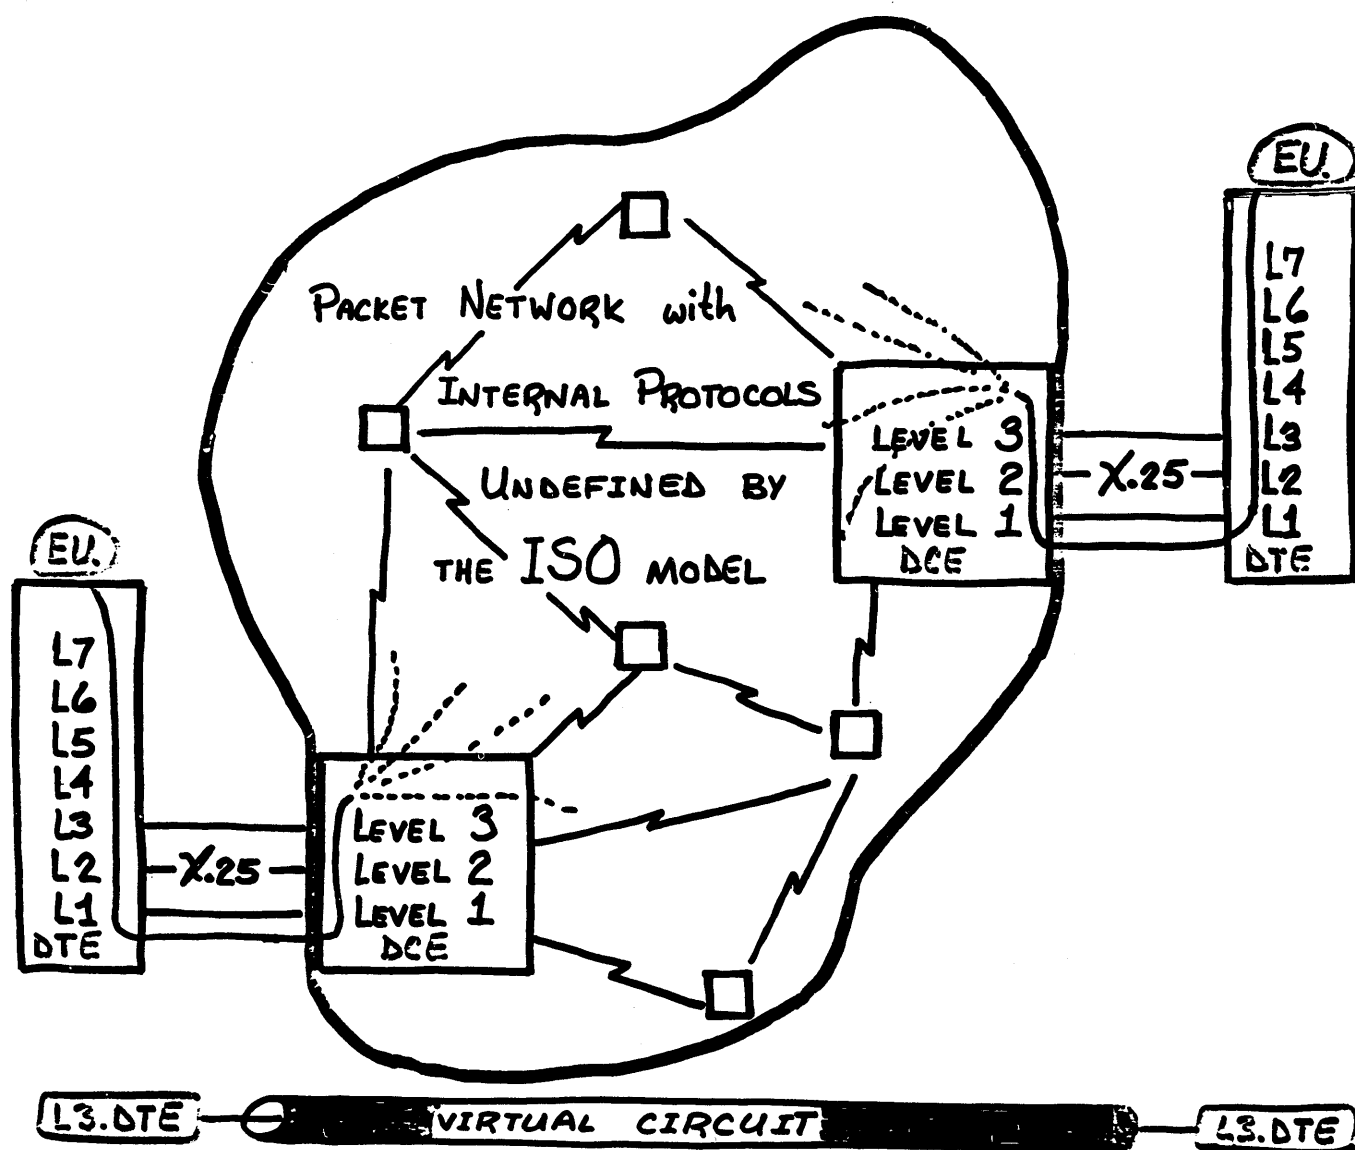
The Network layer contains functions necessary to provide the Transport layer with a firm Network/Transport boundary which is independent of the underlying communications media in all things other than quality of service.

The emphasis on this boundary suggests that it is a candidate for an environmental boundary between dissimilar implementations. This is consistent with current X.25 services. But it does suggest that a service based on the Network/Transport boundary cannot provide a complete transport service; the Transport layer is required to provide the complete service. An example of this would be end-to-end connection service incorporating tandem, dissimilar networks. Services identified with OSI level 4, the Transport layer, are necessary. The burden of providing these is on the user of a service conforming to CCITT X.25 protocols.

However, there are many data communications requirements that can be addressed by a service that provides an OSI level 3 interface. The treatment of the functions of higher levels is defined by the network service available and is deemed adequate by the user.

The OSI Network layer concept of adjacent nodes is illustrated in Figure 7 on page 30. The salient point is that once inside the network with a level 3 interface, the implementations may not be constrained by OSI. This may or may not be significant depending on the need to have access to mechanisms within the packet network.

ADJACENT NODE COMMUNICATION for X.25 INTERPRETATION of ISO MODEL



27
JHR
9/11/81

Figure 8 on page 33 shows a network service that delivers two types of service. This picture depicts an end-user entity executing functions that can be described in two ways:

1. the functions of OSI levels 4, 5, 6, and 7 or,
2. part of the SNA Path Control functions plus the Function Management and NAU Services Manager layer.

At the level of OSI level 4 or SNA Path Control, decisions can be made to use the packet service with the CCITT X.25 interface or to use the SNA-implemented services based on communication controllers interconnected by fixed bandwidth connections. The difference in the two choices is in the function available in looking into the network from OSI level 4 or from SNA Path Control. Both choices represent intelligent networks: the SNA-based services are consistent with all the SNA control structure and all the resources are controllable by the network owner; the transport service based on the CCITT-defined X.25 network is not consistent with the SNA control structure because the resources within the packet net are beyond the control of the network owner.

This may represent an exposure to the user in terms of controlling resource application in response to varying business or external circumstances. On the other hand, the owner of the packet network has accepted responsibility for providing service. This frees the user of that service (the network end user that requests service via OSI level 7 or the SNA NAU Services Manager layer) from concern for how the service is provided.

The functions described for OSI level 2, the Data Link layer, encompass the functions provided by the SNA Data Link Control layer. As described earlier, the data link control vehicles supported under SNA are broader than those currently thought of for OSI implementations. Also, control of transmission groups is part of SNA Path Control rather than a Data Link layer function as OSI suggests.

The functions described for OSI level 1, the Physical Control layer, are the same as those demanded by the SNA Physical Control layer. As described above, the SNA implementations include support for channels as transmission media while OSI does not. OSI does not preclude channel support however.

As stated at the outset, the differences between OSI and SNA begin with their objectives. OSI has the objective of providing a common denominator for interconnection between dissimilar systems. Non-homogeneous systems can be expected to have substantially different internal control structures and management processes. SNA is the foundation for a uniform communications product line. A uniform communications product line implies a consistent control structure and provides the opportunity for standardized network management processes. Products conforming to SNA can have comparatively richer func-

tions for system control, management, and operation than products conforming to the common-denominator requirements of the OSI model.

INTERPRETATION of OSI REFERENCE MODEL with SNA and OTHER PACKET NETWORKS

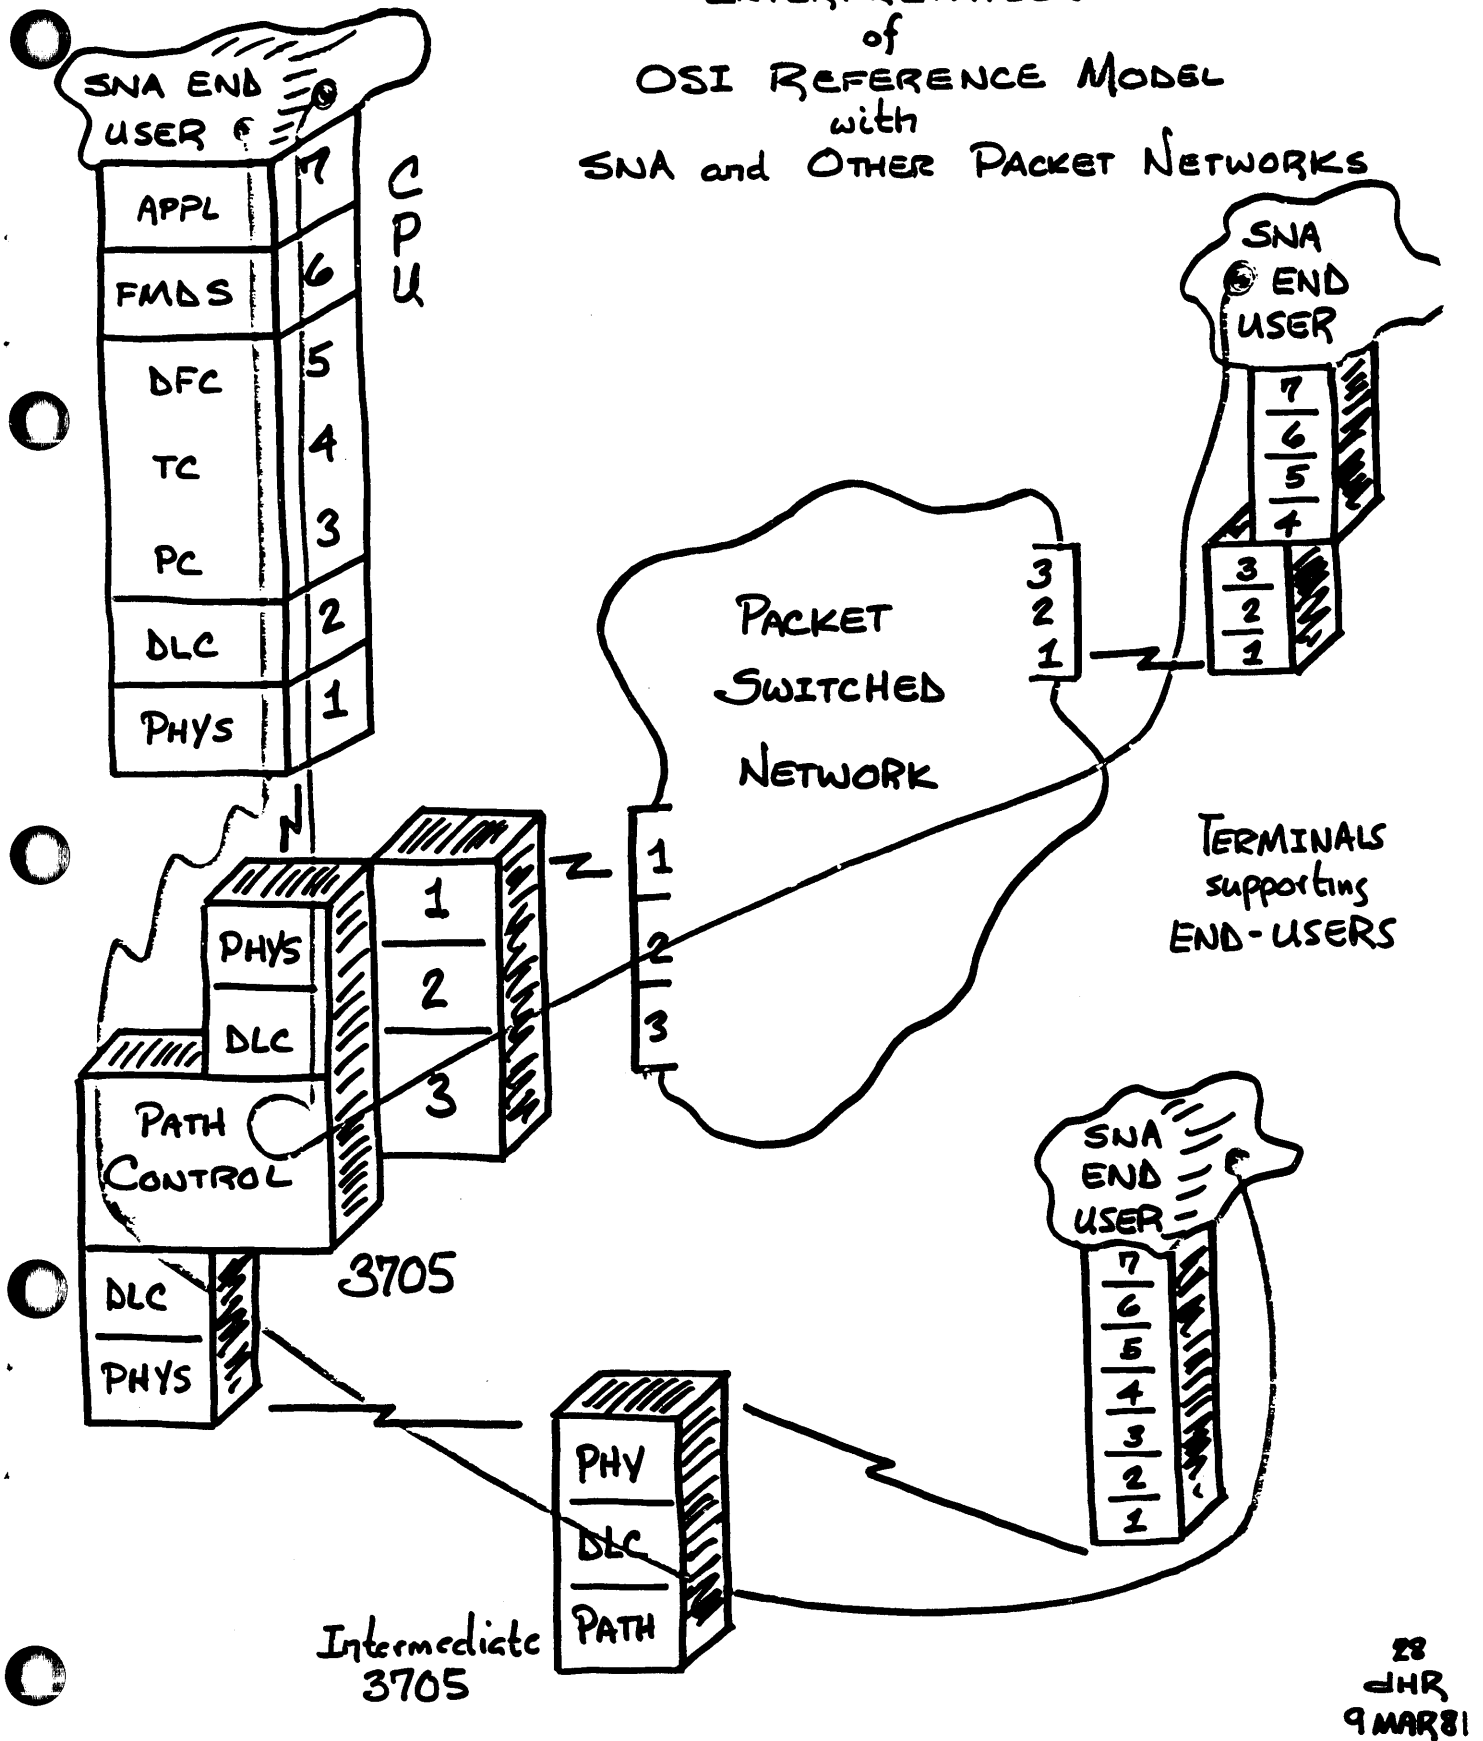


Figure 8

CONCLUSIONS

Predictions supported by meaningful detail are impossible without further OSI architecture and vendor-supplied implementations of that architecture.

In general, a further expansion of options available to address business-related communications issues is a certainty. The degree of adherence to any architecture is a factor of how orderly this expansion may be. The decisions facing a user of data communications offering will be increasingly complex and difficult to execute.

How soon will the potential discipline imposed by OSI be visible? Again, the answer is elusive. But a review of the history packet networks based on the CCITT X.25 recommendation is valuable.

As a proposed standard, CCITT X.25 first became visible in 1973. Network services based on this standard first appeared in 1977. By 1980, CCITT X.25 packet networks have become significant, although not dominant, in the worldwide data communications market. In June, 1979, differences in available implementations of CCITT X.25 packet nets were documented (9). These differences significantly impact products being developed to attach to packet nets. The result has been increased cost and retarded availability of equipment to support packet nets. The community of packet network owners/administrators agreed to improve standardization with a target date of 1982.

From the first visibility of CCITT X.25 to significant standardization of offerings spanned a ten year period. OSI is much broader in scope. As a result, OSI embraces a much larger community of users and vendors. This suggests a lengthy period for the fruition of offerings based on OSI.

What will the IBM role be? IBM can be expected to continue enhancement of SNA and the products that implement it. IBM encourages the use of international standards as the basis for interfaces to public data networks. IBM continues to participate in and contribute to international standards efforts to develop and enhance these interfaces. Additionally, IBM participates in and contributes to the ISO efforts to define, via OSI, a single set of consistent protocols for inter-system communication. Announcement of the capability to support products or functions based on international standards will be based on IBM's technical and business judgment in addressing the requirements of its customers.

APPENDIX - ARCHITECTURE

Consider the following definitions of 'architecture' provided by the dictionary¹:

1. "The art and science of designing and erecting buildings.
2. "A structure or structures collectively.
3. "A style and method of design and construction.
4. "Any design or orderly arrangement perceived by man."

Using the fourth definition, one can assert that any generic process that can be segmented into component processes or methodologies can be addressed by an architecture.

It is widely believed that the data communications process can be segmented into several component processes, usually less than ten. An orderly arrangement of these component processes will constitute a communications architecture.


An initial assumption will be that communications as discussed in this paper will be two-party only. This does not preclude broadcast applications; it simply suggests that a total communications task will be defined as a series of communications between two parties.

As a discussion point, assume that a communications architect has defined five component processes. Together these processes define the complete communications process for one party of a desired two-party communication. This architecture should concern itself with:

1. the logical structure or definition of the component processes,
2. the relationship of each process to the others,
3. the format of data as it passes through the various components, and
4. the peer-to-peer protocols between the component processes.

¹ William Morris, ed., The American Heritage Dictionary of the English Language, 1979, Houghton Mifflin Company, Boston.

PARTY A and the FIVE LAYERS



- The highest level is an end-user port
- The lowest level provides & maintains the physical interface to the transmission medium.
- For LEVEL N , service is requested by LEVEL $N+1$ and provided by LEVEL N using the facilities of lower levels.

6
JHR
7MARS1

Figure 9 depicts the five level interpretation of communications for a single party. Three important notions can be rationalized from this figure.

First is the idea of service represented by a particular level N. Level N provides service to higher levels. Service is requested by level N+1, the next higher level. Level N will honor requests using its own functions and the services provided by the lower levels. There are two special cases of this concept of service requests from higher levels. These are addressed below as the second and third notions.

The second notion is that the highest level, level 5 in this example, will receive service requests from the end user instead of a next higher level. The end user can be manifested in many forms, e.g., a terminal operator, an application program, etc.

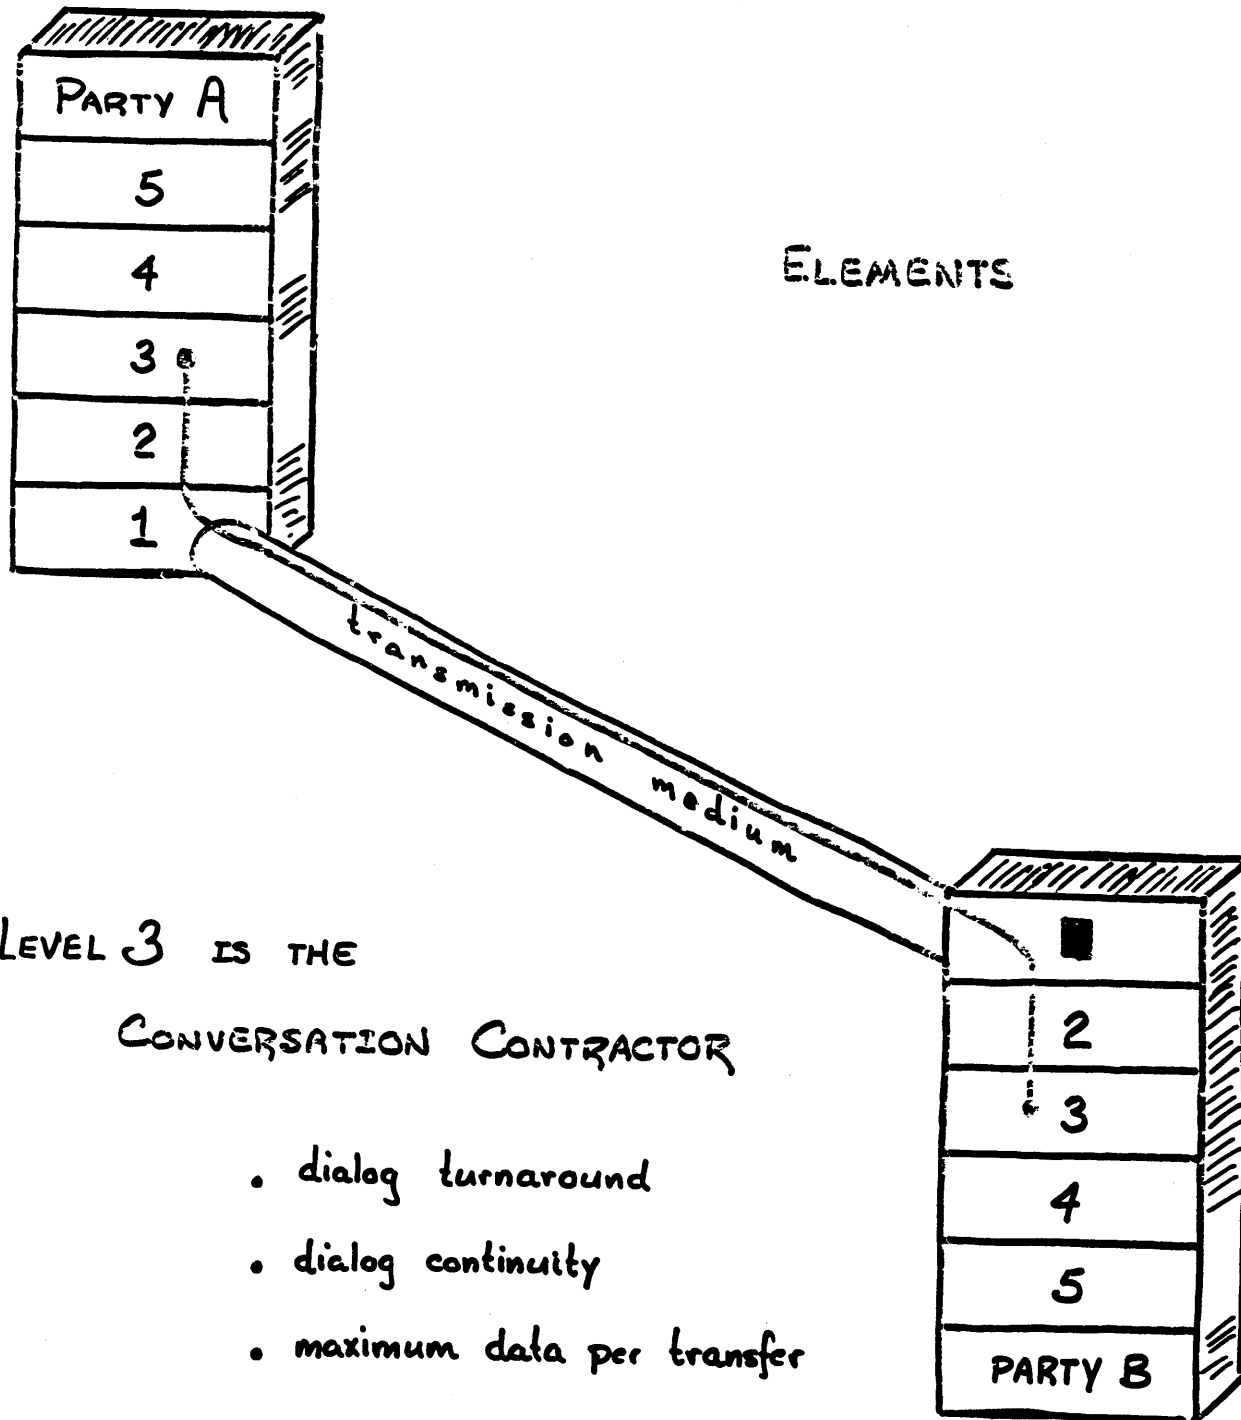
The third notion is that the lowest level interfaces with a physical transmission medium. Within the architecture there are no lower levels from which this level can request service. It does request actual movement of data over the attached transmission medium.

'Peer-to-peer communications' is an essential concept in understanding how these communications architectures are intended to work. See Figure 10 on page 38. Consider Layer 3 to be the conversation contractor, i.e., the control function responsible for agreement between Parties A and B about conditions of dialog turnaround, logical continuity (bracket protocols), and maximum data per transmission. In order to strike an agreement on these parameters, Layer 3 representing Party A may need to communicate with its peer, Layer 3 representing Party B. This communication would be accomplished using the Layer 2 and Layer 1 services available to each Layer 3 entity.

In order for this peer communication to occur, control data associated with a particular component process must be created and transferred. At the sending party, this data will be appended to the data unit received from the next highest level. At the receiving party, used control data will be stripped off as the data unit moves up in levels toward the user party. This process is shown in Figure 11 on page 40.

If an exchange of control information is required between two intermediate levels as described above for Layer 3, then the initial control data unit is created at the originating function. In the above example, this would have been either Layer 3.Party A or Layer 3.Party B depending on who originated the conversation.

PEER-TO-PEER COMMUNICATIONS



7
JHR
7MAR

Figure 10

This pairing of peer functions offers improved opportunities for distribution of function. For example, there could be intelligence at each end of the conversation as described above. Or, if one end of the conversation were implemented with a low function hardware terminal, negotiations on behalf of that terminal could be implemented elsewhere, e.g., fixed parameters or code in a communications controller.

In what form does a component process exist? The component process is an architectural entity, a group of functions associated by decree. This set of functions must be engaged on behalf of any parties desiring to communicate within the rules of the architecture. The medium for providing the functions can be hardware, a program (including microcode), or a human operator. The entity will make decisions on behalf of a desired conversation between two parties.


If code is the medium, executable code will provide the component process. Control blocks associated with each communicating party provide the unique identity of the process with a party or connection.

If the medium is hardware, the component process and its unique identity with an end user are usually locked together.

If there is rigid adherence to formats and protocols specified by the architecture, theoretically communication should be possible between two peer levels implemented by different parties and technologies.

When architectural conformity is desirable, the purveyor of communications products must be concerned with cost alternatives. When the resources to develop and support the products are constrained and there is a choice between providing additional function at the same cost, reducing cost, or adhering to an architecture, there is a strong incentive to either provide more function or reduce cost. The realities of these choices can impede architectural conformity until benefits are clear and functions are not ambiguous.

FORMATS OF DATA & ASSOCIATED CONTROL INFORMATION DURING MOVEMENT FROM B TO A



8
JNR
7 MAR

Figure 11

Even with the cost alternative considerations mentioned above, there are three primary benefits that accrue from the imposition of a data communications architecture. They are:

1. improved flexibility for taking advantage of technological innovations,
2. a reduction of incompatible implementations of the same generic process, and
3. peer-to-peer interface standardization easing interconnections between dissimilar data communication implementations.

REFERENCES AND NOTES

1. Albrecht, Harold R., and Ryder, Keith D. "The Virtual Telecommunications Access Method: A Systems Network Architecture perspective." IBM Systems Journal (1976):53-80.
2. Corr, Frank P., and Neal, Don H. "SNA and emerging international standards." IBM Systems Journal 18 (1979):244-262.
3. Cullum, Phillip G. "The transmission subsystem in Systems Network Architecture." IBM Systems Journal 15 (1976):24-38.
4. Data Processing - Open Systems Interconnection - Basic Reference Model. International Standards Organization Draft Proposal 7498 (December, 1980).
5. Folts, Harold C. "Coming of age: A long-awaited standard for heterogeneous nets." Data Communications, January, 1981: 63-73.
6. Recommendation X.21 - General Purpose Interface between Data Terminal Equipment (DTE) and data Circuit-Terminating Equipment (DCE) for Synchronous Operation on Public Data Networks. CCITT, Geneva (1976).
7. Gray, James P., and McNeill, Tony B. "SNA multiple-system networking." IBM Systems Journal 18 (1979):263-297.
8. Halsey, John R.; Hardy, Leslie E.; and Powning, Leo F. "Public data networks: Their evolution, interfaces, and status." IBM Systems Journal 18 (1979):223-243.
9. Hess, M. L.; Brethes, M.; and Saito, A. "A Comparison of Four X.25 Public Network Interfaces." In ICC '79 Conference Record, pp. 38.6.1-38.6.8. New York: The Institute of Electrical and Electronics Engineers, Inc., 1979.
10. IBM Implementation of X.21 Interface General Information Manual. No. GA27-3287. IBM Corporation, Data Processing Division, White Plains, New York.
11. IBM Synchronous Data Link Control General Information. No. GA27-3093. IBM Corporation, Data Processing Division, White Plains, New York.
12. McFayden, John H. "Systems Network Architecture: An overview." IBM Systems Journal 15 (1976):4-23.
13. Provisional Recommendations X.25 on Packet-Switched Data Transmission Services. CCITT, Geneva (1978).

14. Schay, Peter S. "'Letter to the Editor: Not fair to compare' in 'Newsfront' column." Data Communications, October, 1980: 53-56.
15. Systems Network Architecture Format and Protocol Reference Manual: Architectural Logic. No. SC30-3112-2. IBM Corporation, Data Processing Division, White Plains, New York.
16. There are three types of NAU: the logical unit (LU), the physical unit (PU), and the system services control point (SSCP). The PU and SSCP are associated with network control and are beyond the scope of this paper for discussion. The logical unit (LU) is the access port into the network for the end user or communicating party.
17. The existence of the terms used in these paragraphs is an indication of the depth of the current SNA implementations. However, a feeling for what they constitute is necessary to maintain a sense of reality. The terms are:
 - common session control (CSC),
 - function management data services (FMDS),
 - logical unit (LU) (16),
 - logical unit services and logical unit services manager,
 - network addressable unit (NAU) (16),
 - network control,
 - physical unit (PU),
 - physical unit services and physical unit services manager,
 - system services control point (SSCP) (16), and
 - session control,

These terms denote architectural entities; groups of function that logically associated. These entities are defined implicitly in the SNA Format and Protocols Manual (15). In actual implementation, the protocols and functions are executed by code, and sometimes hardware, that may or may not be structured the same as the architecture. This means there may not be code modules with names corresponding to these terms. One motivation for this may be economy of development. Architectural purity of structure and function is sometimes sacrificed.



READER'S COMMENT FORM

Title: OSI and SNA: A Perspective
Washington Systems Center
Technical Bulletin GG22-9225

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Please state your occupation: _____

Comments:

Please mail to: James H. Rutledge
Washington Systems Center
18100 Frederick Pike
Gaithersburg, MD 20760

O

O

O

O

O

GG22-9225-00

OSI And SNA: A Perspective Printed in U.S.A. GG22-9225-00

IBM[®]

GG22-9225-0

