A STUDY OF
HIGHER LEVEL PROTOCOLS
FOR
OPEN SYSTEM INTERCONNECTION
- FINAL REPORT

B.E. Greenleaf
P.J. Power
B.A. Wootton

Government of Canada / Gouvernement du Canada

MEMORANDUM — NOTE DE SERVICE

'O. Monkewich/hlj

TO / A: Distribution

FROM / DE: Y. F. Lum, DCN

SECURITY - CLASSIFICATION - DE SÉCURITÉ:

OUR FILE/NOTRE RÉFÉRENCE: 9722-28

YOUR FILE/VOTRE RÉFÉRENCE:

DATE: September 10, 1980

SUBJECT / OBJET: Report on a study of higher level protocols for Open System Interconnection

Attached is a copy of the Systemhouse Ltd. report entitled "A Study of Higher Level Protocols for Open System Interconnection". The work was carried out for the Department of Communications of the Government of Canada under the general framework of computer/communication research.

The overall objective of the study was to gain insight into the problem of developing unified Canadian standards in higher level protocols.

The study focused on the high level protocols used in Canadian computer communications industry and their relationship to the Open System Interconnection concept. It examined the ISO Reference Model of Open System Interconnection and compared it to several proprietary architectures used in Canada. Furthermore, it surveyed Canadian industry to ascertain current and future requirements, potential solutions and general disposition towards architectures and international standards.

I would like to point out that this study is of exploratory nature and is not to be construed to reflect in any way DOC's position or thinking.

3-64160

Y. F. Lum

YFL/hlj

Encl.

A STUDY OF HIGHER LEVEL PROTOCOLS FOR OPEN SYSTEM INTERCONNECTION - FINAL REPORT

B.E. Greenleaf
P.J. Power
B.A. Wootton

TABLE OF CONTENTS

TABLE OF CONTENTS

LIST OF FIGURES

NUMBER	TITLE
2.4.1.1.1-1	Association of Transport Entities (ISO)
2.4.2.1.2-1	Half Sessions, Virtual Routes, and an Explicit Route (SNA)
2.4.2.1.3-1	Segmenting and Blocking (SNA)
2.4.2.1.3-2	Potential for Shortened Delay in Intermediate Node, With Segmenting (SNA)
2.4.2.1.4-1	End to End Flow Control (SNA)

SUMMARY

Whereas large organizations had been growing the information systems environment in terms of individual systems, often using diverse hardware, the trend towards large integrated intra-company systems has forced an interconnection of previously unrelated computer systems into homogeneous networks. As inter-organization (as opposed to intra-company) communications becomes more widespread, a similar requirement for interconnection of dissimilar systems occurs.

This market pull has not been satisfied by the push of technology from the systems manufacturers. Rather, the major manufacturers appear to be developing their own specialized but proprietary communications network architectures. IBM's System Network Architecture (SNA) is perhaps the best known, largely due to that company's dominance in the marketplace. By itself, SNA is a very powerful architecture, but makes no provision for communication amongst non-IBM systems.

There has been considerable progress in the development of vendor independent communications standards, of which the most widely known is the X.25 standard for a network access protocol. X.25, however, provides only the transport subsystem and there is no provision for application related protocols.

In order to address the problems of communications at higher functional levels, standards have been under development at the international level for several years. Perhaps the best known is the ISO Reference Model of Open Systems Interconnection, as described in publication ISO/TC97/SC16 N 227.

Although highly conceptual in nature, this model provides the framework for manufacturer independent networks which would allow communication amongst users, applications, and systems tasks.

To aid in the further development and refinement of this model, Canada was asked to provide input regarding the current and projected requirements for interconnection amongst various vendor's equipment. The study, performed for the Department of Communications, under DSS contract No. OSV79-00048, was one vehicle to determine such Canadian Industry requirements.

The Study of Higher Level Protocols for Open System Interconnection was carried out in two phases. Phase I of the study involved the study of the ISO and CCITT reference models for Open Systems Interconnection and three proprietary architectures; specifically, IBM's SNA, DEC's Digital Network Architecture, and Hewlett-Packard's Distributed System Network. The Phase I report presented a definitive discussion of architectures and protocols, and a layer by layer functional description of all considered architectures. A comparative analysis of reference models was also presented.

Phase II of the research involved an extensive survey of Canadian Industry. The survey was intended to acomplish a threefold objective:

- to ascertain the commitment to manufacturer provided architectures and the effect of non-standardization on forward systems planning.

- to ascertain the projected development of custom architectures/protocols to accomplish dissimilar system interconnection.

- to ascertain the Canadian requirements for distributed systems.

The Phase II report first presents the survey methodology, analysis, results, and conclusions. In particular, the Canadian Industry requirements are highlighted.

Next, the functionalities proposed for the various layers of the ISO model, are compared and contrasted with equivalent functionalities provided by SNA and DECNET.

The last section of the Phase II report presents the mapping of Canadian Industry requirements to a functional space. These "required functionalities" are then discussed in terms of being realized by the ISO model.

The report indiciates that further OSI work involves four areas of activity:

- definitional work related to the basic building blocks and interim solutions;

- definitional work related to the ultimate OSI architecture and its protocols;

- implementation of the initial building blocks and interim solutions;

- implementation of a full OSI capability;

Each of these areas is discussed in detail in the Phase II report.

In order to set the stage for future OSI development some immediate concerns must be addressed by DOC. One such concern is the interconnection of similar networks via a Public Data Network (PDN). As discussed in the report, the SNA/PDN interface problem must be resolved as a precursor to OSI evolution. It is suggested that DOC have an active involvement in committees, such as SHARE (IBM's user group), devoted to the resolution of this situation. It is especially important that the Canadian industry requirements be adequately represented to IBM.

Also, from the national perspective, DOC should actively undertake OSI prototype design and development. This will not only promote the Canadian communications industry, as being at the leading edge of state-of-the-art communications architectures, but provide immediate benefits to Canadian industry at large.

From an international perspective, DOC should actively participate in the international standards organizations, as a representative of Canadian industry requirements. Prototype development, because of the associated detailed analysis and design involved, would be valuable to DOC in completely realizing the Canadian industry needs.

0.0 PROJECT INTRODUCTION

The development of standards for the interconnection of computer system components began as soon as these systems became products as compared to one-off research prototypes. The evolution of standards has paralleled the evolution of computer systems themselves, typically lagging the state-of-the-art, since need must always precede solution. The current state of computer systems sees the long awaited development of corporate wide information systems, as well as the beginning of information systems technology to replace conventional modes of communication amongst firms.

Whereas large organizations had been growing the information systems environment in terms of individual systems, often using diverse hardware, the trend towards large integrated company-wide systems, has forced an interconnection of previously unrelated computer systems into homogeneous networks. As communications between organizations (rather than within an organization) becomes more widespread, a similar requirement for interconnection of dissimilar systems occurs.

This market pull has not been satisfied by the push of technology from the systems manufacturers. Rather, the major manufacturers appear to be developing their own specialized but proprietary communications network architectures. IBM's System Network Architecture (SNA) is perhaps the best known, largely due to that company's dominance in the marketplace. By itself, SNA is a very powerful architecture, but makes no provision for communication amongst non-IBM systems, except through the normal trend of many manufacturers to produce IBM-compatible products.

There has been considerable success in the development of vendor independent communications standards, of which the most widely known is the X.25 standard for access to a public packet network. However X.25 does not solve the problem of intersystem communication, since it addresses only the areas of physical interface, link protocol (for network access), and transport or routing of messages to the correct network destination. There is no provision for end-to-end protocol or for specifying the intent of the data (i.e., command, data base update, or inquiry) when it reaches its destination. As an aside, the practical use of X.25 has been somewhat disappointing, perhaps due to the cost and complexity of developing individual interface capabilities to the various computers and terminals in the marketplace.

In order to address the problems of intersystem communications at higher functional levels, standards have been under development at the international level for several years. Perhaps the best known is the ISO Reference Model of Open Systems Interconnection, as descibed in publication ISO/TC97/SC16 N 227. Although highly conceptual in nature, this model provides the framework for manufacturer independent networks which would allow communication amongst users, applications and systems tasks.

To aid in the further development and refinement of this model, Canada was asked to provide input regarding the current and projected requirements for interconnection amongst various vendor's equipment. This study, performed for the Dept. of Communications, under DSS Contract #OSV79-00048, is one vehicle to determine such requirements.

The work will study the Reference Model, compare that to typical proprietary architectures, most notably IBM's SNA, and most important by surveying Canadian industry to ascertain current and future requirements, current and proposed solutions and general disposition towards architectures, whether proprietary or international standards.

0.1 Phase II Executive Summary

Phase I of the Study of Higher Level Protocols for Open Systems Interconnection involved the study of the ISO and CCITT reference models for Open Systems Interconnection and three proprietary architectures; specifically, IBM's SNA, DEC's Digital Network Architecture, and Hewlett Packard's Distributed System Network. The Phase I report presented a definitive discussion of architectures and protocols, and a layer by layer functional description of all considered architectures. A comparative analysis of reference models was also presented.

Phase II of the research involved an extensive survey of Canadian industry. The survey was intended to accomplish a threefold objective:

(i) to ascertain the commitment to manufacturer provided architectures and the effect of non-standardization on forward systems planning.

(ii) to ascertain the projected development of custom architectures/protocols to accomplish dissimilar system interconnection.

(iii) to ascertain the Canadian requirements for distributed systems.

Section 1 of the Phase II report presents the survey methodology, analysis, results, and conclusions. In particular the Canadian Industry requirements are highlighted.

Section 2 compares and contrasts the functionalities proposed for the Presentation, Session, and Transport layers of the ISO Model, with equivalent functionalities provided by SNA and DECNET.

Each sub-section (one per layer) is prefaced by an executive summary which highlights the various architectural congruencies and incongruencies.

Section 3 presents the mapping of requirements discussed in section 1 to the functional space outlined in section 2. The "required functionalities" are then discussed in terms of being realized at a particular layer of the ISO Model.

Recommendations with respect to the Department of Communications' involvement in future OSI activity and areas for further study are presented in the next section.

0.2 Recommendations

0.2.1 Future of Open Systems Interconnection

Future work with respect to the definitions and implementations of OSI models must be carefully considered. Although the long-term gain is theoretically high, and the current interest and requirement by Canadian Industry is apparently strong, there are many complex issues of a technical, logistical, and policy nature that must be addressed.

Future work in OSI must cover four types of activity. Along one axis there is work that should address the ultimate long-term objectives of OSI in its purest form, but at the same time, address the more urgent aspects of the architectural definition first. This will to provide the means for "building-block" definitions, on which the full capability will be based, but which will also provide some immediate benefits and results, and alleviate the problem of dissimilar system interconnection. Work can also be sub-divided along the other axis as definitional (i.e. definition of functional sets) and implementational (i.e. implementation of prototypes, pilots, and products).

This results in four areas of activity, being the combination of two sub-divisions along two degrees of freedom:

- definitional work related to the basic building blocks and interim solutions;

- defintional work related to the ultimate OSI architecture and its protocols;

- implementation of the initial building blocks and interim solutions;

- implementation of full OSI capability.

These are considered in more detail in the following sections.

0.2.1.1 Definitional Work Related to Basic Building Blocks and Interim Solutions

There are a number of areas that warrant immediate attention in terms of providing building blocks for future OSI implementations and which could provide immediate benefits. Most of these areas have been identified by various organizations but have not had sufficient attention paid to their development and implementation.

One such area is the SNA/Public Data Network (PDN) compatability and interface. Although several possible solutions have been developed or defined, these solutions are typically restrictive in nature, do not provide full virtual circuit capabilities (as defined by PDNs), nor are fully integrated within the functional capability of SNA.

In our survey of the private sector of the Canadian marketplace, this matter is of very great concern and one which requires immediate attention if these organizations are to further their forward systems planning, and implement their more immediate requirements. It is anticipated that the public sector, especially the federal government, would have similar if not corresponding requirements and concerns.

Section 1.3.1 gives a more detailed presentation of the SNA/PDN problems, solutions, and benefits.

0.2.1.2 Definitional Work Related to the Ultimate OSI Architecture and its Protocols

Prior to the ultimate realization of an OSI architecture much work is required in the definition of peer-to-peer protocols, and boundary (layer interface) protocols. From the Canadian perspective, a building block approach to the realization of layer protocols would be the definition and specification of functional subsets, representative of the Canadian industry requirements, for the Application, Presentation, Session, and Transport layers, followed by prototype protocol development based on these functional subsets.

Much of this work will be centred around studies of compatibility between SNA and OSI, including investigation of the feasibility of a detailed OSI specification that adequately conforms to SNA. Similar specification must be done with regard to other proprietary architectures, such as Digital's DNA, and finally must consider as a subset interface, the compatibility with non-architectural systems such as a link level interface to a processor performing a dedicated function.

0.2.1.3 Implementation of the Initial Building Blocks and Interim Solutions

The Canadian environment for information processing is quite unique in that most industry represents a branch plant type of organization, and although heavily dependent on communications, few Canadian organizations are of sufficient size to warrant extensive investment in computer technology that may be needed to implement their business applications. The result of the branch plant syndrome, is a set of unique product characteristics. Canadian representatives of equipment manufacturers may well recognize product characteristics that are necessary in the Canadian marketplace and yet not have funds or authority to undertake development. Development is done only at such time as the parent corporation perceives a world-wide business interest. The ramifications on Canadian users of information processing technology, being relatively small in comparison to our U.S. counterparts, is that they are typically forced into a compromised implementation of their business requirements because they by themselves cannot afford to implement their own, for example, communications architecture, but instead must adhere to that which is readily available and are often constrained in the requirement accordingly. Both of the situations contribute to the continuing dominance of information processing technology by the United States and European organizations, even though, on an individual basis, Canadian industry has demonstrated its capability to be at least comparable.

The X.25 standard for PDN interfacing is a good example of the situation. Although Canada has taken a lead in terms of X.25-based networks, thus creating a requirement within the marketplace for interfacable DTE's, most of the development work in this regard was slow in coming, and was ultimately delivered, for the most part, from locations outside Canada. As time goes on, and new ventures become increasingly more complex and costly, this technology will continue if the private industry organization is faced with the entire cost of research and development. In order to implement even the basic building blocks of an OSI architecture, the investment required will certainly exceed that which was required to implement an X.25 interface. Even on a contract basis, Canadian companies, although competent, will have a difficult time establishing sufficient credibility with foreign manufacturers unless they are able to first develop a prototype solution as a firm indication of not only their capability but leading edge expertise in the area.

There are a number of basic building blocks required for a future implementation of an OSI architecture and they may be viewed in two general categories; namely the human gateway and the machine gateway.

The machine gateway refers to the OSI interface to existing proprietary architectures, most notably SNA and DNA. Considerable work is required here, but an attempt has already been made to determine congruencies between architectures. They are presented in section 2 of this report.

At the human gateway is a requirement for a basic virtual terminal capability that will ultimately provide a full OSI interface functionality. As a basic building block for that capability is required a terminal sub-system that provides interface to a PDN and yet is expandable to accomodate future functional layers as they are defined and implemented. If the higher level protocols were not an objective, there would be little need for implementation of the PDN access logic in the terminal since a NIM (network interface machine) concept would be adequate for most situations. As a foundation for later implementation of the higher level functional capabilities, however, one must first implement a user work station capability that will support the first three layers (X.25). This also provides an interim solution to a number of requirements for terminal compatibility, although the full benefit of a virtual terminal service would not be realized until higher level protocols were implemented at both source and destination.

0.2.1.4 Implementation of full OSI Capability

Implementation of higher level protocols, as a prototype implementation, will accelerate the OSI standardization process, and benefit the Canadian industry forward systems planning. Canadian industry is unique and fortunate to have a number of suitable test beds that would provide useful and practical application of higher level protocols. Although there is doubtless a long list of possible candidates, three major application areas are apparent:

- third party access from Telidon;

- the Canadian Trade Information Systems Network from Costpro;

- various departments of the federal government as users of information processing technology.

0.2.2 DOC Participation in OSI

In order to set the stage for future OSI development some immediate concerns must be addressed by the Department of Communications. One such concern is the interconnection of similar networks via a PDN. As discussed in section 1.3.1 the SNA/PDN interface problem must be resolved as a precursor to OSI evolution. It is suggested that the Department of Communications have an active involvement in committees, such as SHARE, devoted to the resolution of this situation. It is especially important that the Canadian industry requirements be adequately represented to IBM.

Also from the Canadian perspective, DOC should actively undertake OSI prototype design and development. This will not only promote the Canadian communications industry, as being at the leading edge of state-of-the-art communications architectures, but provide immediate benefits to Canadian industry at large. From an international perspective, the Department of Communications should actively participate in the international standards organizations, as a representative of Canadian industry requirements. Prototype development, because of the associated detailed analysis and design involved, would be valuable to DOC in completely realizing the Canadian industry needs.

0.2.3 Areas for Further Study (Systemhouse Participation)

Systemhouse's participation in furthering Open Systems Interconnection in Canada can be categorized as:

- representation of Canadian Industry;
- research and development of an OSI prototype.

These are considered in more detail in the following sections.

0.2.3.1 Representation of Canadian Industry

Although a special sub-committee of SHARE has been formed to investigate the SNA/PDN interface requirements, and to make a recommendation to IBM, there are two concerns. First, the chairman of this sub-committee has recently left to pursue other duties, and secondly, because of far-reaching ramifications of such an activity, that it should be carefully monitored, and ideally, actively contributed to by a representative of the Department of Communications. Systemhouse can offer this representation because of its expertise in the area, and because of past participation in this sub-committee.

Systemhouse can also offer representation on the international standards committees, not only because of resident expertise in the area of international standards, but because Systemhouse has acquired a unique understanding and identification with Canadian industry's distributed data processing requirements.

0.2.3.2 Research and Development of an OSI Prototype

Systemhouse can provide the research and development of an OSI prototype along two parallel paths.

One path involves the design and development of the "human" and "machine" gateways discussed in section 0.2.1.3. The "human" gateway would at first accomodate a PDN interface and would be upward expandable to a full OSI implementation.

The "machine" gateway would provide an architectural mapping function between proprietary architectures and the OSI architecture. It should be noted that both gateways may be realizable within the same physical entity.

The second path involves the definition of functional subsets for each of the Application, Presentation, Session, and Transport layers based on Canadian requirements. From these functional subsets, protocol prototypes can be designed, developed, and implemented.

In conclusion, the importance of the work on Open Systems Interconnection must be stressed. Many individuals and organizations have devoted considerable amounts of time and effort to further this work. The work is as much an effort in human communications as in computer communications. The key to success is in understanding the needs and concerns of the various organizations and structuring the solution to meet these needs. Systemhouse understands Canadian industry needs. Systemhouse can provide the solution.

1.0 SURVEY OF CANADIAN INDUSTRY

1.1 Introduction

As part of the investigation into the problem of developing unified Canadian standards for higher level protocols, Systemhouse Ltd. carried out a survey of various private sector Canadian companies. The objectives of the survey were to ascertain the commitment to manufacturer architectures, the development of custom network architectures and protocols, and to identify Canadian requirements for distributed systems.

This section describes how the survey was undertaken to achieve these objectives. It details the results of a statistical analysis of the responses, and explains how the Canadian requirements were deduced from this analysis. The work being carried out in private industry on open system interconnection, as revealed by the survey, along with details of the higher level protocols currently in use and being developed, are also covered.

The section also contains a report of related topics covered at the SHARE 54 Conference held in California.

1.2 Survey Techniques Used

The survey was carried out using high volume distribution of a survey questionnaire and low volume interviews with those companies thought to be most active in this area. A relevant conference was also attended. This method was adopted because of economic and time constraints.

1.2.1 Survey Questionnaire

Appendix A contains a copy of the covering letter and survey questionnaire that were distributed. Approximately 300 of these questionnaires were distributed to companies across Canada. Systemhouse Ltd. contacted several organizations to arrange distribution to their members. The Canadian Manufacturing Association sent out 100 copies to those member companies involved with telecommunications and distributed processing, the Canadian Industry Communications Assembly sent out 50 copies to similar member companies and the Canadian Bankers' Association sent copies to the major banks.

Systemhouse Ltd. distributed 156 copies to companies across Canada. The criteria used were that the company should have more than one branch, use more than one manufacturer's equipment, and have more than one computer. In each case the survey was sent to the company head office (in many cases the branches spanned more than one province). The distribution by province was as follows.

Alberta	14
British Columbia	14
Manitoba	6
New Brunswick	6
Newfoundland	1
North West Territories	0
Nova Scotia	6
Ontario	77
Prince Edward Island	0
Quebec	28
Saskatchewan	4
Yukon	0
TOTAL	156

1.2.2 Survey Interviews

Interviews were carried out with organizations covering many different types of operations and which were considered those most likely to be involved with networking, distributed systems and higher level protocols. Sixteen interviews were held and all companies had branches in more than one province. The breakdown by province was as follows:

Alberta	1
British Columbia	1
Manitoba	0
New Brunswick	0
Newfoundland	0
North West Territories	0
Nova Scotia	0
Ontario	8
Prince Edward Island	0
Quebec	5
Saskatchewan	1
Yukon	0
TOTAL	16

1.2.3 Share 54 Conference

Systemhouse Ltd. sent a representative to the Share 54 conference at Anaheim, California to report on any discussions relevant to this study.

1.3 Survey Analysis

This section contains details of the information obtained through the survey, including:

- a report describing the Share 54 conference with details of relevant topics discussed;
- details of industry response to the survey with general analysis;
- identified industry requirements for distributed systems and communications;
- industry involvement in open systems interconnection and networking;
- industry involvement in the design of higher level protocols.

1.3.1 SHARE 54 Conference

SHARE's primary goal has been to foster joint research and development and provide a forum for the exchange of ideas pertaining to computer science. SHARE provides IBM with documented user requirements for developing new products, and for modifying and enhancing present hardware and software systems. Through SHARE procedures for establishing such requirements and evaluating their relative costs and merits, IBM receives a coherent, unified presentation of real user needs.

The real work of SHARE Inc. is done by the various task forces, projects and their committees. Each project provides a forum for the exploration of critical issues and the participation of assigned IBM representatives opens a unique communication window.

Systemhouse participated in the SHARE 54 SNA/X.25 committee because of a Canadian Industry requirement to interface SNA and Packet Data Networks (PDN). IBM has addressed the PDN/SNA interface support in the IBM Systems Journal which has resulted in favourable responses from both IBM users and PDN carriers. IBM has also stated the need for a technically adequate user input on PDN/SNA interface support, which has created the need to intensify the PDN/SNA activity within SHARE.

1.3.1.1 PDN Interface Standards

The virtual circuit (VC) concept, and the store-and-forward data transmission and control functions of the PDNs have necessitated a new family of "X" CCITT recommendations (i.e. X.3, X.25, X.28, X.29, X.75, X.121, etc.).

The PDN interfaces impose considerable architectural implications on the SNA networking, since they are not fully transparent to the SNA end-to-end protocols. From an IBM user's point of view, the present level of maturity of the SNA concept and the "X" interfaces has created a realistic expectation that a "universal" PDN/SNA interface support will be provided by IBM in the not too distant future.

1.3.1.1.1 CCITT Recommendation X.25

X.25 is the primary device independent interface between a PDN and a DTE (Data Terminal Equipment) operating in the packet-mode. The X.25 interface consists of three functional levels: Physical Control (X.21, V24, or RS232); Link Control (LAP or LAPB subset of HDLC; LAPB is preferred by IBM and CCITT); Network Control (provides the procedure for data and control information exchange between a DTE and a PDN using the VC concept). X.25 is proposed as being the lower three levels of the ISO OSI reference model. Considerable pressure has been put on IBM to support this model by the public carriers and special interest groups (e.g. SHARE SNA/X.25 committee).

1.3.1.1.2 CCITT Recommendations X.3, X.28, X.29

X.3, X.28, and X.29 specify the requirements for attachment of start-stop character mode terminals to PDNs. X.3 specifies the packet assembly/disassembly (PAD) functions and parameters, which include assembly of characters received from a start-stop terminal into packets to be transmitted over a PDN and disassembly of packets received from a PDN into characters to be transmitted to a start-stop terminal. X.28 provides specifications for the PAD-to-terminal protocol, which consists of service initialization, data transmission, and control information exchange. X.29 specifies the end-to-end procedures for data transmission and control information exchange between a PAD supporting a start-stop terminal and a remote X.25-compatible host, controller or terminal which also must be X.29 compatible. The SNA/X.29 host interface support would make the SNA migration more encompassing by allowing migration of a user's existing start-stop terminals to an SNA host over a PDN.

1.3.1.2 PDN/SNA Interface Benefits

The Canadian industry requirement of a PDN/SNA interface is based on the following associated benefits:

(i) Capabilty to implement internationally-based SNA networks, including the countries in which PDNs may provide the only commercially available data transmission facilities (e.g. European countries).

(ii) Lower data communication costs due to higher resource utilization, resulting from sharing of available bandwidth, among multiple users, and user charges which are usage sensitive.

(iii) Wide geographical accessibility to a PDN due to the distributed PDN nodal communications processing, which makes PDNs suitable for implementing distributed data processing (DDP) SNA systems.

(iv) Increases in a SNA system's growth flexibility since additional SNA nodes require only new access links to PDN's (minimal engineering change to the existing SNA network).

(v) Adaptability to a user's SNA network traffic growth without a need to re-engineer the data transmission subsystem, since higher traffic volumes are accommodated by a PDN's nodal communications processing capacity, with the only cost to a user being the charge for a data traffic increment.

(vi) Low starting data transmission costs, which is an attractive proposition for a smaller SNA user who cannot afford to implement a transmission subsystem using leased lines (i.e. smaller SNA single-domain, 8100 peer-to-peer, or 4331 peer-to-peer networks).

(vii) Use of the X.25 virtual circuit services, which can complement SNA capabilities in specialized applications.

(viii) Increased data transmission reliability due to a PDN's error checking, data flow control, dynamic routing, etc.

(ix) Improved overall SNA network management due to a PDN's SNA-complementary network management features.

(x) A PDN's architectural and interface compatability with the OSI concept, since X.25 can be implemented as the lower three levels of OSI which makes PDNs desirable in a user's long term SNA network planning.

1.3.1.3 PDN/SNA Interfacing

The SNA/X.25 interface possibilities depend on the type of the communicating DTEs. First, if both DTEs are SNA-based (e.g. SNA Host and SNA cluster), the necessary SNA/X.25 accommodations must be provided at both ends. Second, if one DTE is SNA-based (e.g. VTAM-driven host), and the other DTE is X.25 based, the SNA/X.25 accommodation must be provided in the SNA-based DTE. In this case, the communicating DTEs do not use the SNA end-to-end protocol. Third, if one DTE is SNA-based (e.g. VTAM-driven host), and the other DTE is an ASCII start-stop terminal, the SNA-based DTE must accommodate X.25 and X.29. The PDN serving node must accommodate X.3, X.28 and X.29. Figure 1.3.1.3-1 illustrates PDN/SNA interface possibilities.

Fig. 1.3.1.3-1 PDN/SNA INTERFACE POSSIBILITIES

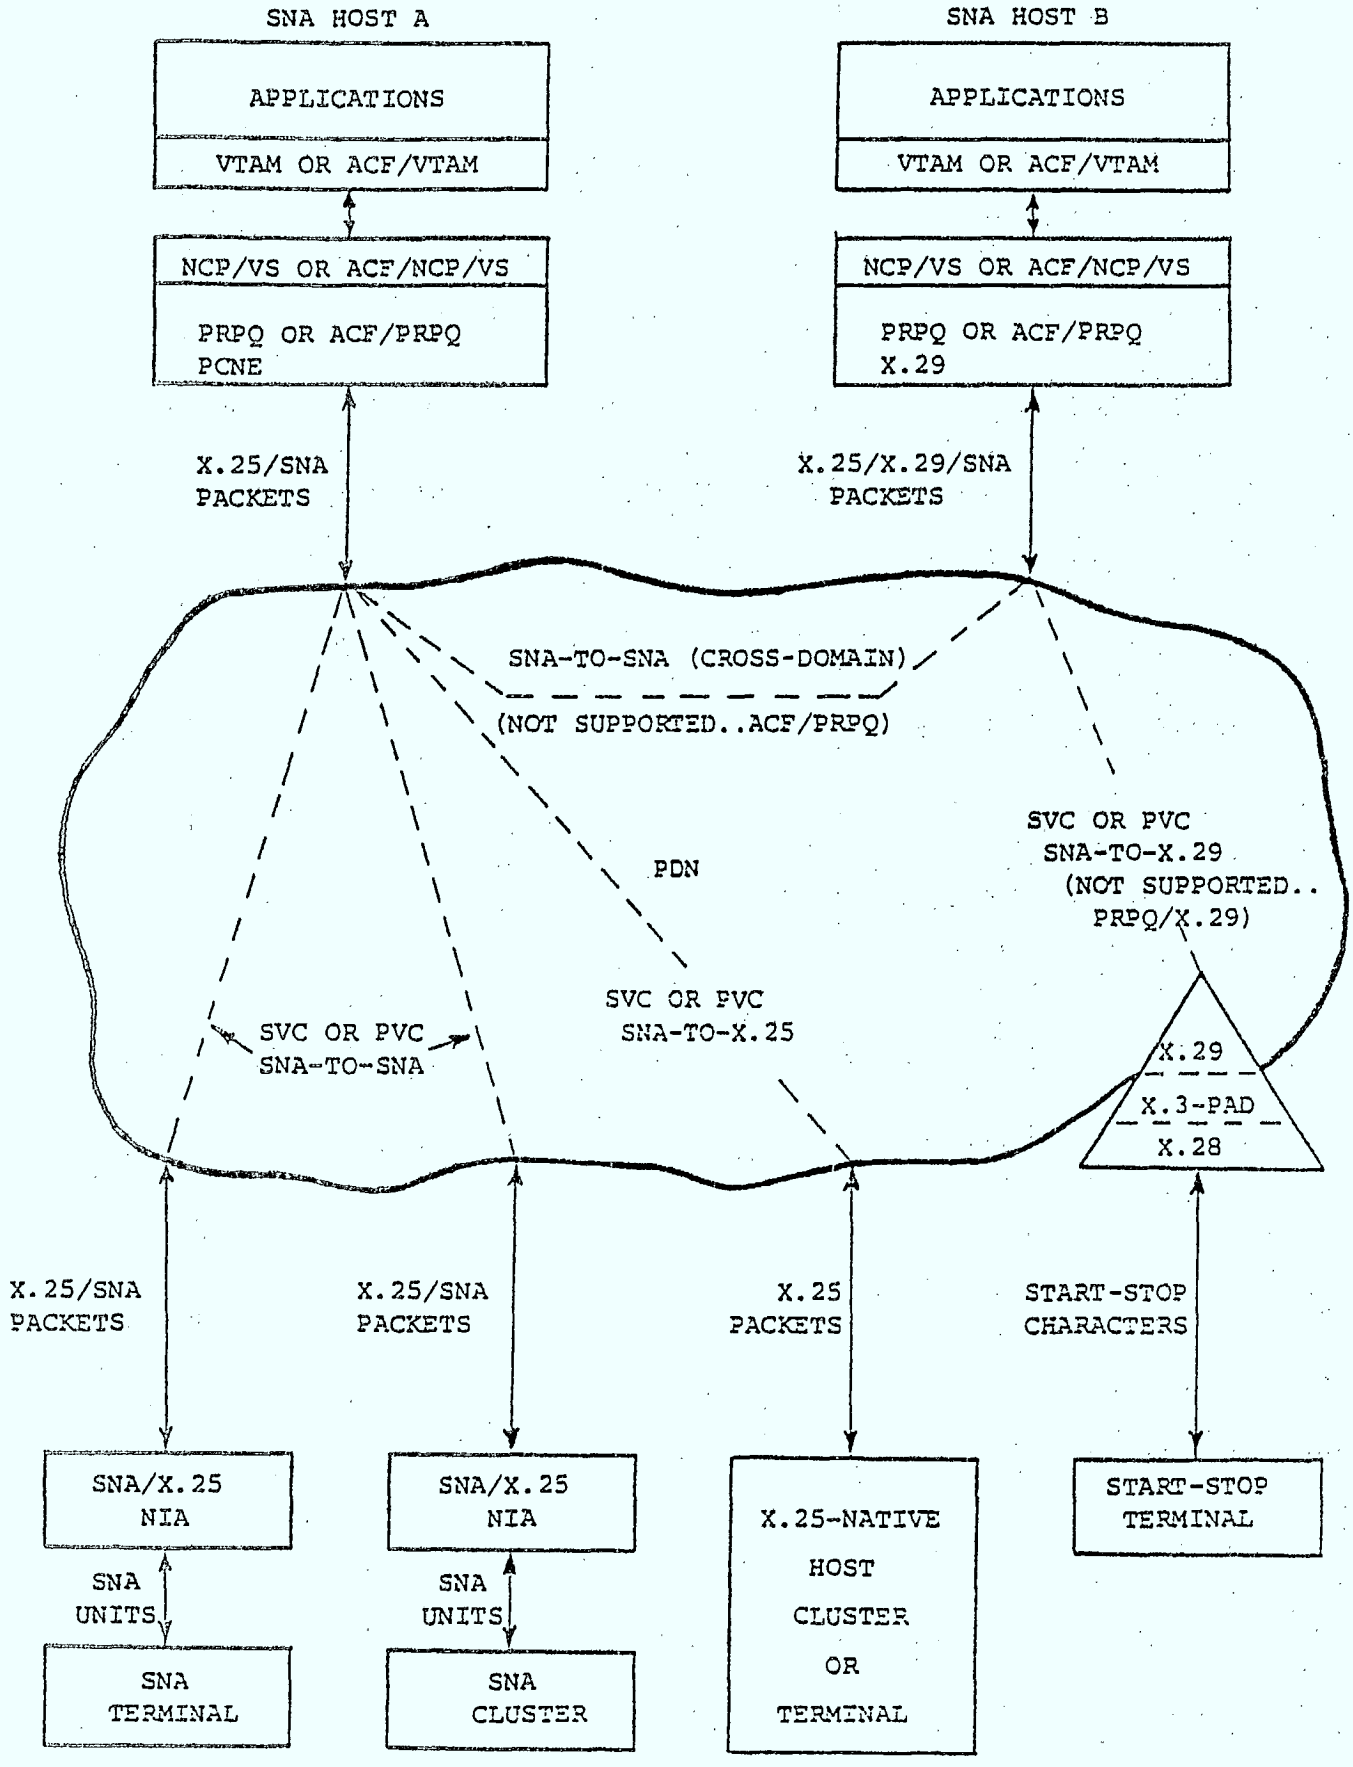

1.3.1.3.1 Current SNA/X.25 Solutions

(a) SNA-to-SNA Solution

The existing solution, offered in Canada and France only, consists of the PRPQ ZA4239 software modification of the NCP5 in the IBM 3705-II, and the IBM 5973-L02 NIA (Network Interface Adapter). The PRPQ/NIA solution allows an SNA cluster/terminal to communicate with an SNA host over a PDN using a link level (level 2) SNA/X.25 mapping. An X.25 virtual circuit appears as a private line replacement to the communicating SNA DTEs. The SNA/X.25 conversions are not visible to VTAM, or to the SNA cluster/terminal. This means that the X.25 VC services are not used by the communicating SNA DTEs.

(b) SNA-to-X.25 Solution

The PCNE (Protocol Conversion for Native Equipment) feature of the PRPQ allows an SNA host to communicate with an X.25-based (native) DTE which appears as an IBM 3767. This X.25 accommodation is provided in the NCP and is not visible to VTAM. It is possible, however, to define an X.25 DTE via VTAM macros.

1.3.1.3.2 PDN/SNA Requirements

(i) The provision of SNA/X.25 support for all existing PDNs (not only Canadian Datapac and French Transpac), in order to satisfy the international interconnectability requirements in SNA networks.

(ii) Upgrading the existing SNA/X.25 interface solution(PRPQ/NIA) to allow use of PDNs in multi-domain SNA networks.

(iii) Providing an integrated X.25 interface support (adapter) in the SNA cluster controllers and DDP products in order to eliminate the need for additional hardware (separate adapter) in the SNA host-to-SNA cluster and peer-to-peer communications of a PDN.

(iv) Implementing an SNA/X.25 interface support for the SNA host-to-X.25 native equipment communications at the terminal, cluster controller, communications controller, and host levels. The existing PRPQ-PCNE feature is implemented only at the SNA terminal level (IBM 3767 emulation).

(v) Providing an SNA/X.25 interface support for the SNA-to-SNA communications at the communication controller and host level (PU4 and PU5), allowing multiple virtual circuits between multiple PUs/LUs (e.g. SNA host-to-SNA host cross-domain, or 8100-to-8100 or 4331-to-4331 peer communications).

(vi) Including the X.29 protocol in the SNA/X.25 host interface support in order to allow use of PDNs in the SNA migration of start-stop character mode terminals. The SNA-to-X.25/X.29 interface requirement is considered to be an important SNA migration step, due to the vast population of ASCII terminals.

(vii) Providing the SNA/X.25 three-level mapping in accordance with the OSI/X.25 mapping, with the necessary SNA level boundary functions changes to take full advantage of the X.25 virtual circuit services.

1.3.2 Response to Survey

A total of 23 responses to the survey questionnaire were received in time to be included in this report. With the 16 interviews held, this provides a sample of 39 organizations on which analysis was performed.

The distribution of the sample, by province, was as follows:

Province	Questionnaire	Interview	Total
Alberta	2	1	3
British Columbia	2	1	3
Manitoba	1	0	1
New Brunswick	0	0	0
Newfoundland	0	0	0
North West Territories	0	0	0
Nova Scotia	0	0	0
Ontario	9	8	17
Prince Edward Island	0	0	0
Quebec	1	5	6
Saskatchewan	1	1	2
Yukon	0	0	0
Not Known	7	0	7
	23	16	39

All analyses in the following sections deal with the sample sub-categorized by organization type. The organization type distribution within the sample is shown below:

Type of Organization	Questionnaire[1]	Interview	Total
Banks	5	1	6
Other Financial	4	0	4
Computer Services	3	2	5
Manufacturing	5	3	8
Retail	1	2	3
Natural Resources	2	4	6
Transportation	0	2	2
Communications	0	2	2
Educational	2	0	2
Other	1	0	1
	23	16	39

Those questionnaire replies that did not include an organization type, and which, by and large, had no OSI requirement, were excluded from the following analyses.

1.3.3 Requirements

This section details the distributed data processing and communications requirements identified from the responses to the survey questionnaire and from the survey interviews. The section is sub-divided by organization type.

The identified requirements are categorized as follows:
- extent of communications;
- data traffic types;
- end user services;
- associated requirements.

1.3.3.1 Banks

Questionnaire responses -	4
Interviews -	1
TOTAL	5

1.3.3.1.1 Extent of Communications

Both international and intercontinental open system interconnection between banks are required for applications such as on-line banking and credit card verification. There is the requirement to support large networks with many terminals and processors.

1.3.3.1.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications, which involves the transmission of large volumes of data at high speeds;
- terminal to processor communications, which involves the transmission of low volumes of data requiring a fast response time;
- automatic teller machine support, which involves low volume intermittent but secure communications;
- transaction transmission capabilities;
- low priority bulk data transmission, including remote printing capabilities.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.1.3 End User Services

All types of data transfer are required including:

- inquiry/update file access;
- print file and bulk data transmission;
- remote job entry.

These imply the requirement for a file transfer service and a job transfer service.

Many terminal types are required including automatic teller terminals and credit card verification terminals (to be installed at all locations where credit cards may be used). This implies the requirement for a virtual terminal service.

1.3.3.1.4 Associated Requirements

The following additional requirements are associated with the banking applications:

- availability and reliability, of the network is critical for some of the required communications;

- error rate, of the communications service must be low;

- security, of transmission is an important factor in some cases;

- sequential delivery, of data is an additional requirement for some applications;

- assured delivery, must be provided for some applications.

1.3.3.2 Other Financial Organizations

Questionnaire responses -	4
Interviews -	0
TOTAL	4

Of the four replies received, only one of the companies did not have the requirement to connect multiple systems and was therefore satisfied by the current manufacturer offerings.

1.3.3.2.1 Extent of Communications

Both international and in one case intercontinental, including satellite, communications are required. There is a requirement to interface telex and teletype terminals to computer data communications networks. An additional requirement is to use the same communications link for both voice and data communications.

1.3.3.2.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications, to allow data communications with other companies;

- terminal to processor communications, to allow file inquiry/update, on line programming, and low usage terminal transactions;

- terminal to terminal communications, which require low speed transmission for teletype and facsimile terminals;

- transaction transmission capabilities;

- low priority bulk data transmission, including remote printing capabilities.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.2.3 End User Services

The end user services required include all types of data transfer

- data entry/data collection;
- print file and bulk data transmission;
- inquiry/update file access;
- on-line programming;
- remote job entry;
- resource sharing and remote device control.

These imply the requirement for a file transfer service and a job transfer and manipulation service.

Other end user services required include many terminal applications:

- data entry/data collection support;
- message switching;
- terminal independance for applications;
- provision for low usage terminal services;
- facsimile terminals for documents, graphs etc.
- teletype terminals.

These imply the requirement for a virtual terminal service.

1.3.3.2.4 Associated Requirements

No additional requirements were found to be associated with the other financial organization applications.

1.3.3.3 Computer Services Organizations

Questionnaire responses -	3
Interviews -	2
TOTAL	5

1.3.3.3.1 Extent of Communications

A major concern is open system interconnection between Canada and the United States. Communication with Europe is a lesser concern of the computer services organizations.

All companies involved would prefer to use one master architecture although this is a high priority only in some cases. Communications must be available between users and branch offices, via a Canadian or U.S. public data network. There is a requirement for internetworking. Satellite communications is an economic requirement of most organizations.

1.3.3.3.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications, to allow data communications with other companies;

- terminal to processor communications, to include the capability for low speed terminals;

- terminal to terminal communications, to include telex;

- transaction transmission, task to task communication;

- low priority bulk data transmission, including remote printing;

- digitized voice transmission.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.3.3 End User Services

The end user services required include the following data transfer services:

- print file and bulk data transmission;
- inquiry/update file access;
- remote job entry;
- remote device control.

These imply the requirement for a file transfer service and a job transfer and manipulation service.

Terminal services required include:

- interactive processing;
- support for multiple terminal types;
- message switching;
- facsimile terminal support;
- word processor support;
- graphics support.

These imply the requirement for a virtual terminal service. Other services required include:

- provision of network services to outside users;
- network management functions.

1.3.3.3.4 Associated Requirements

The following additional requirements are associated with the computer services applications:

- accountability, of all network and device services used by any network user is an additional requirement associated with a computer services company operation;

- support, for existing terminals, protocols, etc. must be provided by any implementation of standards;

- compatibility, with the developments in the United States must be provided by any implementation of standards;

- open border, must exist for data traffic between Canada and the United States in either direction;

- competition, by Canada in U.S. markets should be supported by standards.

1.3.3.4 Manufacturing

Questionnaire responses -	5
Interviews -	3
TOTAL	8

Of the eight companies involved only one did not have the requirement to connect multiple systems. Some of the others felt that the current manufacturer offerings would meet their requirements, if they used one master architecture and made other manufacturer equipment emulate the protocols.

1.3.3.4.1 Extent of Communications

The majority of companies require communications in Canada, or between Canada and the United States. One company also requires communications with Europe. Internetworking is an additional requirement.

1.3.3.4.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications;
- terminal to processor communications, to support both interactive and low speed terminals;
- transaction transmission, task to task communications;
- low priority bulk data transmission, including remote printing capabilities.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.4.3 End User Services

The end user services required include the following data transfer services:

- data entry/data collection;
- print file and bulk data transmission;
- inquiry/update file access;
- remote job entry.

These imply the requirement for a file transfer service and a job transfer service. Terminal services required include:

- support for multiple terminal types;
- data entry/data collection support;
- interactive processing;
- message switching.

These imply the requirement for a virtual terminal service.

1.3.3.4.4 Associated Requirements

The following additional requirements are associated with the manufacturing organization applications:

- internetwork communications, to allow company networks to interface with service bureau networks must be provided by any standards;

- interfacing, to process control computers was identified as an additional requirement for one company's applications;

- availability, and resilience of the network (23 hours a day) was identified as an additional requirement for one company's applications;

- standardization of terminals, for one company, in the business of manufacturing automatic teller terminals, would remove many restrictions imposed by the current non-standardization of terminals.

1.3.3.5 Retail Organizations

Questionnaire responses -	1
Interviews -	2
TOTAL	3

One of the companies is not currently concerned with networking.

1.3.3.5.1 Extent of Communications

All of the companies had requirements for communications within Canada only.

1.3.3.5.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications;
- terminal to processor communications, to include support for point-of-sale terminals;
- transaction transmission;
- low priority bulk data transmission, including remote printing.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.5.3 End User Services

The user services required include the following data transfer services:

- print file and bulk data transmission;
- data entry/data collection;
- inquiry/update file processing;
- remote job entry.

These imply the requirement for a file transfer service and a job transfer service. Terminal services required include:

- data entry/data collection support;
- interactive processing;
- message switching.

These imply the requirement for a virtual terminal service.

1.3.3.5.4 Associated Requirements

No additional requirements were found to be associated with the retail organization applications.

1.3.3.6 Natural Resources Organizations

Questionnaire responses -	2
Interviews -	4
TOTAL	6

Only one company had no requirement to interconnect systems.

1.3.3.6.1 Extent of Communications

Both international and intercontinental communications are required. The ability to interconnect multiple networks is also required. One company is developing its own private value added network for data communications, that will support voice communications in the future.

1.3.3.6.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications;
- terminal to processor communications;
- terminal to terminal communications;
- transaction transmission;
- low priority bulk data transmission, including remote printing.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.6.3 End User Services

The end user services required include the following data transfer services:

- print file and bulk data transmission;
- inquiry/update file access;
- data entry/data collection;
- remote job entry.

These imply the requirement for a file transfer service and a job transfer service. Terminal services required include:

- data entry/data collection support;
- interactive processing;
- word processor support;
- graphics support.

These imply the requirement for a virtual terminal service. Other services required include:

- electronic mail;
- provision of network services to outside users;
- network management functions.

1.3.3.6.4 Associated Requirements

The following additional requirements are associated with the natural resources organization applications:

- high priority, must be selectable for some transactions;
- assurance, of delivery or confirmation of delivery must be provided for certain transactions by any implemented standards;
- compatability, of systems to allow internetworking must be provided by any standards.

1.3.3.7 Transportation Organizations

Questionnaire responses -	0
Interviews -	2
TOTAL	2

1.3.3.7.1 Extent of Communications

Both international and intercontinental communications are required. The use of one master architecture would be the preferred way to go.

1.3.3.7.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications;
- terminal to processor communications;
- transaction transmission, task to task communication;
- bulk data transmission and remote printing.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.7.3 End User Services

The end user services required include the following data transfer services:

- data entry/data collection;
- inquiry/update file access;
- print file and bulk data transmission;
- remote job entry.

These imply the requirement for a file transfer service and a job transfer service. Terminal services required include:

- data entry/data collection support;
- message switching;
- word processing support;
- facsimile terminal support.

These imply the requirement for a virtual terminal service. Other services required include:

- electronic mail.

1.3.3.7.4 Associated Requirements

The following additional requirements are associated with the transportation organization applicatons:

- fast response, will be required for some short transactions;

- end to end sequencing, of data must be provided;

- assurance, of delivery or confirmation of delivery must be provided for some transactions;

- application device independance, to allow implementaton of an application on dissimilar devices must be provided by any implementation of standards;

- internetwork communications, must be provided for some applications;

- public databases, must be accessible using any implemented standards.

1.3.3.8 Communications Organizations

Questionnaire responses -	0
Interviews -	2
TOTAL	2

1.3.3.8.1 Extent of Communications

Data communications is required nationally, and internationally between Canada and the United States. Internetworking is also a requirement. Any attempts to introduce standards will be supported by the communications organizations.

1.3.3.8.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications;
- terminal to processor communications;
- terminal to terminal communications;
- transaction transmission, task to task communications;
- bulk data transmission and remote printing.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.8.3 End User Services

The end user services required include the following data transfer services:

- print file and bulk data transmission;
- inquiry/update file processing;
- remote job entry;
- remote device control.

These imply the requirement for a file transfer service and a job transfer and manipulation service. Terminal services required include:

- interactive processing;
- support for low speed terminals;
- support for other specialized terminal services provided by communications companies.

These imply the requirement for a virtual terminal service.

1.3.3.8.4 Associated Requirements

The following additional requirements are associated with the communications organization applications:

- internetwork communications, must be provided to support the facility to communicate between networks and with the networks of other companies. This includes a requirement to communicate with government department networks;

- telephone switching, implies an additional specific requirement to connect telephone switching machines to minicomputers via a public packet switched network. These machines produce large volumes of data irregularly and the requirement may be compared to that of the connection of process control computers.

1.3.3.9 Educational Establishments

Questionnaire responses - 2
Interviews - 0

TOTAL 2

1.3.3.9.1 Extent of Communications

The majority of data communications will be within Canada although there is a minor requirement for international communications.

1.3.3.9.2 Data Traffic Types

The data traffic types required include:

- open processor to processor communications;
- terminal to processor communications;
- bulk data transmission.

These data traffic types must eventually be supported by both private and public networks.

1.3.3.9.3 End User Services

The end user services required include the following data transfer services:

- data entry/data collection;
- inquiry/update file access;
- print file and bulk data transmission;
- remote job entry.

These imply the requirement for a file transfer service and a job entry service. The terminal requirements include:

- data entry/data collection support;
- interactive processing;
- support for multiple terminal types.

These imply the requirement for a virtual terminal service.

1.3.3.9.4 Associated Requirements

The following additional requirements are associated with the educational establishment applications:

- terminal support, must include support for both low speed terminals and microprocessor-based intelligent terminals.

1.3.4 OSI Involvement

This section details the involvement of the various organization types in open systems interconnection, as identified from the responses to the survey questionnaire and from the survey interviews.

1.3.4.1 Banks

Questionnaire responses -	4
Interviews -	1
TOTAL	5

The banking organizations are involved in open systems interconnection between similar systems, i.e. single manufacturer networks. In the majority of cases these networks use SNA, or are migrating to SNA. There is some use of DNA for smaller system networks separate from the main systems. The attitude towards future involvement in open systems interconnection seems to be that some form of standards are required. The best way to co-ordinate the implementation of standards would be through the Canadian Bankers' Association.

1.3.4.2 Other Financial Organizations

Questionnaire responses -	4
Interviews -	0
TOTAL	4

Some of the other financial organizations are involved with open systems interconnection using SNA. The attitude towards future involvement in open systems interconnection is that it will not be possible to improve the existing networks unless a standard interface between SNA and the public packet networks is developed.

1.3.4.3 Computer Services Organizations

Questionnaire responses -	3
Interviews -	2
TOTAL	5

The computer services organizations are involved in open systems interconnection between similar systems, i.e. single manufacturer networks. The architectures in use are mainly customized IBM, or migrating to SNA. Some DEC services are provided on separate networks. The attitude towards future involvement in open systems interconnection seems to be to follow manufacturer offerings and customize according to individual requirements.

1.3.4.4 Manufacturing

Questionnaire responses -	5
Interviews -	3
TOTAL	8

The manufacturing organizations are involved to varying degrees in open systems interconnection. The majority are between similar systems, i.e. single manufacturer networks. Some connection between dissimilar systems is done using lower level protocols. The attitude towards future involvement in open systems interconnection seems to be to extend the use of manufacturer offerings.

1.3.4.5 Retail Organizations

Questionnaire responses -	1
Interviews -	2
TOTAL	3

The retail organizations have some involvement with open systems interconnection between dissimilar systems currently using lower level protocols. The attitude towards future involvement in open systems interconnection seems to be to wait for new standards.

1.3.4.6 Natural Resources Organizations

Questionnaire responses -	2
Interviews -	4
TOTAL	6

The natural resources organizations are involved with open systems interconnection between similar systems using, or migrating to, existing manufacture architectures. There is some involvement with interconnection between dissimilar systems currently using a magnetic tape interface. The attitude towards future involvement in open systems interconnection seems to be to wait for standards to allow open interconnection between manufacturer products.

1.3.4.7 Transportation Organizations

Questionnaire responses -	0
Interviews -	2
TOTAL	2

The transportation organizations are involved with open systems interconnection to connect dissimilar systems, currently using lower level protocols. The attitude towards future involvement in open system interconnection seems to be that a standard architecture should be used, with manufacturer supported standardized protocols at all but the Application Layer.

1.3.4.8 Communications Organizations

Questionnaire responses -	0
Interviews -	2
TOTAL	2

The communications organizations are involved in open systems interconnection between dissimilar systems using customized higher level, or lower level manufacturer protocols. The attitude towards future involvement in open systems interconnection seems to be to actively support implementation of new standards.

1.3.4.9 Educational Establishments

Questionnaire responses -	2
Interviews -	0
TOTAL	2

The educational establishments have some involvement in open systems interconnection between dissimilar systems currently using lower level protocols. The attitude towards future involvement in open systems interconnection seems to be to follow any new implementation of standards if it is supported by manufacturer products.

1.3.5 Higher Level Protocols

This section details the involvement of the various organization types in the design and implementation of higher level protocols, as identified from the responses to the survey questionnaire and from the survey interviews.

The banks, other financial organizations, manufacturing organizations, retail organizations, transportation organizations and educational establishments surveyed are not currently involved in the design or development of higher level protocols.

1.3.5.1 Computer Services Organizations

Questionnaire responses -	3
Interviews -	2
TOTAL	5

Some computer services organizations are involved with the design and development of higher level protocols, but this involvement is either in the customization of manufacturer protocols, or in producing implementation specific protocols.

1.3.5.2 Natural Resources Organizations

Questionnaire responses -	2
Interviews -	4
TOTAL	6

The majority of the natural resources organizations are not currently involved in the design or development of higher level protocols. There is some involvement in the design and development of application specific protocols.

1.3.5.3 Communication Organizations

Questionnaire responses -	0
Interviews -	2
TOTAL	2

The communications organizations are involved in the design and development of higher level protocols both for specific implementations and as input to the design of prototype protocols for international standards.

Particular examples include:

- a higher level protocol has been developed by Bell Canada to interface telephone switching machines, via DATAPAC, with minicomputers. The requirement is similar to that for interfacing process control computers.

- prototype protocols have been designed by Bell Northern Research for the Session and Transport Layers.

1.4 Conclusions

This section summarizes the survey analyses detailed in section 1.3. It provides a general overview of Canadian industry, categorized as follows:

- Canadian industry requirements, a summary of the requirements identified in Section 1.3.3;

- Canadian industry involvement in open systems interconnection, a summary of the analyses detailed in Section 1.3.4;

- Canadian industry involvement in higher level protocols, a summary of the analyses detailed in section 1.3.5.

1.4.1 Canadian Industry Requirements

The requirements identified are categorized as follows:

- extent of communications;
- data traffic types;
- end user services;
- associated requirements.

These are considered in more detail in the following sections.

1.4.1.1 Extent of Communications

The majority of companies have requirements for data communications with the United States. Many of those companies are interested in ensuring that no restrictions are imposed on data communications between the two countries. Some companies have the requirement to communicate with Europe, but the survey indicated that this requirement is of a lower priority than communication with the U.S. Some companies feel that the introduction of a satellite data service would improve data communications with the United States and should be supported by any future standards.

Many of the companies involved in the survey use telex services and some feel that it would be useful to provide some means of connecting telex services to data communications networks. There is also a minor requirement, which will probably expand in the future, to provide a hybrid voice and data communications network, possibly using digitized voice transmission techniques (e.g. packetized voice).

A small proportion of the companies surveyed, notably the banks, have the requirement to support very large networks. The majority of the companies have an internetworking requirement. This includes the interconnection of tandem private networks, the connection of private networks across public networks, the connection of tandem public networks, of a combination of these.

1.4.1.2 Data Traffic Types

The most common types of data transmission required are processor to processor and terminal to processor transmission. They will support transaction and bulk data transmission.

Communication between dissimilar processors is required both within a network (i.e. intranetwork) and between networks (i.e. internetwork) to allow inter-company, inter-branch, etc. communication. This type of communication typically requires high-speed exchange of large volumes of data.

Communications between terminals and processors is required both within private networks and across tandem networks to allow a network user to communicate with a remote processor. Many different types of terminals with varying data characteristics will be used. Typically this type of data traffic is low volume and low speed, but requires a fast response.

Transaction transmission is an intranetwork and internetwork requirement to allow a transaction to be sent from a terminal or processor to a remote processor. This type of data service is characterized by intermittent but high volume traffic. Some transactions may have particular constraints imposed by the application (i.e. high level of security, delivery confirmation, etc.).

Bulk data transmission is required both intra and inter-network to allow for data file transmission between two processors or between a terminal and a processor. Bulk data transmission is characterized by the intermittent flow of large volumes of low priority data. Therefore, bulk data transmission imposes no particular constraints on network resources and can share available bandwidth (as opposed to requiring dedicated bandwidth).

Terminal to terminal data transmission is also a communications requirement, and will allow special communications such as telex, facsimile etc. The data transmission service required will normally be low speed, although voice communication, an example of terminal to terminal communications, requires a relatively high bandwidth.

All data traffic types must eventually be supported by both private and public communications networks. This implies that networks must be "open" to each other for data exchange and therefore must offer similar services, that is, implement an open systems interconnection standard.

1.4.1.3 End User Services

The most common end user services required are those mentioned within the ISO Model:

- virtual terminal service;

- virtual file service;

- job transfer and manipulation service.

The details of the requirements for particular facilities within these end user services are contained within the following sections.

1.4.1.3.1 Virtual Terminal Service

The major requirements of the virtual terminal service are support for multiple terminal types, provision of interactive processing and support for message switching.

Support is required for multiple terminal types including asynchronous (e.g. teletype) facsimile, etc. This must allow the connection of similar terminals from different manufacturers.

Interactive processing must be supported to allow terminal users to interact with applications on remote processors via single or tandem networks.

Message switching must be supported to allow messages to be sent from one end user to another via single or tandem data communications networks.

Other virtual terminal service requirements are support for word processing systems, data entry systems, and specialized graphics systems.

1.4.1.3.2 Virtual File Service

Most companies have identified a requirement for a virtual file service of some description. The major requirements are inquiry/update file access, bulk data transmission and data entry/data collection.

Inquiry/update file access must be supported to allow a terminal user or an application on a local processor to extract information from, or add information to, a remote data file across single or tandem networks.

Bulk data transmission must be supported to allow a terminal user or an application on a local processor to transmit a data file to another terminal or processor across single or tandem networks. The file may be a downline system load, a print file being spooled or transmitted directly to a printer, or any data file.

Data entry/collection must be supported to allow information gathered at one point to be transmitted to a remote data store across single or tandem networks. The data may be stored temporarily at the entry point and transmitted, as a data file, at intervals to the remote data store using bulk data transmission; or transmitted on entry using update file access.

Other less common requirements associated with the virtual file service include support for on-line programming to allow a terminal user to access and process a data file containing program code. The file may be held locally and processed remotely using bulk data transmission, or held and processed on a remote machine using inquiry/update file access.

1.4.1.3.3 Job Transfer and Manipulation Service

The major requirement associated with the job transfer and manipulation service is to support remote job entry. This will allow a terminal user or application process to transmit a batch job across single or tandem networks for entry to the batch job queue or the remote processor. This is a special case of file transfer.

A subsidiary requirement associated with the job transfer and manipulation service is the facility to share resources and control devices at the remote processor. This is required for more complicated remote job processing applications.

1.4.1.3.4 Other End User Services

Other end user services identified by the survey as requirements include:

- electronic mail, to allow a terminal user to send information to a remote location;

- the provision of network services to outside users, to allow a terminal user access to information and resources within a value added systems network, both for processing and networking;

- the provision of network management functions, to allow network controllers visibility to the network and to make network control simple and, where possible, non-disruptive, i.e. dynamic.

1.4.1.4 Associated Requirements

The most common additional requirements identified as being associated with various applications were:

- internetworking;

- assured delivery;

- sequential delivery;

- availability and reliability;

- ability to interface with process control type computers.

These are considered in more detail in the following sections.

1.4.1.4.1 Internetworking

The ability to connect networks of different companies, or different networks of the same company, across single or tandem public networks was identified by many companies as an additional requirement associated with their particular applications. The differing reasons included the ability to connect a company network to a service bureau, or to connect a company network to a government network, or to connect dissimilar remote networks of the same company.

1.4.1.4.2 Assured Delivery

Some applications were identified as requiring assured delivery of all information sent. This assurance must be provided end to end across single or tandem public or private networks.

1.4.1.4.3 Sequential Delivery

Most applications required sequential transmission across single or tandem public or private networks, and in some cases the requirement included sequential delivery of data from multiple sources.

1.4.1.4.4 Availability and Reliability

The ability of a network to be available to its users for long periods of time and to be reliable enough to withstand node/route failures etc. is a requirement identified by some companies.

1.4.1.4.5 Ability to Interface with Process Control Type Computers

The ability to interface process control type computers with data communications networks is a requirement identified by some companies.

1.4.1.4.6 Other Associated Requirements

The other additional requirements identified in the survey as being associated with specific applications include the following:

- the ability of future networks to support existing terminals and protocols;

- the ability to support both low speed terminals and microprocessor controlled terminals;

- provide device independence for applications to permit hardware changes without software changes;

- standardization of automatic teller terminals to remove the current restriction of emulating the de facto IBM standard;

- provision of a single or tandem network service with a low end to end error rate to reduce retransmissions;

- ability to access public databases from private networks via single or tandem public networks;

- ability to provide high security for certain types of data in particular applications;

- ability to select a high priority for certain types of data in particular applications;

- ability to provide a fast response for some transactions in particular applications;

- ability to account for all usage of network resources, and processing resources within the network, to individual users;

- provision of networking standards compatible with those in the United States.

- provision of an open data communications border with the United States;

- provision of standards allowing Canadian companies to compete with the United States.

1.4.2 Canadian Industry Involvement in Open Systems Interconnection

The interconnection of similar systems is being achieved using lower level protocols or manufacturer architecture offerings.

The majority of companies involved in connecting dissimilar systems use lower level protocols, i.e. make one processor emulate a terminal when talking to another processor. Some customized solutions have been produced, but these do not provide full open systems interconnection as addressed by the International Standards Organization.

Most companies would prefer to use manufacturer offerings, but some would use new standards if they were supported by manufacturer products.

In the majority of cases, the ability to interconnect systems over public networks is of major importance, and many companies would be satisfied, in the short term, with a solution of the SNA/PDN interface problem.

1.4.3 Canadian Industry Involvement in Higher Level Protocols

The majority of organizations surveyed are not involved in the design or development of higher level protocols. Of those organizations that are involved, most of them are concerned with customizing manufacturer protocols, or developing implementation or application specific protocols.

Only the communications organizations are actively involved in the design and development of higher level protocols for non-specific open systems interconnection.

2. LAYER FUNCTIONALITY AND PROTOCOLS

2.1 Introduction

This section presents the results of further research into the detailed functionality of the Presentation, Session and Transport Layers of the ISO's Reference Model, IBM's SNA, and Digital's DNA . Particular emphasis was placed on specific user services within the Presentation Layer, and on the boundary between the Presentation and Session Layers.

This section is presented as three sub-sections (one per layer). Each sub-section is prefaced by an executive summary which highlights the functional relationships between the ISO reference model and the manufacturer architectures. The remainder of each sub-section details the functionalities and protocols provided, at that layer, by the ISO Model and the manufacturer architectures. Comparison to the ISO Model are made in the discussions of the manufacturer architectures.

IBM's System Network Architecture (SNA) was considered because of its prevalence in the communications marketplace, as shown in section 1, and its architectural compatibility with ISO's Open System Interconnection Model.

DEC's DECNET or Digital Network Architecture (DNA) was considered not only because of its popularity, but because of recent announcements by DEC to eventually support SNA, X.25, and provide enhanced routing capabilities in DNA.

HP's Distributed System Network (DSN) was not considered, as it was in the interim report of this study, because of time constraints, lack of detailed functional specifications, and its extensive divergence from the OSI philosophy. It was felt that any comparative analysis would be worthless. The first sub-section, which follows, details the functionalities of OSI's Presentation Layer.

2.2 Presentation Layer

Executive Summary

The functions provided by the network services and presentation services of SNA (function interpreters for function management data (FI.FMD)) have been found to correspond closely to the functions defined within the Presentation Layer of the ISO Model. In particular, the following correspondences could be made:

- session establishment, session recovery and session termination functions, defined by the ISO Model, correspond to similar functions provided by the network services FI.FMD of SNA;

- expedited flow, defined by the ISO Model, corresponds to the expedited flow provided by the network services FI.FMD of SNA. It is made available to the SNA end user for interrupts and is used by the network services for transmission of commands;

- purge of session, defined by the ISO Model, corresponds to the purge functions provided by the network services and presentation services FI.FMD's of SNA. The presentation services allow abnormal termination of data flow and the network services allow abnormal termination of sessions. These functions together correspond to the purge operation;

- selection of initial presentation options, defined by the ISO Model, corresponds to similar functions provided by the network services and presentation services FI.FMD's of SNA. The network services allow selection of those options defined by the session establishment procedures. The presentation services allow selection of those options determined by the use of control headers, and for these there is a simple negotiation procedure available in some cases as suggested by the ISO Model;

- renegotiation of options during a session, defined by the ISO Model, can be performed using the presentation services FI.FMD of SNA on those options that are specified using control headers. The negotation procedure available during the selection of initial presentation options may also be used here;

- transformation of transmitted data, as defined by the ISO Model, is performed by the presentation services FI.FMD of SNA;

- control access to data formats, defined by the ISO Model, corresponds to functions within the presentation services FI.FMD of SNA which do not allow data formats to be altered by one of the communicating partners unless the change is acceptable to the other;

A direct correspondence cannot be made between the ISO Model and SNA for the following functions:

- Special Function Options, those defined by the ISO Model are provided within SNA by the compression and compaction functions performed by the presentation services FI.FMD. Encryption and decryption functions, which the ISO Model also defines as special function options, are available within SNA, but are performed by data flow control.

The end user services described by the ISO Model are all provided by SNA. The detailed correspondence for each of the three services mentioned in the ISO Model are as follows:

(a) Virtual Terminal Service

- selection of terminal class, defined by the ISO Model, corresponds to the selection of session type, which is defined by the logical unit (LU) type of SNA. Not all of the terminal classes identified by the ISO Model are currently defined within SNA, but new LU types can be defined without altering the protocol;

- negotiation of profile, and data and command transfer, defined by the ISO Model, correspond to similar functions provided by SNA;

- forms management functions, mentioned in the ISO Model, are provided by SNA. These functions have been developed in some detail for SNA and contain facilities not covered by the ISO Model;

- control of operation, defined by the ISO Model, corresponds to a similar function provided by SNA.

(b) Virtual File Service

- encoding/decoding of internal/external attributes, defined by the ISO Model, corresponds to the exchange of status information within SNA;

- formatting and communication of file commands and data, defined by the ISO Model, corresponds to similar functions provided by SNA.

(c) Job Transfer and Manipulation Service

Control of record structures and devices, command formatting, and data formatting, defined by the ISO Model, correspond to similar functions provided by SNA. This is treated as a special case of file transfer by SNA.

The Presentation Layer functions defined by the ISO Model correspond to functions provided by the Data Access Protocol (DAP) of DNA. In particular, the following correspondences can be made:

- session establishment, session recovery, and session termination, as defined by the ISO Model, correspond to similar functions provided by DAP;

- selection of initial presentation options, defined by the ISO Model, corresponds to the identification of system type, protocol version, and generic capabilities of each partner immediately after link establishment as provided by DNA. No negotiation facilities are currently available other than the yes/no type of response;

- renegotiation of options during a session, defined by the ISO Model, can be provided by DAP using the functions mentioned for selecting the initial presentation options. The presentation options may be redefined during a session but no negotiation facilities are currently available;

- purge of session, as defined by the ISO Model, is provided by the DAP abort command function;

- provision and use of an expedited flow, defined by the ISO Model, corresponds to the use of two sub-channels by DAP. The expedited flow corresponds to the interrupt sub-channel.

A direct correspondence cannot be made between the ISO Model and DNA for the following functions:

- Transformation of Transmitted Data, the functions defined by the ISO Model, are similar to the conversion functions provided by DAP. These are restricted by DAP and can only be performed by one of the communicating partners;

- Control Access to Data Formats and Special Function Options, the functions as defined by the ISO Model are not available within DAP.

As DNA requires that each of the nodes in a DNA network is a processor, and because the data access protocol is designed primarily for access to remote data, the virtual terminal service as defined by the ISO Model is not available within DNA. The terminal to local processor communication is outside the DNA network.

The other end user services as described in the ISO Model are provided by DNA. The detailed correspondence for each of the other two services is as follows:

(a) Virtual File Service

- encoding/decoding of internal/external attributes as defined by the ISO Model corresponds to the exchange of attributes messages in DNA;

- formatting and communication of file commands and data as defined by the ISO Model corresponds to similar functions provided by DNA.

(b) Job Transfer and Manipulation Service

The job transfer and manipulation service defined by the ISO Model is provided to some extent by a remote batch job submission facility provided by DNA. This is a special case of the DNA file transfer service and the functions defined within the ISO Model Presentation Layer are not provided.

The functions provided by SNA correspond more closely to those defined by the ISO Model than do those of DNA. All the functions defined by the ISO Model are provided to some extent by SNA, but generally are not available in DNA.

Detailed descriptions of the functions defined by the ISO Model and the functions and protocols of SNA and DNA are contained in the following sub-sections.

2.2.1 ISO Model Description

2.2.1.1 Functions Within the Layer

The ISO Presentation Layer provides to the Application Layer, management of formats and management and performance of transformations.

The functions include:-

- session establishment
- selection of initial presentation options
- renegotiation of options during a session
- transformation of transmitted data
- control of access to data formats
- special function options
- purge of session
- session recovery
- provision and use of an expedited flow
- session termination

These proposed functions are described in more detail in the following pages.

2.2.1.1.1 Session Establishment

Session establishment is requested by the applications through the Presentation Layer. During establishment an agreement is reached on the available transformations between the two sets of presentation services. The initial presentation service at a node may be implicit, for example, in the address.

2.2.1.1.2 Selection of Initial Presentation Options

The presentation options can be selected by value, by name or by prior agreement. The method of negotiation allows for simple use and is therefore structured in a hierarchical fashion. In simple cases, predefined presentation options can be used, and in complex cases the selection may include an evaluation of the best match of presentation services. To do this each entity declares which values are acceptable with an indication of the 'cost' of each option.

2.2.1.1.3 Renegotiation of Options During a Session

Where complex selection procedures are used for presentation options, renegotiation of those options may take place during a session. This means that the presentation-entities may switch from data transfer, to transfer of control information, and back again. This implies the capability of the applications to switch from sending one type of data to sending another type of data, possibly with different blocking or priority etc., without closing down and then restarting the session.

2.2.1.1.4 Transformation of Transmitted Data

Once the combination of presentation services has been agreed on, transmission of data can begin. The sending presentation-entity performs it's agreed half of the required transformation and it labels the data items transmitted with an identifier, meaningful to the receiving presentation-entity, which specifies the format of the data. On checking this identifier, the receiving presentation-entity then performs it's agreed half of the required transformation before passing the data to the application.

2.2.1.1.5 Control Access to Data Formats

When either application process wishes to update the format of the data, mechanisms within the Presentation Layer are invoked which ensure that these changes are carried out in an orderly manner. These changes may also result in a change in the presentation options being used.

2.2.1.1.6 Special Function Options

Some special purpose transformations may be available within the Presentation Layer, possibly to improve the efficiency of the data transfer. If a special purpose transformation is available to both presentation-entities, then that transformation may be applied by the sending presentation-entity and removed by the receiving presentation-entity. Examples of these special purpose transformations are:

- compression/compaction and expansion
- encryption and decryption etc.

2.2.1.1.7 Purge of Session

Purging of the session may be requested by either presentation-entity and returns both entities to a predefined state with possible loss of information.

2.2.1.1.8 Session Recovery

After a reset operation or any abnormal behaviour, session recovery functions will be used within the Presentation Layer to set the session to a defined state.

2.2.1.1.9 Provision and Use of an Expedited Flow

The Presentation Layer makes available to the Application Layer an expedited flow for interrupts etc. The expedited flow is also used by the Presentation Layer itself for exchanging short signalling or control messages. The expedited flow will bypass any queues of normal data within the lower layers and an example of its use would be for the purge request. See figure 2.2.1.1.9-1.

2.2.1.1.10 Session Termination

For normal session termination the application would request termination through the Presentation Layer. In the case of abnormal termination the application would be informed of abnormal end of session and be given a reason.

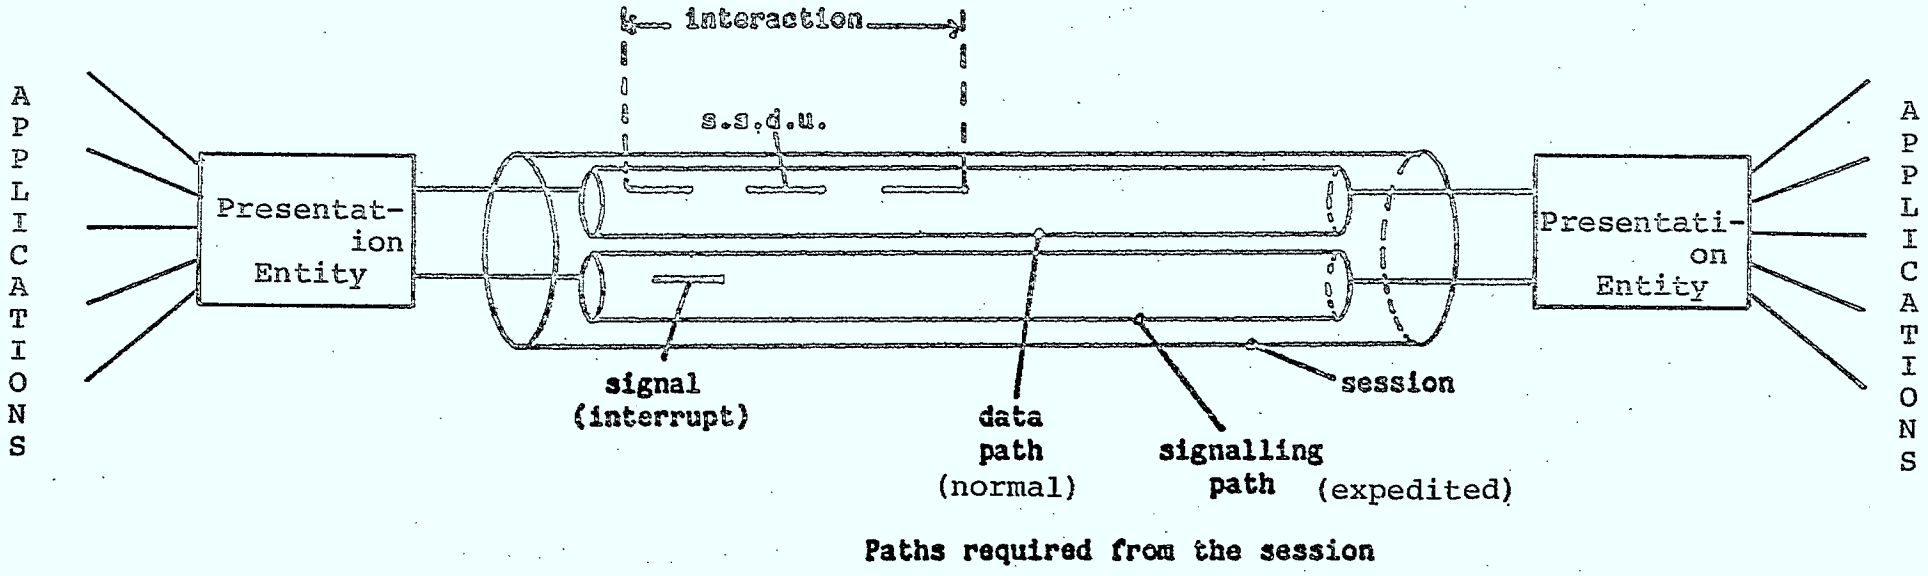

Paths required from the session

Fig 2.2.1.1.9 - 1 Provision and Use of Expedited Flow

2.2.1.2 End-User Services v.s. Layer Protocols

An end user service is an application meaningful to the end-user. It may require functions from both the ISO Presentation Layer and the ISO Application Layer. A layer protocol, however, provides the functions required within a particular layer. There is therefore a distinction to be made between an end-user service 'protocol' and a protocol defined within the architecture as a peer or layer protocol.

The Presentation Layer also contains some functions common to different end user services. Any functions, specific to an end user service, not provided by the Presentation Layer will be performed within the Application Layer. This implies that the Session Layer functions are application independent.

The three types of end user services considered are:

- Virtual Terminal Service
- Virtual File Service
- Job Transfer and Manipulation Service

The functions required by these three application services will be provided by both Application Layer protocols and Presentation Layer protocols. See figure 2.2.1.2-1.

Fig 2.2.1.2 - 1 APPLICATION SERVICES PROTOCOL DISTRIBUTION

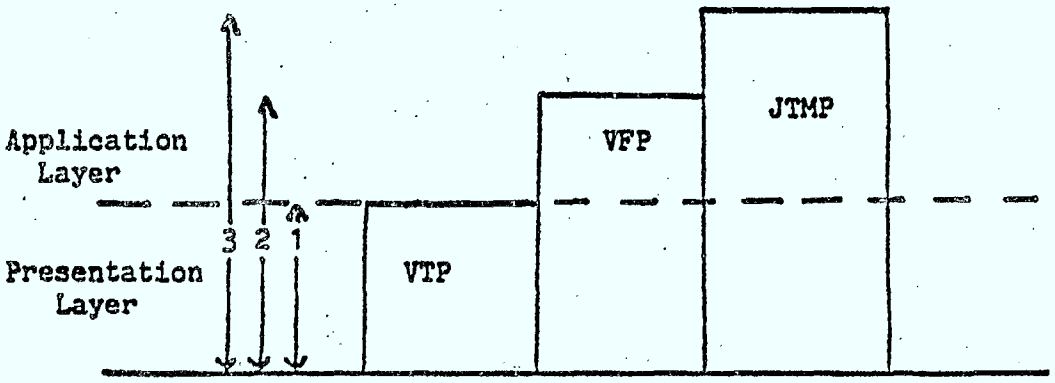

Functional Distribution of VTP, VFP and JTMP

Note : This figure does not preclude common provision of format management and transformations for VTP, VFP and JTMP.

In this diagram, the functions in (1) are data and command structuring according to transformation rules, in (2) are file manipulation, management + (1) during file control and transfer, and in (3) are station management and command language + (2) + (1).

VTP - Virtual Terminal Protocol

VFP - Virtual File Protocol

JTMP - Job Transfer and Manipulation Protocol

2.2.1.2.1 Virtual Terminal Service

The virtual terminal service will allow human partners to access application partners through a large variety of terminals. This involves eliminating incomptabilities between terminals of the same class. Several classes of terminals will be identified by ISO, each of which will have a single set of presentation services that can be standardized and is called the 'virtual terminal'. See figure 2.2.1.2.1-1.

The ISO Model defines the following functions within its Presentation Layer:

(a) Selection of Terminal Class

Several terminal classes have already been identified by ISO and include:

- stream class
- page class
- data entry class
- word processing class
- graphics class

(b) Negotiation of Profile

During the initiation of the virtual terminal the two presentation-entities negotiate a virtual terminal profile that represents the effective characteristics of the real terminal. A minimum profile for the class will be assumed when the negotiation does not take place. The methodology to be used for this negotiation has not yet been defined by ISO.

Fig 2.2.1.2 - 1 APPLICATION SERVICES PROTOCOL DISTRIBUTION

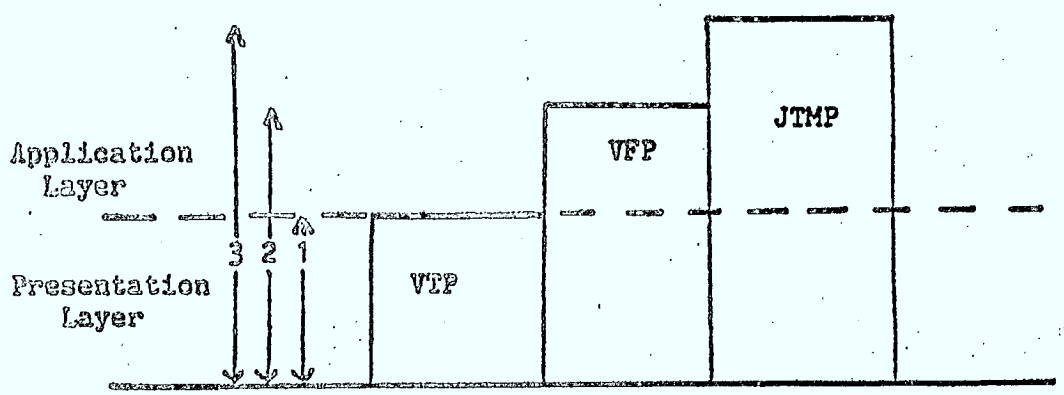

Functional Distribution of VTP, VFP and JTMP

Note : This figure does not preclude common provision of format management and transformations for VTP, VFP and JTMP.

In this diagram, the functions in (1) are data and command structuring according to transformation rules, in (2) are file manipulation, management ÷ (1) during file control and transfer, and in (3) are station management and command language ÷ (2) ÷ (1).

VTP - Virtual Terminal Protocol

VFP - Virtual File Protocol

JTMP - Job Transfer and Manipulation Protocol

(c) Data and Command Transfer

The exchange of data is governed by the subset of the peer Presentation Layer protocol corresponding to the selected virtual terminal. The protocol will encompass, for each terminal class, a different set of commands. Some commands however will be common to several classes of terminal.

(d) Forms Management

If the class of terminal selected involves the use of forms, then the management of those forms, including form description and form assignment, will be performed by the Presentation Layer.

(e) Control of Operation

During the terminal operation there will be several different states. Each of these states will allow a different subset of the set of commands available to this terminal class.

2.2.1.2.2 Virtual File Service

The virtual file service will allow the representation of files and file activity in a standard fashion. This is done by expressing all file manipulation requirements in terms of a virtual filestore, which would contain only files described by standard attributes with names presented in a standard manner, and would permit only standard sets of operations. See figure 2.2.1.2.2-1.

Fig 2.2.1.2.1 - 1 THE VIRTUAL TERMINAL SERVICE

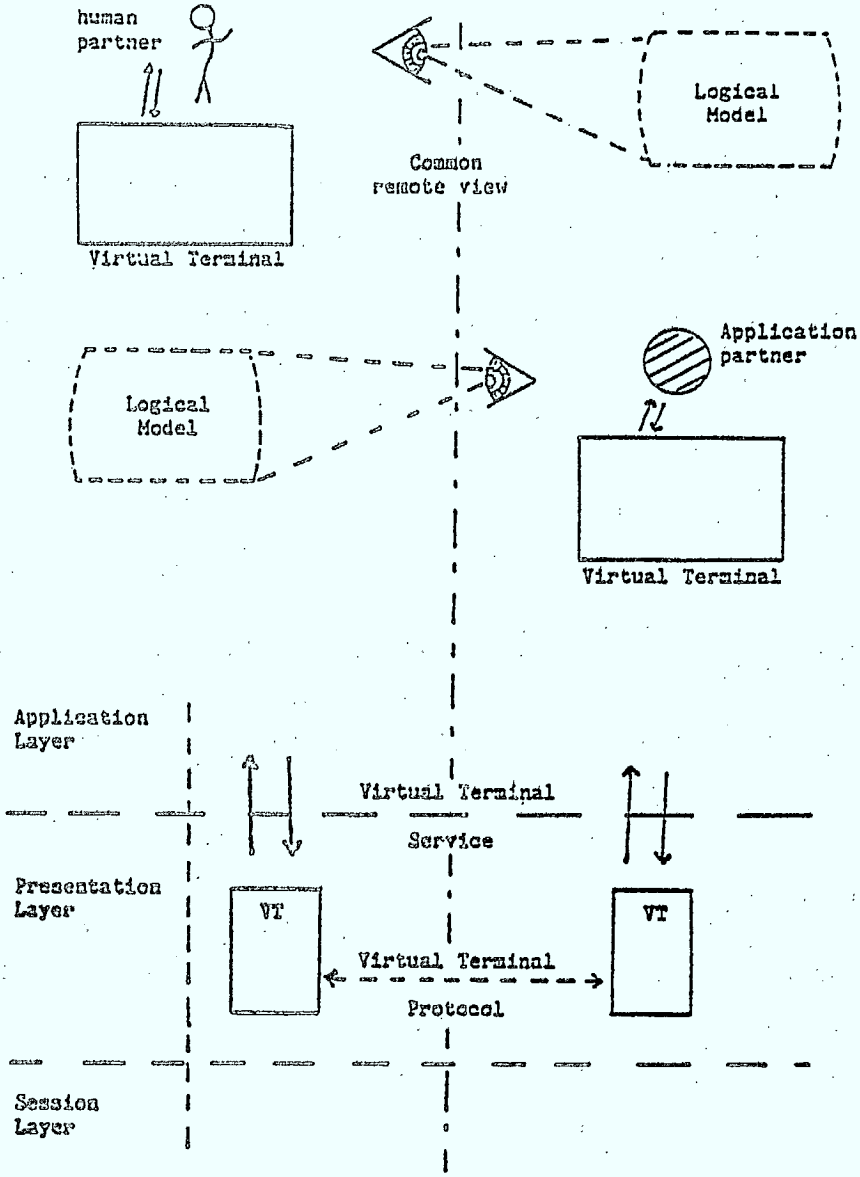

Logical models and Virtual Terminal Service

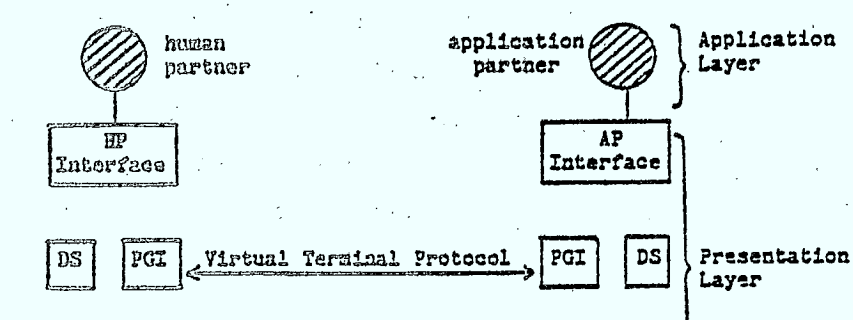

DS data store HP human partner AP application partner
VT virtual terminal PGI protocol generator and interpreter

Example of interaction between human partner and application partner

The ISO Model defines the following functions within its Presentation Layer:

(a) Encoding/Decoding Internal Attributes

Internal attributes of a file are concerned with the internal structure and properties of the file contents and deal with size, structure, encoding representation, data types etc. The presentation-entity will encode these attributes into the standard virtual filestore format so that they may be understood by the partner presentation-entity which will, if required, decode the attributes so that they may be understood by the application entity.

(b) Encoding/Decoding External Attributes

External attributes of a file are concerned with its properties as an indivisible object for administration purposes. The Presentation Layer deals with these in the same way as the internal attributes.

(c) Formatting Virtual Filestore Commands

All file access is requested by using commands that the presentation-entity converts to standard format virtual filestore commands before passing to the Session Layer.

(d) Communication of File Data and Commands

The presentation-entity will convert file data and file management commands into standard format before passing them to the Session Layer.

Fig 2.2.1.2.2 - 1 THE VIRTUAL FILE SERVICE

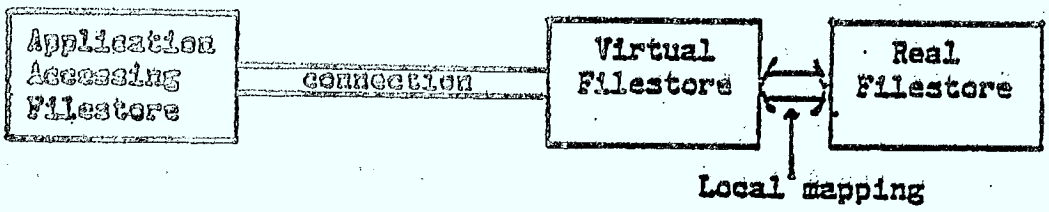

Virtual File Service

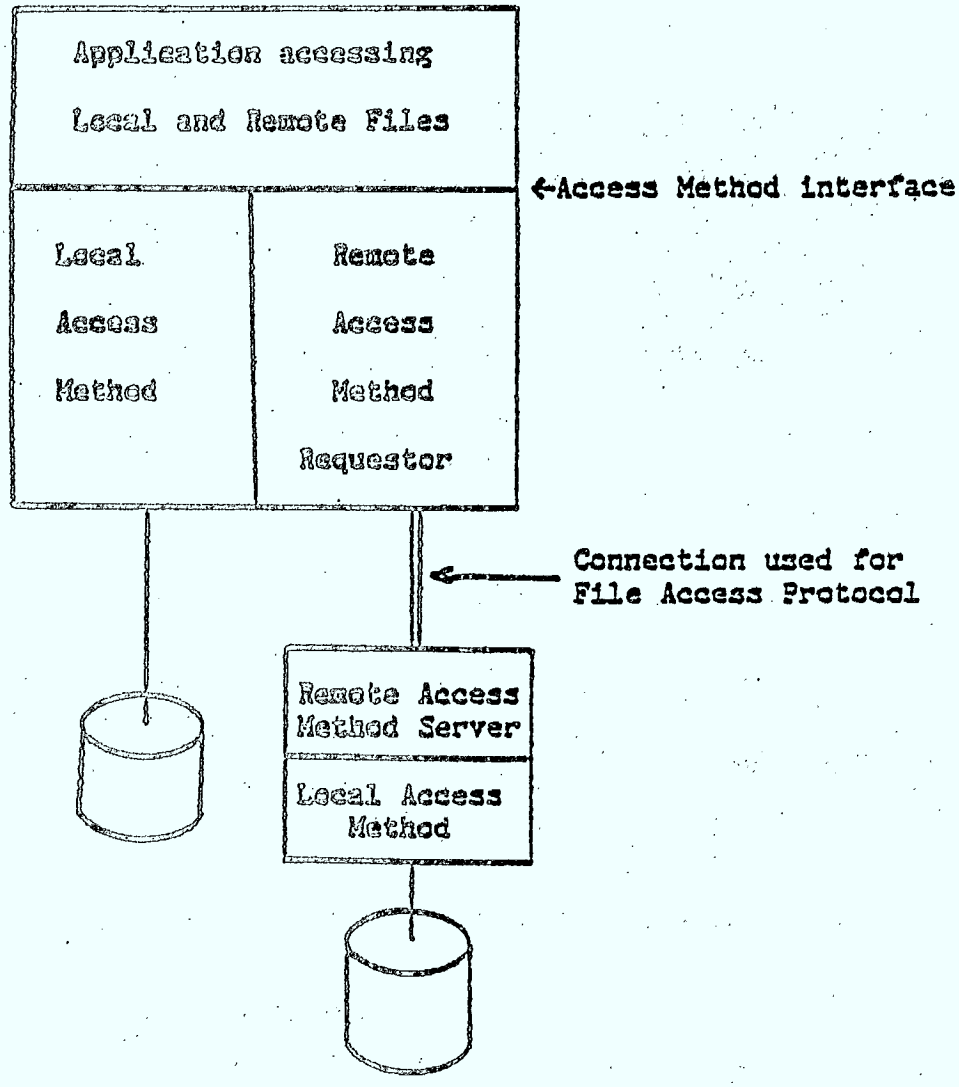

File access

2.2.1.2.3 Job Transfer and Manipulation Service

The job transfer and manipulation service will allow a user to submit a job to a remote host using either the JCL of the remote host, if it is known to the end user, or a network wide JCL, freeing the user from differences between hosts. See figure 2.2.1.2.3-1.

The ISO Model defines the following functions within its Presentation Layer:

(a) Control of Record Structures and Devices

Record structures and device control commands are converted to a standard format before being passed to the Session Layer.

(b) Command Formatting

Management, access and job commands are converted to a standard format before being passed to the Session Layer.

(c) Data Formatting

During the data transfer phase the data is converted to a standard structure and format before being passed to the Session Layer.

Fig. 2.2.1.2.3 - 1 THE JOB TRANSFER AND MANIPULATION SERVICE

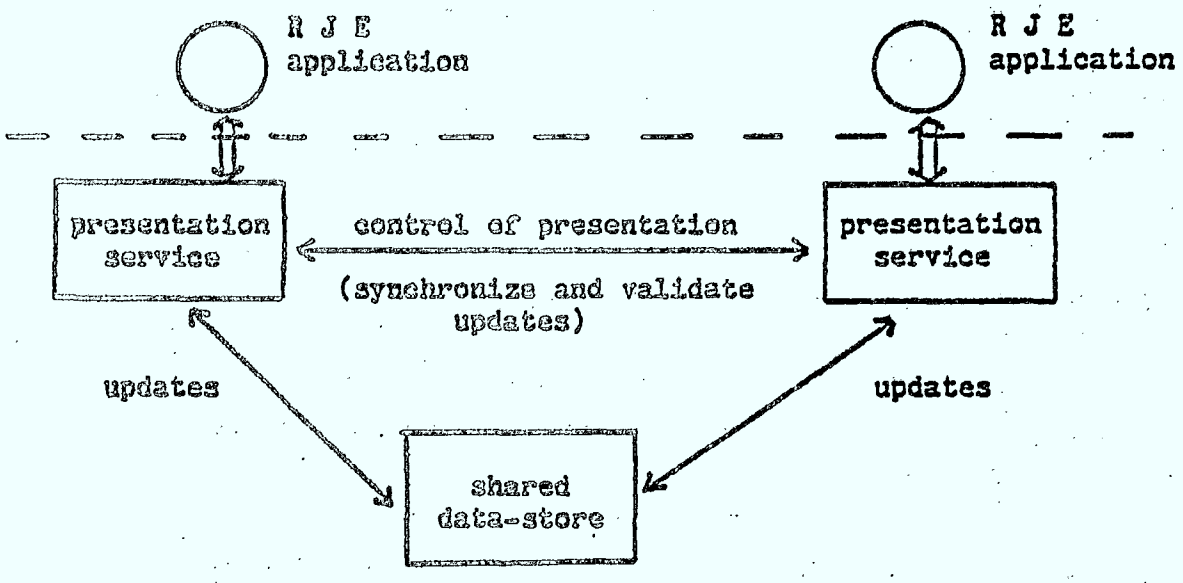

Shared data store

2.2.1.3 Functions at the Session Layer Boundary

At the boundary of the Presentation Layer and the Session Layer, the Presentation Layer requests services from the Session Layer to fulfill the requirements placed on the Presentation Layer by the applications. The Presentation Layer passes to the Session Layer the following data:

(a) From the Application Layer

- session establishment requests/responses
- class of service requests
- optional facilities requests
- session termination requests

(b) From the Presentation Layer

- control data
- transformed data
- data transfer mode request
- flow type request

These are described in more detail in the following sections.

2.2.1.3.1 Session Establishment Requests/Responses

Session establishment is requested by the application entity. The Presentation Layer passes this request directly to the Session Layer to initiate the session establishment procedure at the lower levels. Responses from the sink application are handled in the same way.

2.2.1.3.2 Class of Service Requests

The class of service is requested by the application entity. The Presentation Layer passes this request directly to the Session Layer either to be used in session establishment or to initiate class of service changes at lower levels during a session.

2.2.1.3.3 Optional Facilities Requests

Optional facilities may be requested by the application entity. The Presentation Layer passes these requests directly to the Session Layer to activate the selection mechanisms used to determine optional session parameters to implement the requested optional facilities.

2.2.1.3.4 Session Termination Requests

The session termination is requested by the application entity. The Presentation Layer passes this request directly to the Session Layer to initiate the session termination procedure at the lower levels.

2.2.1.3.5 Control Data

Control messages generated in the Presentation Layer are passed to the Session Layer as data for transmission.

2.2.1.3.6 Transformed Data

Data from the Application Layer is passed to the Session Layer as data for transmission after it has been transformed by the Presentation Layer. Data for transmission may be passed to the Session Layer in unlimited variable length blocks.

2.2.1.3.7 Data Transfer Mode Request

The data transfer mode is requested by the Presentation Layer. The Session Layer will provide the requested mode of data transfer, which may be one-way, two-way alternate or two-way simultaneous.

2.2.1.3.8 Flow Type Request

The flow type is requested by the Presentation Layer. The Session Layer will provide the requested flow type which may be normal or expedited.

2.2.2 IBM-SNA Functions and Protocol

The Presentation Layer of the ISO Model corresponds to the Function Interpreters for Function Management Data (FI.FMD) of SNA. Refer to figure 2.2.2-1 for the mapping of ISO and SNA layers.

The SNA protocol at this layer is implemented using two different FI.FMD types. One provides network services and the other provides presentation services.

2.2.2.1 Functions Within the Layer

The functions performed by the Function Interpreters for Function Management Data are split into two different categories:

- network services, which are involved in:
  - session establishment
  - selection of initial presentation options
  - purge of session
  - session recovery
  - expedited flow
  - session termination
- presentation services, which are involved in:
  - selection of initial presentation options
  - renegotiation of options during a session
  - transformation of transmitted data
  - control of access to data formats
  - special function options
  - purge of sessions

Some functions require the use of both of these service FI.FMD types.

Fig 2.2 .2-1 TABLE SHOWING THE MAPPING OF THE OSI MODEL ONTO SNA

OPEN SYSTEM INTERCONNECT MODEL		IBM SYSTEMS NETWORK ARCHITECTURE
LAYER	NAME	
7	APPLICATION LAYER	NETWORK ADDRESSABLE UNIT SERVICES MANAGER
6	PRESENTATION LAYER	FUNCTION INTERPRETERS FOR FUNCTION MANAGEMENT DATA
5	SESSION LAYER	DATA FLOW CONTROL
		TRANSMISSION CONTROL
4	TRANSPORT LAYER	PATH CONTROL
3	NETWORK LAYER	
2	DATA LINK LAYER	DATA LINK CONTROL
1	PHYSICAL LAYER	PHYSICAL LAYER

2.2.2.1.1 Network Services

The network services provide functions through which the application-entities can monitor and control the network resources. The amount of control performed by a node will depend on its type (system services control point (SSCP), logical unit (LU) or physical unit (PU)). Figure 2.2.2.1.1-1 shows how network services are distributed between the application and presentation layers, i.e. between the Network Addressable Unit Services Manager and the FI.FMD, in each of three types of unit. In each case only a small part of the network services functions are performed by the FI.FMD.

The interaction between the three types is based on the division of the network into domains, each of which contains one SSCP and its associated PU's, LU's and links that the SSCP controls by having the capability to activate them. The interaction within a domain is shown in figure 2.2.2.1.1-2.

Fig 2.2.2.1.1-1 DISTRIBUTION OF NETWORK SERVICES

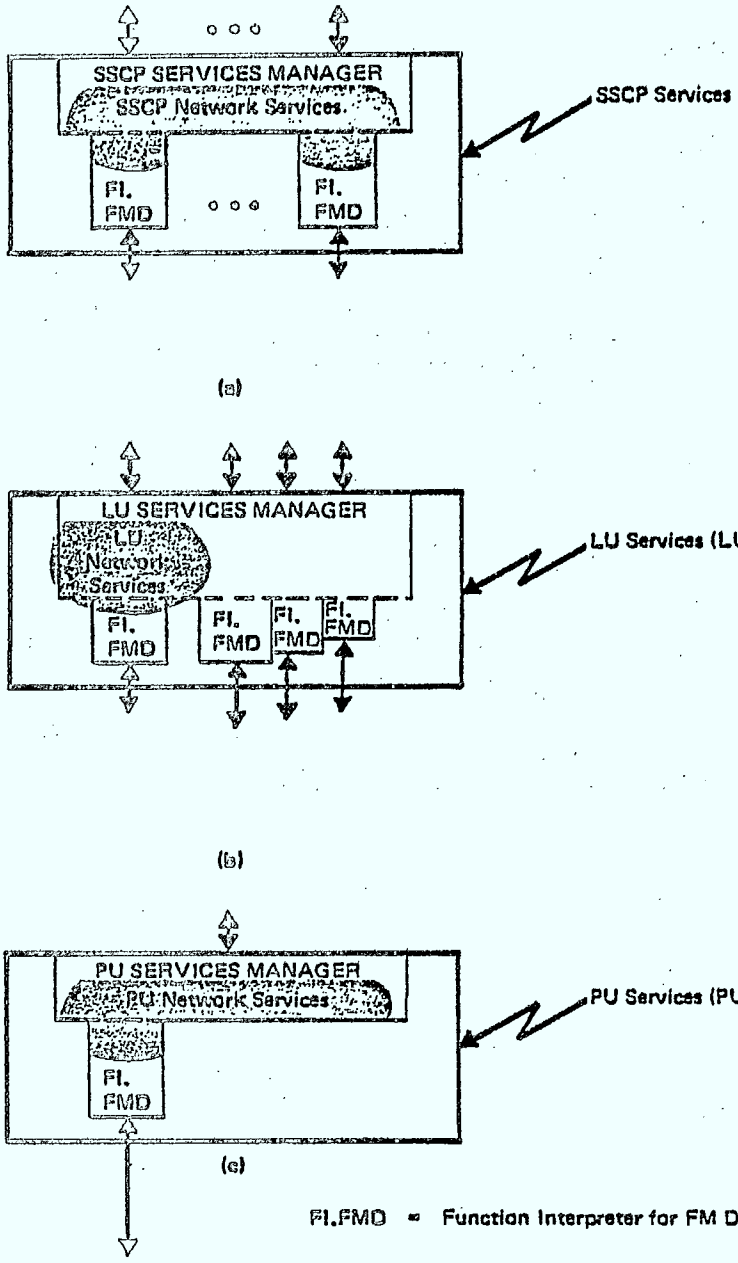

Typical sets of higher-level services in the three types of NAU: (a) in SSCP; (b) in LU; and (c) in PU.

Each logical unit has the capability for supporting at least two sessions concurrently, one with another logical unit and one with the SSCP. This is shown in figure 2.2.2.1.1-3.

The functions provided by network services may be better described by first considering its structure.

Each network services FI.FMD consists of the following protocol machines:

- FMD-RQ-RCV request, receive
- FMD-RQ-SEND request, send
- FMD-RSP-RCV response, receive
- FMD-RSP-SEND response, send

These protocol machines are coupled by a set of finite state machines (FSM) as shown in figure 2.2.2.1.1-4.

The network services functions may be classified as follows:

- configuration services
- maintenance services
- measurement services
- session services
- network operator services

The distribution of these different service types throughout the network, dependent upon the node type, is shown in figure 2.2.2.1.1-5.

Fig 2.2.2.1.1-2 NETWORK SERVICES INTERACTION WITHIN A DOMAIN

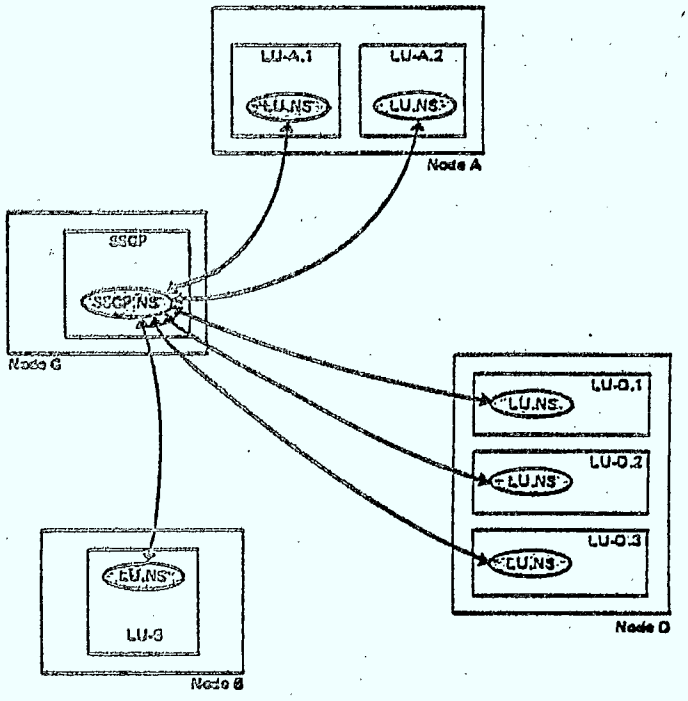

Concept of having some network services distributed among the LUs and the SSCP.

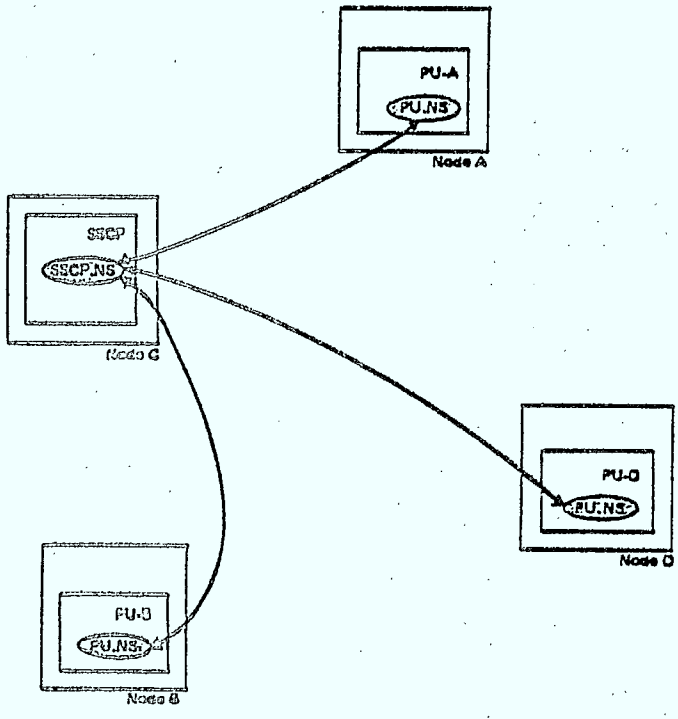

Concept of having some network services distributed among the PUs and the SSCP.

Fig 2.2.2.1.1-3 MINIMUM LOGICAL UNIT SESSIONS

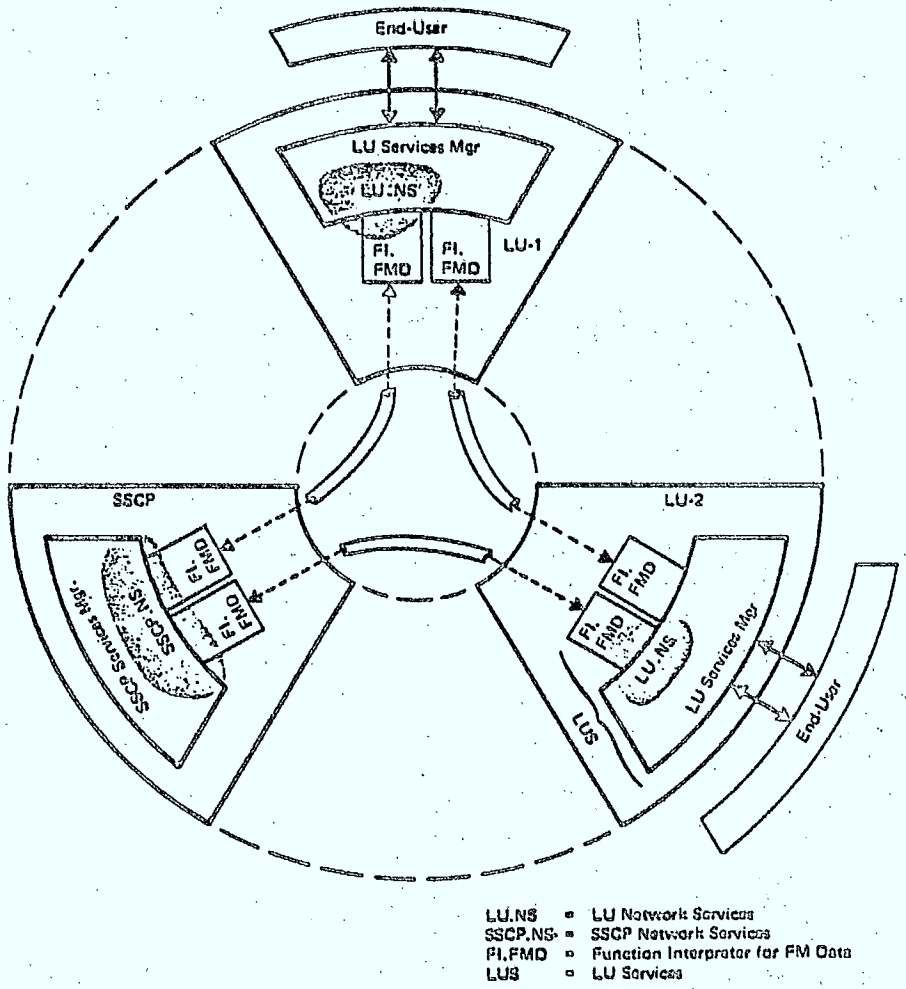

Network services and FI.FMD services in LU-LU and LU-SSCP sessions.

Fig 2.2.2.1.1-4 PROTOCOL MACHINES IN

NETWORK SERVICES

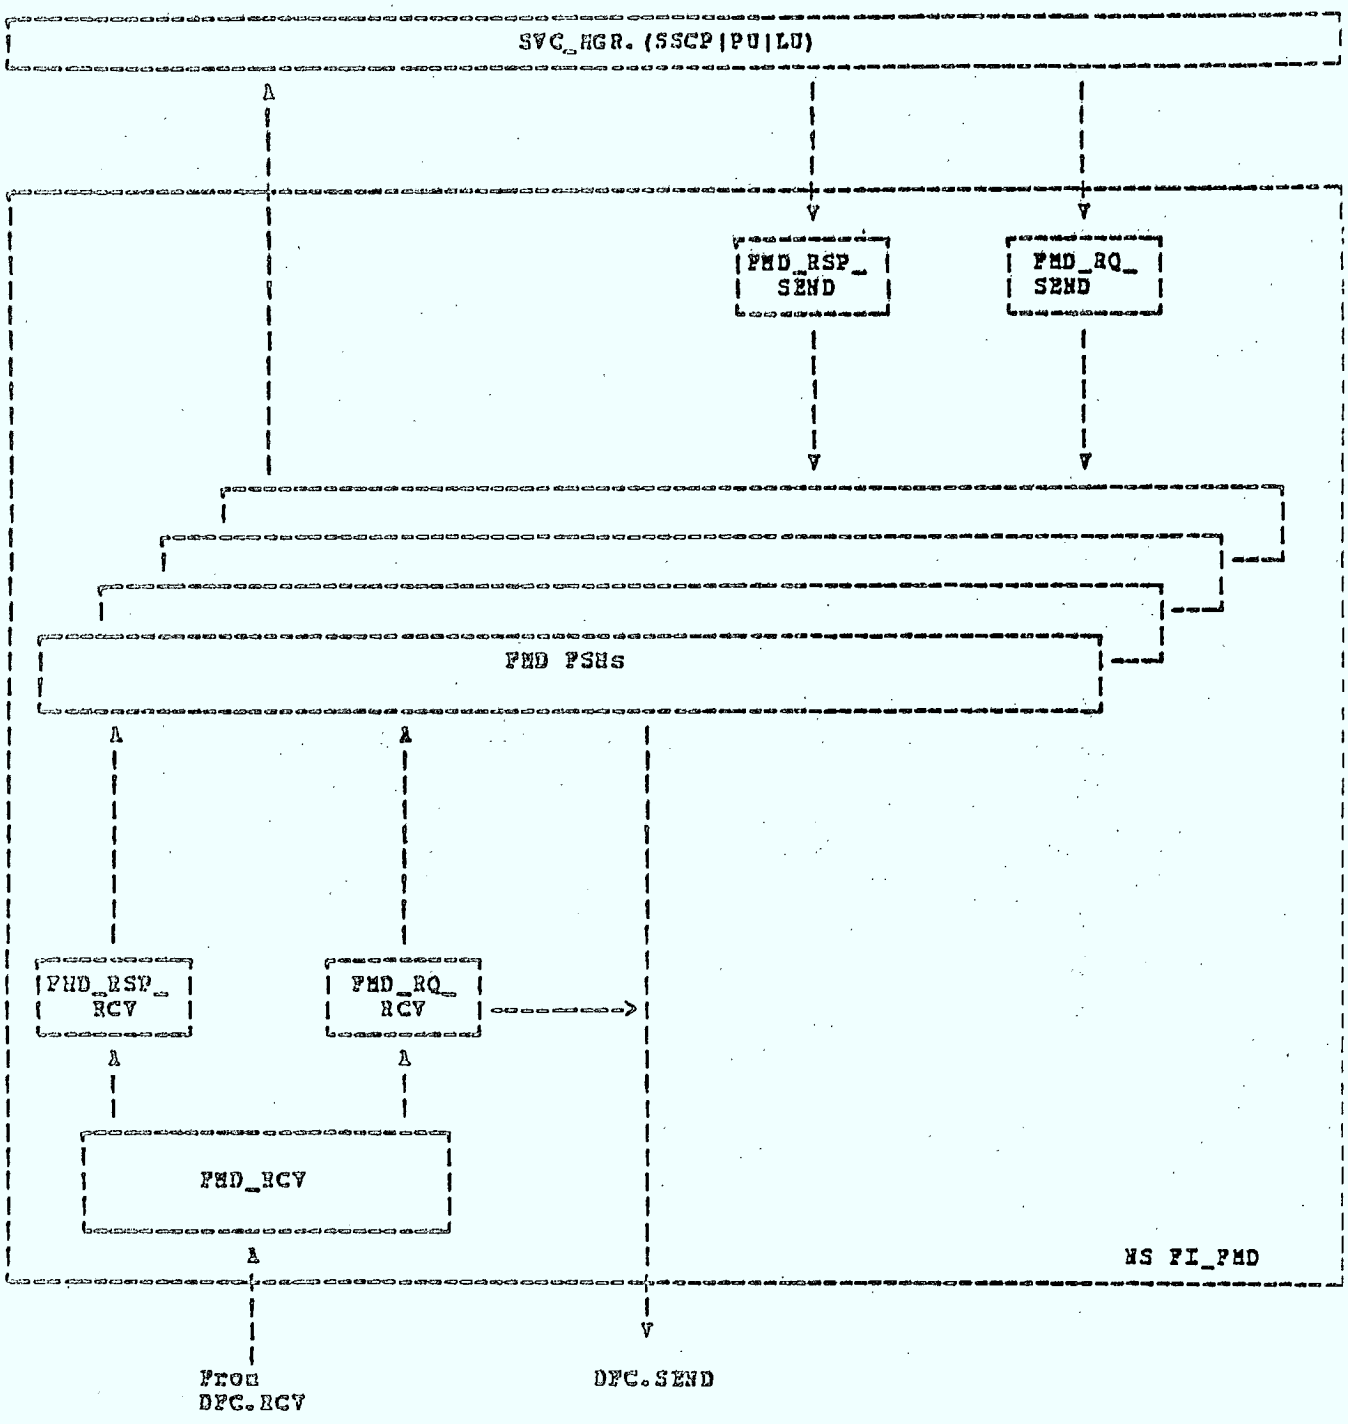

Structure of NS FI_FMDs

Fig 2.2.2.1.1-5 DISTRIBUTION OF NETWORK SERVICE TYPES

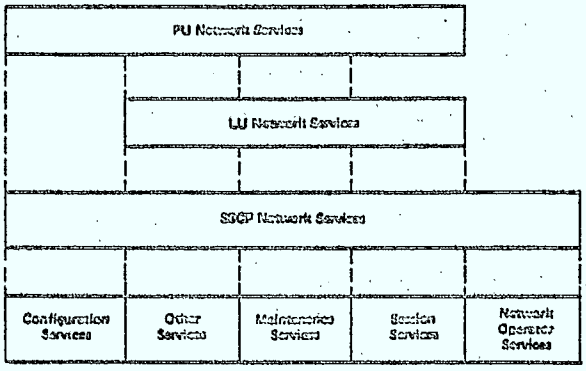

Potential distribution of network services among the SSCP, PU, and LU.

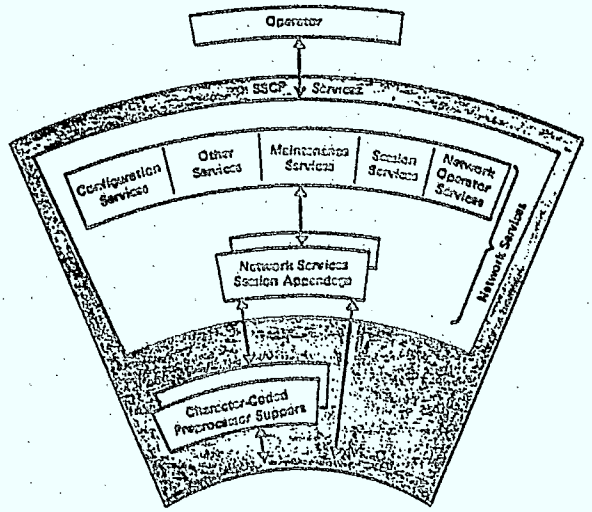

Types of services within the SSCP.

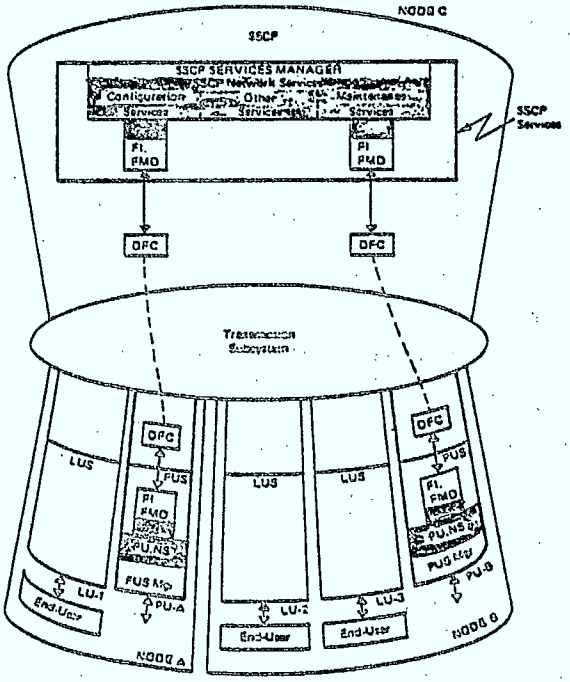

NS = Network Services
FI,FMD = Function Interpreter, FM Data
PUS = Physical Unit Services
LUS = Logical Unit Services
DFC = Data Flow Control

Distributed network services (shaded areas) among the SSCP and PUs in multiple nodes.

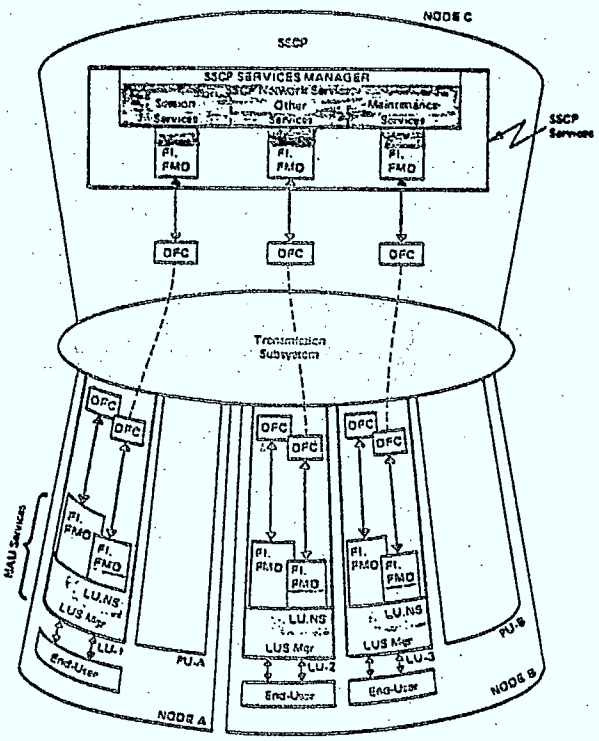

NS = Network Services
FI,FMD = Function Interpreter for FM Data
DFC = Data Flow Control

Network services (shaded areas) distributed between LUs and the SSCP.

The majority of network services functions are performed by the NAU Services Manager (which corresponds to the Application Layer of the ISO Model) and are supported by the FI.FMD protocol. Network session services functions, however, are provided by the network FI.FMD, and support the activation of a logical connection. These network session services functions include the following as defined in the ISO Model Presentation Layer:

(a) Session Establishment

Several different request/response unit (RU) formats are concerned with the establishment of the session, most of them contain different types of control information; for example, to be used in the case of session establishment failure. Different RU types are used depending on whether the session is with a logical unit in the same domain or another domain. Conversion of remote device names to addresses is also allowed.

(b) Selection of Initial Presentation Options

The session establishment RU formats include an implementation dependent specification of rules and options to be used during the session, the content of which is also dependent on the type of device. If the devices are intelligent enough, then this process may include negotiation of the best options to be used.

(c) Purge of Session

The session termination RU formats include the facility to specify a forced termination which deactivates the session immediately and unconditionally with possible loss of information. Reason for purge information may be given.

(d) Session Recovery

The session recovery RU formats allow the recovery procedures to be initiated by the user or the network manager. The reason for the use of the recovery procedures may be given.

(e) Expedited Flow

Those commands requiring use of the expedited flow use RU formats that include information identifying whether the command is to be part of the normal flow, i.e. to join any queues of data for the session, or to be part of the expedited flow and bypass any queues of data for the session.

(f) Session Termination

The session termination RU formats allow normal termination, as requested by the applications, providing orderly termination which may include the execution of an end-of-session procedure, and also allow identification of abnormal termination.

2.2.2.1.2 Presentation Services

The presentation services provide functions for presenting information to the end user in a usable form. These functions include handling of data streams, encoding and compressing data, and error recovery procedures.

The functions provided by presentation services may be better described by first considering its structure.

The abbreviations used are:
- PS = Presentation Services
- DFC = Data Flow Control
- FMD = Function Management Data
- RCV = Receive
- SEND = Send

Figure 2.2.2.1.2-1 shows the interrelationships of the protocol machines associated with presentation services.

The PS.DFC-RCV and PS.DFC-SEND provide the interface with Data Flow Control for passing information to the lower layers.

PS.RCV supports session functions on the receive side. It is responsible for controlling and routing the flow of function management data to the end user.

Fig. 2.2.2.1.2-1 STRUCTURE OF PRESENTATION SERVICES

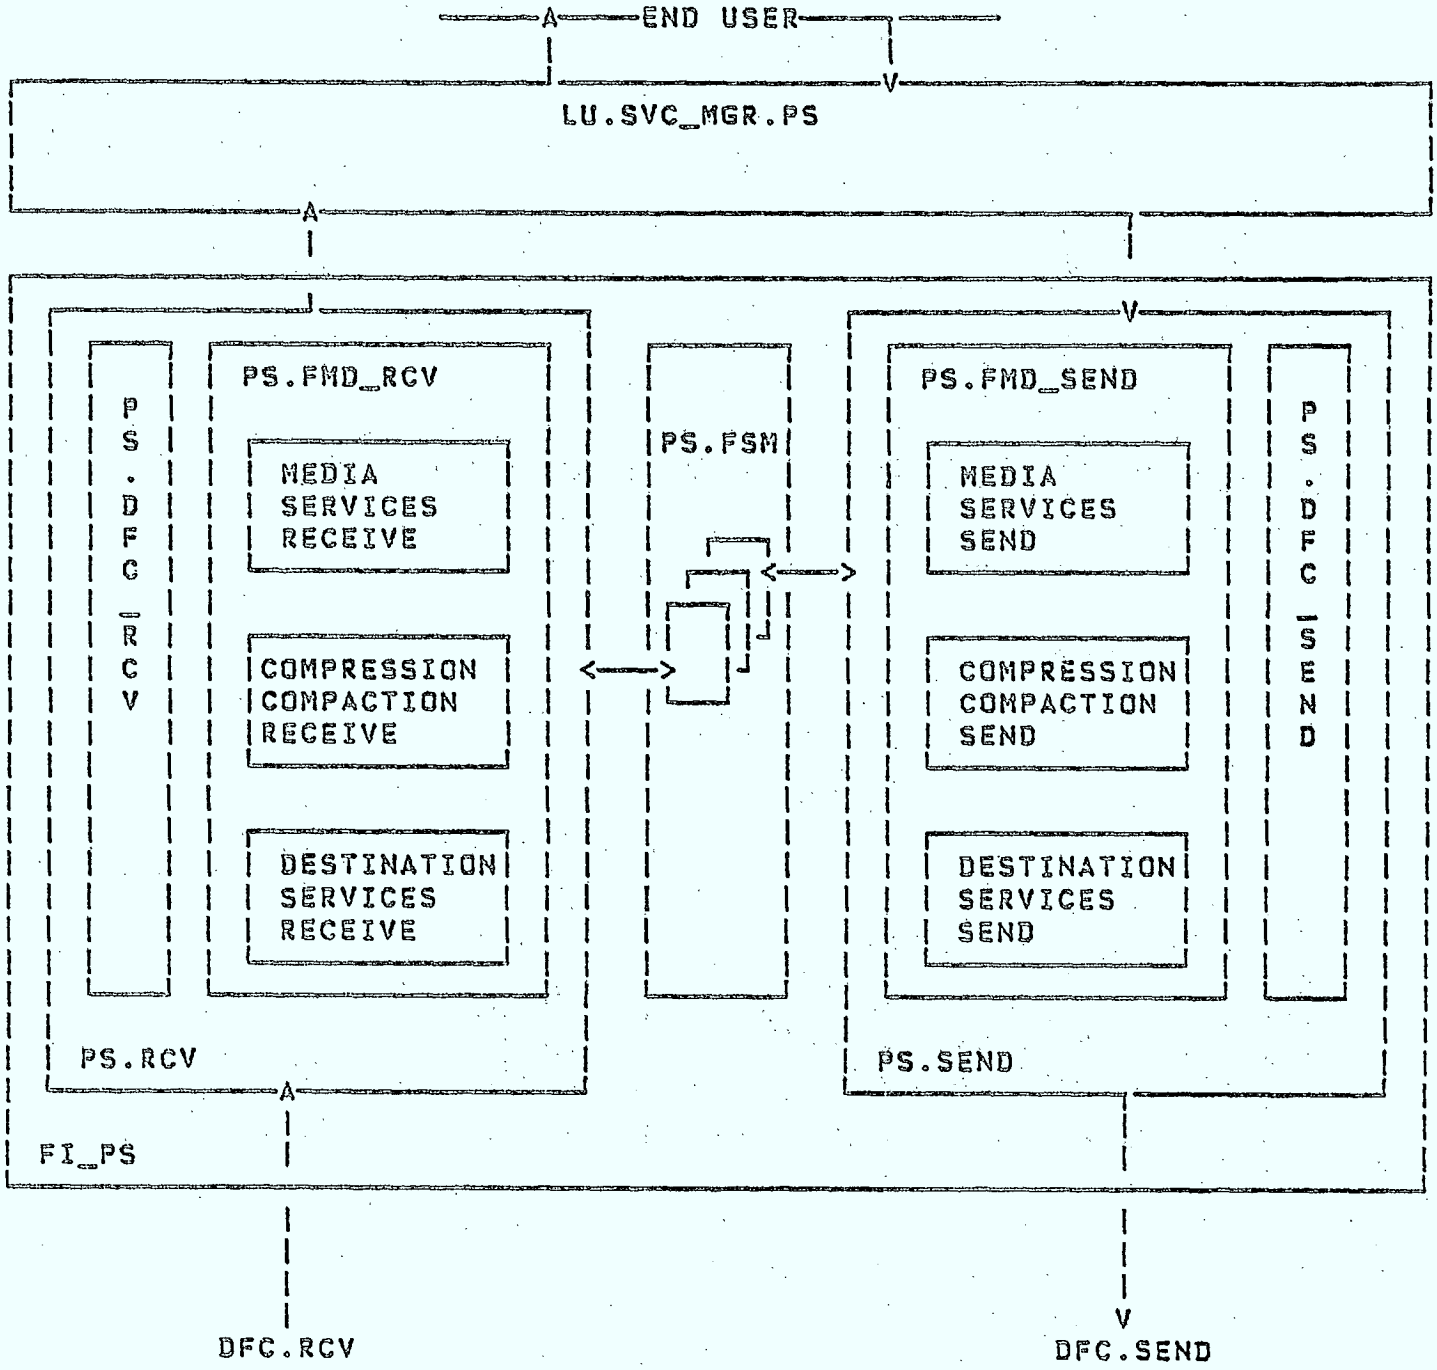

Structure of Function Interpreter for Presentation Services (FI_PS).

FI_PS consists of the following components:

Receive	Common	Send	
DS.RCV		DS.SEND	- Destination Services
CC.RCV		CC.SEND	- Compression/Compaction
MS.RCV		MS.SEND	- Media Services
PS.RCV		PS.SEND	- Presentation Services
	PS.FSM		- Finite State Machines

PS.SEND supports session functions on the send side. It is responsible for controlling and routing the flow of function management data from the end user. It notifies data flow control of the error status of requests and responses processed by PS.RCV.

PS.FSM is a collection of finite state machines that support presentation services functions.

DS.RCV processes control header information. It detects format errors and notifies PS.RCV.

DS.SEND checks for control header information format violations in the data stream being sent to the session partner by PS.SEND.

CC.RCV performs data decompression or data decompaction as requested by PS.RCV. It also detects format errors and notifies PS.RCV.

CC.SEND performs data compression or data compaction as requested by PS.SEND. It generates control information, inserts it in the data stream, and returns the compressed or compacted data stream to PS.SEND.

MS.RCV deblocks the data streams and validates its format. If the data stream format is faulty it notifies PS.RCV of the error.

MS.SEND validates data stream controls and formats the data stream received from the end-user. If format errors are detected it notifies PS.SEND.

The presentation services perform device control, data management, and data compression and compaction. They change the way the information is presented to match the needs or language of each end user. The presentation services protocol sub-set of the FI.FMD protocol provides these functions by the use of control information. The request/response unit (RU) carries control information and data between half sessions. The control information may take up a whole RU - a control RU, which contains a request or an acknowledgement. It may be included in a data RU as a control header, within the body of the data, or both.

Control headers enable an LU to send a data stream to a specific destination and control the way the data is presented at the destination. Control headers are the mechanism used by one LU to select some of the functions it wants the presentation services component of its session partner to perform.

During session establishment, decisions are made as to whether control headers will be used and, if so, which control header functions will be permitted. More than one control header may be present in the RU in which case a concatenation flag is set on in each control header that has another control header following it.

The control information is used to support the following functions mentioned in the ISO Model Presentation Layer description:

a) Selection of Initial Presentation Options

Initial presentation options are defined during session establishment by the parameters contained within the control header. The LU requesting the session specifies a set of parameters in a control header. These can be accepted or rejected by the receiving LU. If the receiving LU is capable of negotiating options, then it will accept the session request but specify a different set of parameters in another control header, otherwise it will be limited to a Yes/No type response to the session request.

b) Renegotiation of Options During a Session

A control header may be sent that changes the control parameters describing the data to be transmitted. These may be negotiated, as in the selection of the initial presentation options, thus allowing for renegotiation of options during a session.

c) Transformation of Transmitted Data

Once a set of presentation options is agreed upon, the transformations to be performed at each LU are defined, and are applied to the transmitted data until a new set of options are negotiated, or the session is ended.

(d) Control of Access to Data Formats

Data formats are defined within the presentation options that are specified using control headers. They may not be altered in either half-session unless the new options requested by one half-session are acceptable to the other half-session.

(e) Special Function Options

The special function options currently defined by SNA include data compression and data compaction. They are used within SNA to shorten network transmissions.

Data is compressed by eliminating gaps, empty fields, redundancies, or unnecessary data. The technique used within SNA is to select one character, called the prime compression character, and replace repetitive sequences of that character with a control character. A control character can also define strings of like characters that are not the prime compression character, but the control character must be followed by a byte containing the non prime character so that the receiver knows which character was compressed. The prime compression character is assumed to be the character blank. When another character is desired, it can be defined using a control header.

Data can also be compacted to shorten the length of transmissions. The technique used by SNA is to compact two bytes of data into one byte. Up to 16 characters from the total character set can be compacted. Again control characters are used in the data stream to define the beginning and end of each block of compacted data.

(f) Purge of Session

A control header may be used to abnormally end data traffic to the receiving half-session. This is used in the purge operation.

The types of presentation services required are defined by the presentation services profile and usage fields. The profile field identifies the session type, and the usage field describes the use of control headers and defines the data streams for the session. The location of these within the LU is shown in figure 2.2.2.1.2-2. The session type restricts the scope of the presentation services that may be used in that session, and the usage field provides further detailed information about specific presentation services to be used.

Fig. 2.2.2.1.2-2 PRESENTATION SERVICES PROFILE AND USAGE

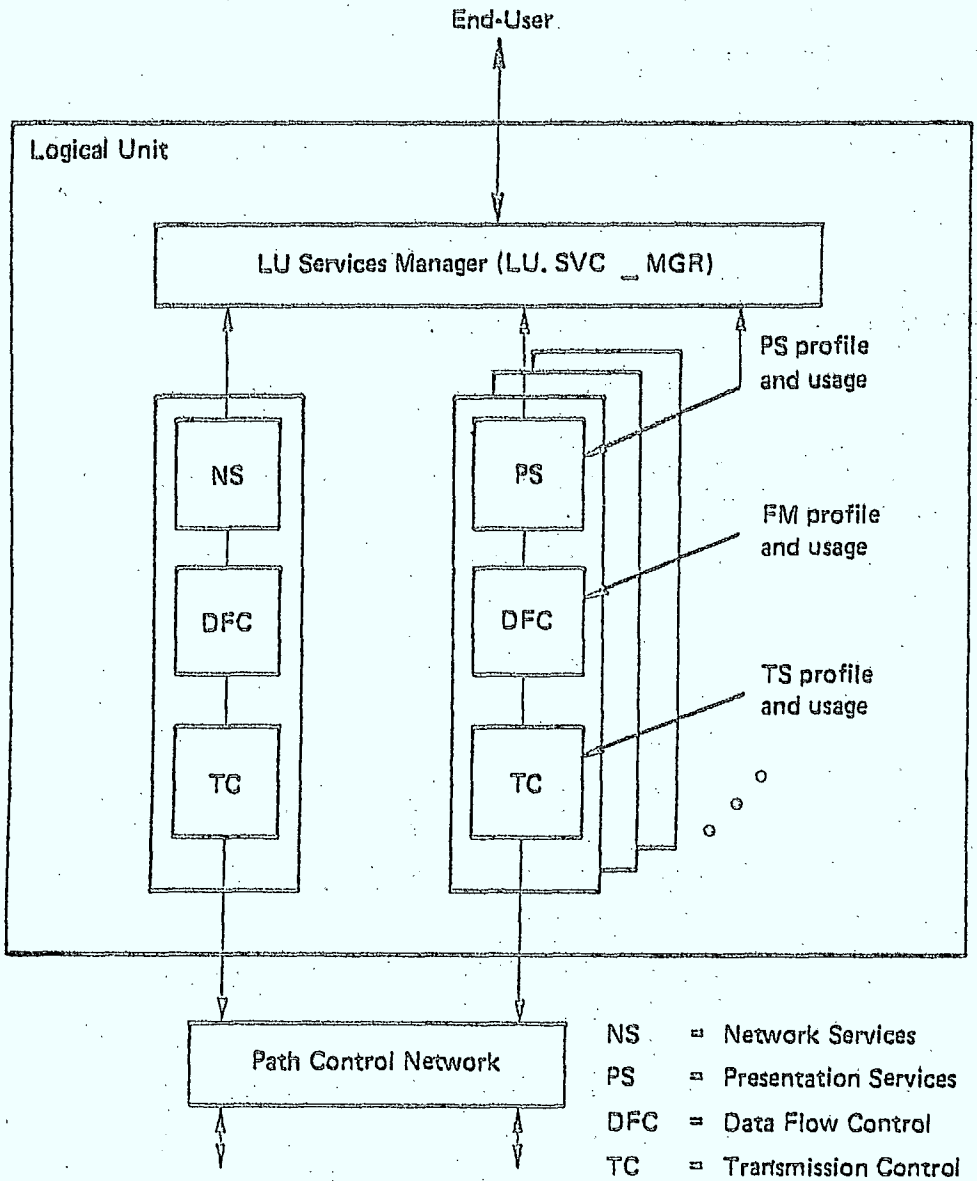

Half-Session Component	Selectable Functions (not a complete list)	BIND Fields Used
Presentation Services (PS)	◦ SNA Character String Usage ◦ Code Repertoire (EBCDIC, ASCII, other) ◦ Attended or Unattended Mode ◦ FM Header Usage	PS Profile and Usage

2.2.2.2 End User Services

The three types of end-user services as described in the ISO section are:

- virtual terminal service
- virtual file service
- job transfer and manipulation service

The SNA support for these services is discussed in more detail in the following sections.

2.2.2.2.1 Virtual Terminal Service

The classes of terminals are defined within SNA by different logical unit types, each of which defines a set of presentation services, and also a subset of the session functions that will be required. The LU type defines the presentation services profile, identifying the session type. Several different LU types have been defined within SNA, but the architecture does not restrict the number of LU types that may eventually be specified. Each LU may support one or more session types - i.e. a logical unit may be able to be classified as more than one LU type, and may therefore support one or more session types.

The correspondence with functions defined in the ISO Model is as follows:

(a) Selection of Session Type

The session type is selected (or pre-selected) when the session is established. The functions available within a particular session type may be further restricted during session establishment.

(b) Negotiation of Profile

The presentation services profile is defined by the LU type, but further refinement of this profile takes place. At session establishment the presentation services usage field describes the use of control headers and defines the data streams for the session. During the session the control headers are used to select presentation options within the profile.

(c) Data and Command Transfer

The transfer of data and commands is governed by the presentation services profile that has been selected. Different session types use different subsets of the total function list.

(d) Forms Management

Once a device has been selected using a control header, the forms management on that device is controlled using other control headers. This includes:

- forms mount
- electronic forms control load
- train mount
- copy functions

  which are features of many line printers.

For text processing, device control codes can be made common for the network. The presentation services FI.FMD scans each RU to locate the control codes and then processes them. It is possible to control line formatting, word hyphenation, page formatting, etc. remotely. It is also possible to leave all formatting to be performed by the receiving presentation services FI.FMD, thus permitting different numbers of characters per line and lines per page in different devices. Whichever distribution of control is adopted, it is carried out within presentation services.

For any output device an address space is used for the creation of the image. A degree of device independence can be obtained by thinking in terms of a logical presentation space that is not tied to physical size or characteristics; the same logical presentation space can then represent different types of devices. The logical presentation space must have a defined extent in each dimension but does not define such things as character size, line spacing, or page size. Because of this factor, a given logical presentation space can be transferred to physical images of different sizes on different devices. Figure 2.2.2.2.1-1 illustrates a simple logical presentation space for alphamerics; the co-ordinates are columns and lines rather than units of measurement.

Different devices may use the same, or a different, logical presentation space. They might use the same type of logical presentation space with different extents, or possibly a combination of different logical presentation spaces.

Presentation services would first map a given RU chain to a logical presentation space suitable for a class of devices, and then apply a final transform to suit the particular dimensions of one device. Paging and scrolling could be used to map a large logical presentation space onto a smaller physical presentation space.

Fig. 2.2.2.2.1-1 SIMPLE PRESENTATION SPACE

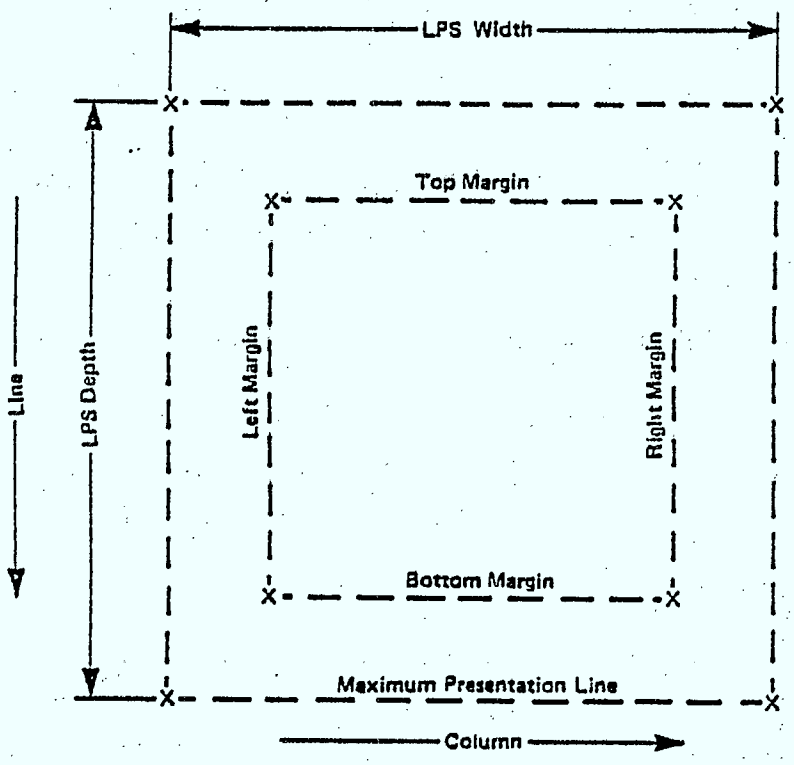

Alphameric logical presentation space. in rows and columns.

The transformation of an RU chain to a logical presentation space can be implemented at either end of the session. The sending presentation services FI.FMD may either:

- transform the field-formatted application data to a byte string with embedded field delimiters adding formatting control characters where necessary, or

- transmit a field-formatted RU chain to the destination presentation services FI.FMD, together with a header identifying the format of the data.

In the first case the receiving presentation services FI.FMD writes the data to the presentation space whereas in the second case the receiving presentation services FI.FMD must have a predefined format transform. If processing power is available at the remote site, the second alternative is more attractive because

- the data transmitted is reduced
- response time may be improved
- host cycles are off-loaded
- host application programs and application subsystems become more independent of the characteristics of the devices to which data is sent.

(e) Control of Operation

The type of operation is dependent upon the LU type and may be primary-secondary or peer-peer. In most cases control headers are used to control the communication states.

2.2.2.2.2 Virtual File Service

Access to remote data objects is provided in SNA by the use of an application program interface, which hides the location of the data objects from the application. The application's request for access to a data object is shipped to the site of the data object, where it is executed. This procedure is shown in figure 2.2.2.2.2-1. The data management operations required can be performed using control headers. permitted data set organizations include:

- sequential, providing sequential access to all records

- addressed direct, a relative data set using a record number for retrieval

- keyed direct, a relative data set using a key for retrieval

- keyed indexed, an indexed data set using a key for retrieval

An LU uses control headers to create and modify these data sets. Changes can be performed on a record or a group of records. In addition, the use of several control headers allow the LU to send and receive status information about the data set.

Fig. 2.2.2.2.2-1 VIRTUAL FILE SERVICE

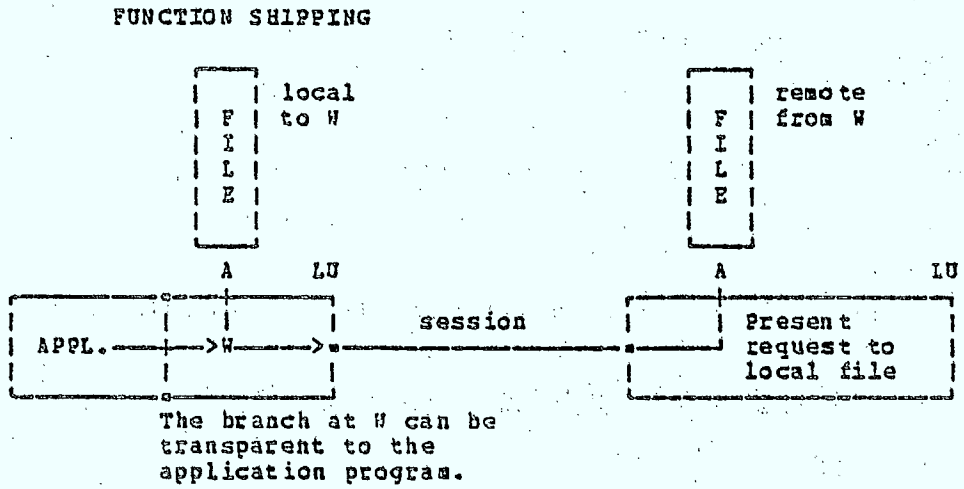

In some cases more than one control header can be used to perform an operation. One header can define the operation to be performed, while another tells where or how to perform it.

The correspondence with functions defined in the ISO Model is as follows:

(a)&(b) Encoding/Decoding Internal/External Attributes

This information is passed between LU's in the form of status information.

(c) Formatting Virtual Filestore Commands

The file commands are formatted according to standard rules. These commands are independent of particular implementations and presentation services will, if necessary, provide the translation between these common forms and forms that are understood by particular end users.

(d) Communication of File Data and Commands

The commands are transmitted in standard formats. The data will be transmitted according to the file organization in a standard form. This is dependent upon the data stream profile defined in the control header.

2.2.2.2.3 Job Transfer and Manipulation Service

Complete remote job entry services are provided which include:

- Translation of multiple I/O streams at the same time. These are transmitted on multiple sessions, one session per active stream.

- Establishment of sessions directly between the batch applications RJE spooling sub-system (an LU) and the remote job entry terminal. This avoids the processor overhead required to use spool to spool transfers to connect the terminal and the application's spool LU through intermediate spool file LU's.

- When the terminal supports it, multiple copies of the output can be created from a single transmitted copy. This can proceed automatically, under the control of the copies parameter on the processor's job control statements, or, it can be done under the control of the terminal operator.

Figure 2.2.2.2.3-1 shows how job entry systems can link together in an SNA network.

The correspondence with functions defined in the ISO Model is as follows:

(a) Control of Record Structures and Devices

Device control and record structure commands are sent using control headers.

Fig. 2.2.2.2.3-1 JOB ENTRY SYSTEMS NETWORK

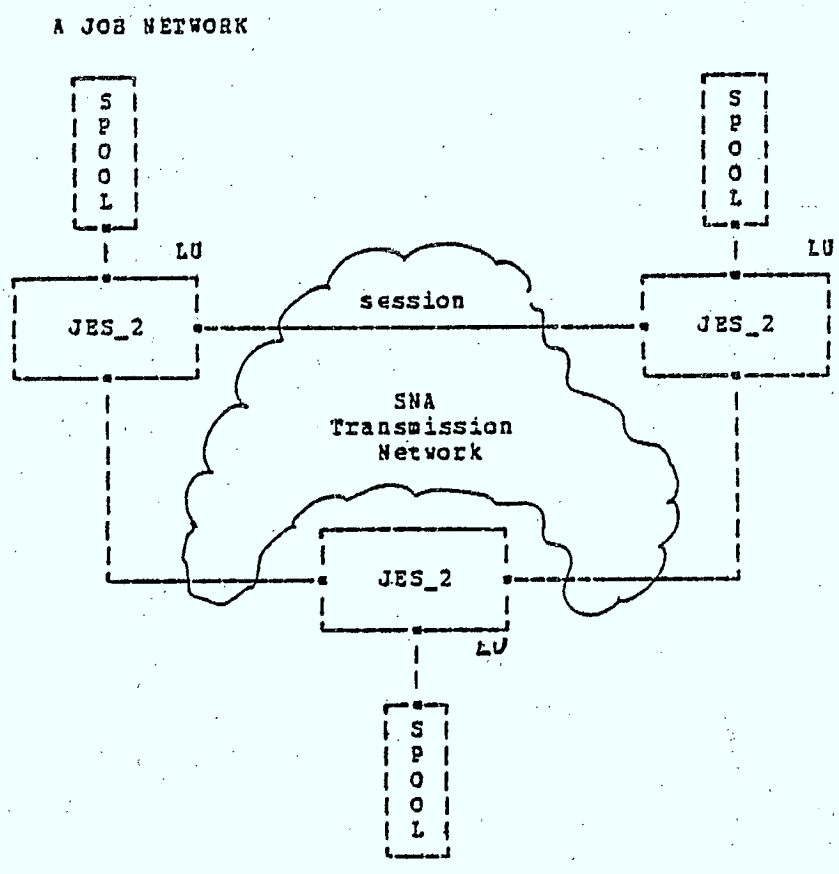

(b) Command Formatting

Management, access and job commands are converted to standard format control headers. This is a special case of file transfer.

(c) Data Formatting

The data will be converted to a standard structure before it is transmitted. This is a special case of file transfer.

2.2.2.3 Functions at the Session Layer Boundary

SNA does not define the protocol for the interface between the function interpreters for function management data (FI.FMD) and the data flow control/transmission control layer (DFC/TC), as this may vary between different product implementations. The FI.FMDs (corresponding to the ISO Presentation Layer) must however pass certain information to the DFC/TC's (corresponding to the ISO Session Layer) that defines those functions required from the lower levels that are visible to the FI.FMD's. This information may be from the NAU Services Manager (corresponding to the ISO Application Layer) or from the FI.FMD's themselves.

The selection of a presentation profile, or LU type, in itself restricts the functions that may be used at the lower levels. Usage fields further qualify the functions that are not part of the profiles.

2.2.2.3.1 Session Establishment

The session establishment request is made by the end user. The NAU Services Manager passes the request to the half-session, which corresponds to the Presentation and Session Layers of the ISO Model. The FI.FMD's supply a presentation services profile which, if acceptable to the correspondent partner, will define the presentation services to be used. This is passed to the DFC/TC layer and is used as input to the layer's establishment procedure, as the presentation profile restricts the choice of session services. The DFC/TC layer then supplies its own profiles, to correspond with the presentation profile, and to be accepted by the session partner.

2.2.2.3.2 Class of Service Requests

Network users are able to select a class of service for each session as it is started. The class of service name is used to assign performance parameters to the session. These parameters determine the transmission priority for the session, the availability of the network and some of the security provided for the session.

The transmission priority will identify a particular transmission group, which may consist of one or more virtual routes through the network. The transmission group scheduling algorithm detects errors that occur on a given line in the group and copies affected packets over onto another line in the group. This copy-on-error algorithm ensures that the transmission group will successfully transmit all packets routed over it, as long as at least one of its lines remains in service. The duplicate packets that can be created by this copy operation are detected and discarded at the receiving side of the transmission group.

The availability of the network is increased by the selection of alternate routes with the class of service name, during session establishment and at session restart. Sessions that do not require increased availability can use a class of service that does not specify any alternate routes, or, that specifies fewer alternate routes than are actually configured.

Security for the session may be enhanced by the selection of a class of service that causes the session to be assigned to physically secure routes.

2.2.2.3.3 Optional Facilities Requests

Any optional facilities available within the lower layers that are visible to the FI.FMD's and the NAU Services Manager will be requested by the application when required. This request will be passed to the DFT/TC layer as control parameters for the establishment of the session.

2.2.2.3.4 Session Termination Requests

Normal session termination is requested by the end user. The NAU Services Manager passes the request to the half session which carries out the appropriate termination procedures. The FI.FMD's will pass to the DFC/TC layer the termination request which includes indicators to show what type of termination is required - orderly, immediate etc.

2.2.2.3.5 Control Data

Commands generated by the FI.FMD's are passed to the DFC/TC layer in the form of control RUs, with no control headers, containing only the command.

2.2.2.3.6 Transformed Data

Data output from the FI.FMD's, for transmission, is in the form of RUs which may contain control headers if required by the presentation profile used.

2.2.2.3.7 Data Transfer Mode Request

The data transfer mode is requested at session establishment and may be defined by the presentation profile.

2.2.2.3.8 Flow Type Request

Data requests, some command requests and their respective responses are sent by normal flow between half sessions. Selected command requests and their responses use expedited flow between half sessions.

The FI.FMD's must identify to the DFC/TC layer which of the RUs that are to be transmitted require the expedited flow.

2.2.3 DEC - DNA Functions and Protocol

Functions within the Presentation Layer of the ISO Model are performed by the Dialogue Layer of DNA using the Data Access Protocol (DAP). DAP provides standardized formats and procedures for accessing and passing data between a user process and a file system existing in a network environment.

2.2.3.1 Functions Within the Layer

The functions provided by DAP that are defined by the ISO Model as belonging to the Presentation Layer are as follows:

- session establishment
- selection of initial presentation options
- renegotiation of options during a session
- transformation of transmitted data
- purge of session
- session recovery
- provision and use of an expedited flow
- session termination

These functions are described in more detail in the following pages.

2.2.3.1.1 Session Establishment

The originating process issues a connect initiate command requesting the creation of a logical link to the DAP process at the destination node. This request may specify an actual process name or the generic DAP object type.

Following link establishment the DAP processes exchange configuration messages to determine the initial presentation options to be used. This conversation also establishes system type and generic system capabilities.

2.2.3.1.2 Selection of Initial Presentation Options

The presentation options are selected by identifying system type, protocol version and generic capabilities:

- system type, which is used when it is necessary to know the type of both the operating system and the file system at the other end of the link.

- protocol version, which is used to identify the DAP version used at each end of the link.

- generic capabilities, which is used to determine the type of file support offered by a remote system without resorting to trial and error techiques.

2.2.3.1.3 Renegotiation of Options During a Session

The configuration messages exchanged after link establishment, to select the initial presentation options, may be exchanged in the middle of a session. This is done to start a different access when different presentation options are required from those currently being used.

2.2.3.1.4 Transformation of Transmitted Data

DNA supports several different record formats and attributes, but not all systems allow all of them. Transferring data between systems using different record formats and attributes will sometimes require conversions. All conversions must be done by the accessing (user) process and not by the accessed (server) process.

Any conversions not supported by DAP must be carried out by the application process.

2.2.3.1.5 Purge of Session

Purge of session may be initiated, by either dialogue process, by issuing an abort command which will always complete successfully.

2.2.3.1.6 Session Recovery

When an error is detected, recovery functions are available to restart the session from the previous correct message.

2.2.3.1.7 Provision and Use of an Expedited Flow

Any interrupts generated by DAP, or by the end user, are transmitted via the interrupt sub-channel, corresponding to the expedited flow defined by the ISO Model.

2.2.3.1.8 Session Termination

For normal session termination the data is first terminated and then the disconnect request is issued.

2.2.3.2 End User Services

The three types of end user services considered in the ISO Model description are:

- virtual terminal service
- virtual file service
- job transfer and manipulation service

The data access protocol version 4.1 is specifically designed for remote file access via a file system resident in the remote processor. Unit record devices and terminals can be accessed if supported by a file system. When unit record devices and terminals are supported by a file system in a device-independent manner, the device control features are not supported by DAP.

As all nodes of a DNA network are processors, the only way a terminal can use the network is through an application, or control program, in its local processor. The terminal itself does not support DNA. The virtual terminal service, as described by the ISO Model, is not provided by DNA.

The virtual file service and the job transfer and manipulation service, as described by the ISO Model, are considered in the following sections.

2.2.3.2.1 Virtual File Service

The data access protocol permits remote file access independently of the I/O structure of the operating system being accessed.

Data access protocol performs the following functions defined within the Presentation Layer of the ISO Model.

(a)&(b) Encoding/Decoding Internal/External Attributes

Information concerning the internal and external attributes of a file (as defined by the ISO Model) are exchanged using attributes messages. These provide information on how data is being structured in the file being accessed. It includes information on file organization, data type, format, record attributes, record length, size, device characteristics, and security.

(c) Formatting Virtual Filestore Commands

All remote file access is peformed using standard control messages. These define a set of functions that are independent of particular systems. The data access protocol will convert file commands to these standard formats.

(d) Communication of File Data and Commands

The commands are transmitted in the standard formats. Data is transmitted in data messages which are sent along data streams established using the control commands. Transmitted data, within the data message, may be in any format subject to the general file organization constraints imposed by the data access protocol.

2.2.3.2.2 Job Transfer and Manipulation Service

The data access protocol includes commands for the transfer and submission of files to a batch processing facility or command interpreter.

A command file will be transmitted as a temporary file, together with an instruction that informs the receiving system to submit the temporary file to the batch job processing facility on access completion, i.e. after the whole file has been transmitted and stored successfully. The file may be deleted, following execution by the batch facility, if required. The JCL within the command file must be that of the remote processor.

The functions defined within the ISO Model Presentation Layer, to support the job transfer and manipulation service, are not available within the data access protocol as the job transfer service provided by DNA is a special case of file transfer.

2.2.3.3 Functions at the Session Layer Boundary

The data access protocol must provide the network services protocol (NSP) with information it requires to complete the transmission. The types of data passed across this boundary are described in the following sections, categorized as within the ISO Model description.

2.2.3.3.1 Session Establishment Requests and Responses

The end user initiates the session establishment. The data access protocol converts the session establishment request into a link establishment and configuration dialogue. Only link establishment information is visible to NSP. The configuration messages are passed to NSP as data messages.

2.2.3.3.2 Session Termination Requests

The end user initiates the session termination. The data access protocol will normally convert the session termination request into an access completion and link termination dialogue. Only the link termination dialogue information is visible to NSP. The access completion messages are passed to NSP as data messages.

2.2.3.3.3 Control Data

Control information generated by DAP is sent either as control messages, which are passed to NSP as data messages for normal transmission, or as interrupts, which are visible to NSP as they are sent on the interrupt sub-channel.

2.2.3.3.4 Transformed Data

Any data for transmission will have passed through any conversion routines required at the sending node before passing to the NSP. If a conversion is required at the receiving node, it will be done after it has passed from NSP to DAP.

2.2.3.3.5 Flow Type Requests

The data transfer mode required for a particular message is predefined by the message type. The message type information is passed from DAP to NSP.

2.3 Session Layer

Executive Summary

The combined functions of the data flow control and transmission control layers of SNA have been found to correspond closely to the functions defined within the Session Layer of the ISO Model. In particular the following correspondences could be made:

- dialogue management functions, defined by the ISO Model, are provided by data flow control in SNA. The full-duplex and half-duplex states described in the ISO Model are provided. The special requirements for half-duplex communication are provided in SNA by the use of two types of half-duplex mode. The one way dialogue described by the ISO Model is, in SNA, a special case of half-duplex.

- data delimiting functions, defined by the ISO Model, are provided by data flow control in SNA. The basic information unit defined within SNA corresponds to the session service data unit described by the ISO Model. The quarantine service described by the ISO Model is provided by the data flow control of SNA, using chains. Additional flow controls are provided on these chains by defining request and response modes. The session interaction unit defined by the ISO Model is equivalent to a bracket as defined by SNA. Other units defined by the ISO Model are not visible to this layer of SNA.

- session establishment and context management functions, defined by the ISO Model, correspond to similar functions provided by the function interpreter for session control, in the transmission control layer of SNA.

- session release functions, defined by the ISO Model, are provided by the function interpreter for session control, in the transmission control layer of SNA. These include normal session release, session abort and the ability to quiesce the session.

A direct correspondence cannot be made between the ISO Model and SNA for the following functions:

- Session recovery, Although SNA provides functions to set the session to a defined state, the use of checkpoints as defined by the ISO Model is not provided by SNA.

- Session mapping, The mapping of session connections onto transport connections as defined by the ISO Model is not fully supported by SNA. The specifics of SNA virtual routes are not available at this time.

- Session identification, The two identifiers defined by the ISO Model do not correspond directly to the session identifiers used by SNA.

The Session Layer functions defined by the ISO Model correspond to functions provided by the network services protocol (NSP) of DNA. In particular the following correspondences can be made:

- session identification, defined by the ISO Model, requires the provision of two identifiers; a session identifier, which is equivalent to the link identifier used by NSP, and a global identifier which is equivalent to the remote node identifier used by NSP.

- session release functions, defined by the ISO Model, are provided by NSP. These include normal termination and abort.

- data delimiting functions, defined by the ISO Model, are provided to some extent by NSP. The data segment, defined within DNA, corresponds to the session services data unit described by the ISO Model. The quarantine service, described by the ISO Model, is provided partly by the NSP layer and partly by the higher layers. The dialogue message of NSP, is equivalent to the quarantine unit defined by the ISO Model. Other units defined by the ISO Model are not visible to the NSP of DNA.

A direct correspondence cannot be made between the ISO Model and DNA for the following functions:

- Session Establishment and Context Management, Although NSP allows for some negotiation of session options, no provision is made for the negotiation of higher level protocols. This is because each node within a DNA network is a processor and will implement specific high level protocols that are not visible to NSP.

- Session Recovery Functions, Those provided by NSP are based on a segment acknowledgement scheme which allows a session to be restarted at the last valid segment transmitted. The defined state and checkpoints specified by the ISO Model provide a more complete recovery mechanism.

- Session Mapping, The mapping of session connections onto transport connections as defined by the ISO Model is not supported by DNA.

- Dialogue Management Functions, Those defined by the ISO Model are not supported by NSP. The NSP layer provides to the higher layer a full duplex data flow and provides a separate flow control for transmitter and receiver operations. Flow control is considered in more detail in the DNA sub-section.

The functions provided by SNA correspond more closely to those of the ISO Model than do those of DNA. The session identification provided by DNA is closer to that defined by the ISO Model than that used by SNA. Both SNA and DNA do not provide the complete mapping of session connections onto transport connections for data exchange.

Detailed descriptions of the functions provided by the ISO Model, and the functions and protocol of SNA and DNA, are contained in the following sub-sections.

2.3.1 ISO Model Description

2.3.1.1 Functions Within the Layer

The Session Layer provides session administration services and the control of data exchange.

The functions include:-

- session establishment
- context management
- session identification
- session recovery
- session release
- session mapping
- data delimiting
- dialogue management

These are described in more detail in the following pages.

2.3.1.1.1 Session Establishment

The establishment request is received from the higher layers. The Session Layer provides selection mechanisms to determine the initial session protocol options to be used during data transfer.

2.3.1.1.2 Context Management

Supports the negotiation of the high level protocol to be used at the Presentation Layer. This includes

- identification of high level protocols;
- negotiation of high level protocol including renegotiation during a session and the use of a default when no negotiation takes place;
- dynamic specification of session layer protocol attributes;
- authentication.

2.3.1.1.3 Session Identification

The session entity provides the presentation-entity with two identifiers:

- a session identifer which the presentation-entity uses to refer to the session during data transfer;
- a global identifier which is unique within the open system.

2.3.1.1.4 Session Recovery

The session recovery functions support the presentation-entities in re-establishing normal operation after a failure of session services. They include:

- functions to set the session to a defined state;
- functions to synchronise on the previous checkpoint.

2.3.1.1.5 Session Release

The session release functions allow the Presentation Layer to perform purge operations, involving session abort, and normal termination which releases the session in a non-destructive way.

2.3.1.1.6 Session Mapping

Involves the mapping of a session connection onto a transport connection. This may be done in several ways:-

- one transport connection supporting several consecutive sessions;
- several consecutive transport connections supporting one session;
- several simultaneous sessions on one transport connection;
- one session supported by several transport connections simultaneously.

Figure 2.3.1.1.6-1 diagrammatically represents these mappings.

2.3.1.1.7 Data Delimiting

There are several ways to group data for transmission between two presentation entities into data units. They include:-

- session service data unit, is the logical unit of user data exchanged between session users and the session service;

Fig. 2.3.1.1.6-1 MAPPINGS OF SESSION CONNECTIONS ONTO TRANSPORT CONNECTIONS

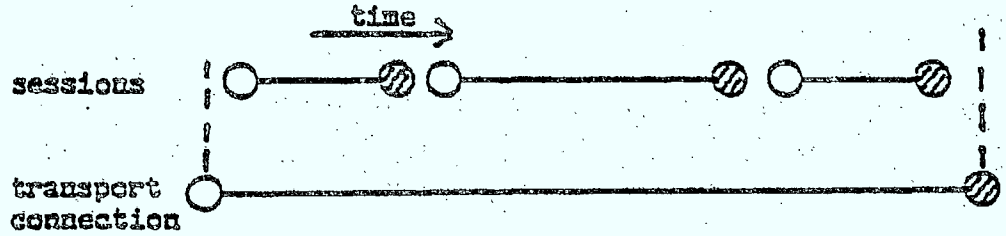

○ : Establishment

◍ : Release

Several consecutive sessions

This case will be useful when the cost or overhead of establishing and releasing transport-connections is high compared with those of sessions.

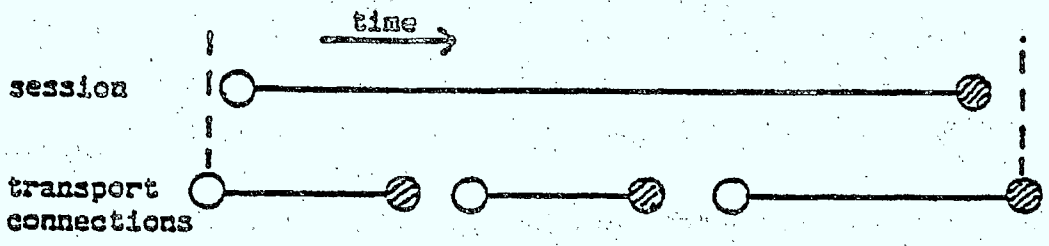

○ : Establishment

◍ : Release

Several consecutive transport-connections

This case could be useful when, within a session, there is a long period without data traffic, and when it is expensive to maintain the transport-connection.

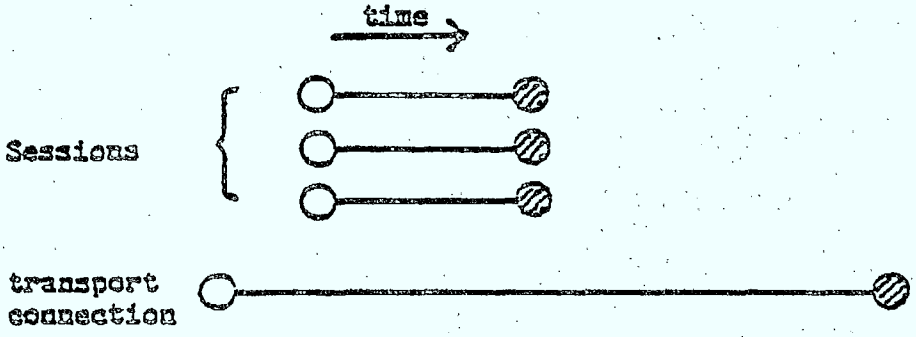

○ : Establishment

◍ : Release

Many to one mapping

- quarantine unit, consists of one or more session service data units which are only meaningful if all of them are available. If the Session Layer is providing the quarantine service, then the data units within the quarantine unit are not delivered to the Presentation Layer until all of them are available. If the quarantine service is provided by the Presentation Layer, then the Session Layer only provides recognition and transmission of quarantine unit delimiters.

- quarantine unit cancellation, allows for an incomplete quarantine unit to be ignored. If the Session Layer is providing the quarantine service, then the data will be purged without consequence to the Presentation Layer. If the quarantine service is left to the Presentation Layer, then the Session Layer simply transmits the cancellation request.

- session interaction unit, consists of one or more quarantine units, is used by presentation-entities to control whose turn it is to perform certain functions. The session interaction unit delimiter activates interaction control mechanisms to perform the transfer of turn in an agreed upon way.

- recovery unit, consists of one or more session interaction units, is defined and managed by presentation-entities for the purpose of recovery (rollback), which is supported by the session services.

- commitment unit, consists of one or more recovery units, is defined and managed by presentation-entities for the purpose of commitment. When a commitment unit is terminated no previous dialogue from the session is recoverable.

2.3.1.1.8 Dialogue Management

The Session Layer supports the following types of interaction:

- two way simultaneous, allows both presentation entities in the session to send and receive at the same time.

- two way alternate, allows the turn to alternate between the two partners, the three requirements for managing this interaction being:

  (i) when the session is established it is necessary to establish who has the first turn.

  (ii) When a change of turn is needed the sender must be able to indicate to the receiver that he may now send.

  (iii) The receiver must also be able to indicate to the sender that the receiver wishes to send. The turnover may take place in an orderly manner or on demand. The receiver may only send control information over the expedited data flow route.

- one way, allows one presentation entity to always be sending and the other to always be receiving.

2.3.1.2 Functions at the Presentation Layer Boundary

At the boundary of the Presentation Layer and the Session Layer, the Session Layer passes received information to the Presentation Layer. This includes the following:

- session establishment requests/responses;
- control data;
- identifiers;
- session failure/recovery information;
- session release information;
- data for transformation;
- data grouping information;
- data transfer mode indication.

These are dealt with in greater detail in the following sections.

2.3.1.2.1 Session Establishment Requests/Responses

The session establishment requests and responses are delivered to the receiving presentation-entity after passing through all the lower layers.

2.3.1.2.2 Control Data

Control information, for Presentation Layer functions, is delivered to the receiving presentation-entity after passing through all the lower layers.

2.3.1.2.3 Identifiers

The session-entity provides the presentation-entity with a session identifier, which is used to refer to the session during data transfer, and a global identifier, which is unique within the open system environment and is used for management purposes such as diagnostics or accounting.

2.3.1.2.4 Session Failure/Recovery Information

When session services fail the session recovery services are invoked. These support the presentation-entities in re-establishing the normal operation of the session.

2.3.1.2.5 Session Release Information

The Session Layer interacts with the Presentation Layer, during session release, to release the session in an orderly and non-destructive fashion. Session release information will also be provided to the Presentation Layer after an abort request has been actioned.

2.3.1.2.6 Data for Transformation

Data from the sending presentation-entity is passed to the receiving presentation-entity after passing through all the lower layers, this data will then be transformed by the half set of selected presentation services in the receiving presentation-entity.

2.3.1.2.7 Data Grouping Information

Quarantine unit delimiters will be passed to the Presentation Layer if it is providing the quarantine service. If the quarantine service is provided by the Session Layer, then this is transparent to the Presentation Layer.

Grouping information concerning interaction units, recovery units and commitment units will be passed to the Presentation Layer by the Session Layer.

2.3.1.2.8 Data Transfer Mode Indication

Information concerning the type of transfer mode to be used, and control information concerning that transfer mode, will be passed between the Session Layer and the Presentation Layer. The data transfer will take place, over the normal flow routes, according to the rules of the transfer mode being used, and control information will be sent along the expedited flow routes.

2.3.2 IBM - SNA Functions and Protocol

The Session Layer of the ISO Model corresponds to the Data Flow Control and Transmission Control (DFC/TC) layer of SNA.

2.3.2.1 Functions Within the Layer

The functions defined by the ISO Model as being within the Session Layer are performed by two layers within SNA. These are:

- Data Flow Control, which is involved in,
  - dialogue management
  - data delimiting
- Transmission Control, which is involved in,
  - session establishment
  - context management
  - session recovery
  - session release
  - data exchange
  - session identification

2.3.2.1.1 Data Flow Control (DFC)

The function of the data flow control layer is to control the flow of function management (FM) data requests and responses between function interpreters for function management data (FI.FMD) pairs within sessions. Data flow control handles only FM data and DFC requests; network control and session control requests do not flow through DFC. This means that DFC is not involved with session establishment, but relies on the FI.FMD to provide the parameters defined, during session establishment, for flow control. DFC is also not concerned with context management, session identification, session recovery (other than to request shutdown), session release (other than to quiesce normal data flow), or data exchange.

DFC can therefore be looked upon as a front-end process, within the DFC/TC layer (corresponding to the ISO Session Layer), which is only used during normal data transfer. This is shown in figure 2.3.2.1.1-1.

The functions provided by data flow control may be better described by first considering its structure. Figure 2.3.2.1.1-2 is a structural overview of data flow control.

Data flow control includes the following protocol machines:

- DFC.RCV, handles receiving of requests and responses;

- DFC.SEND, handles sending of requests and responses;

- OPEN-CHECK.Q-CPMGR-RCV, handles dequeuing of requests and responses from the Q-CPMGR-RCV queue in transmission control;

- SESSAD.DFC-INITIALIZE, initializes, at the activation of each session, information used by DFC (from the FI.FMD's);

- SESSAD.DFC-RESET, resets all DFC finite state machines and correlation tables;

- DFC FSMS, lie in the reset hierarchy;

Fig. 2.3.2.1.1-1 DATA FLOW CONTROL WITHIN THE SESSION LAYER

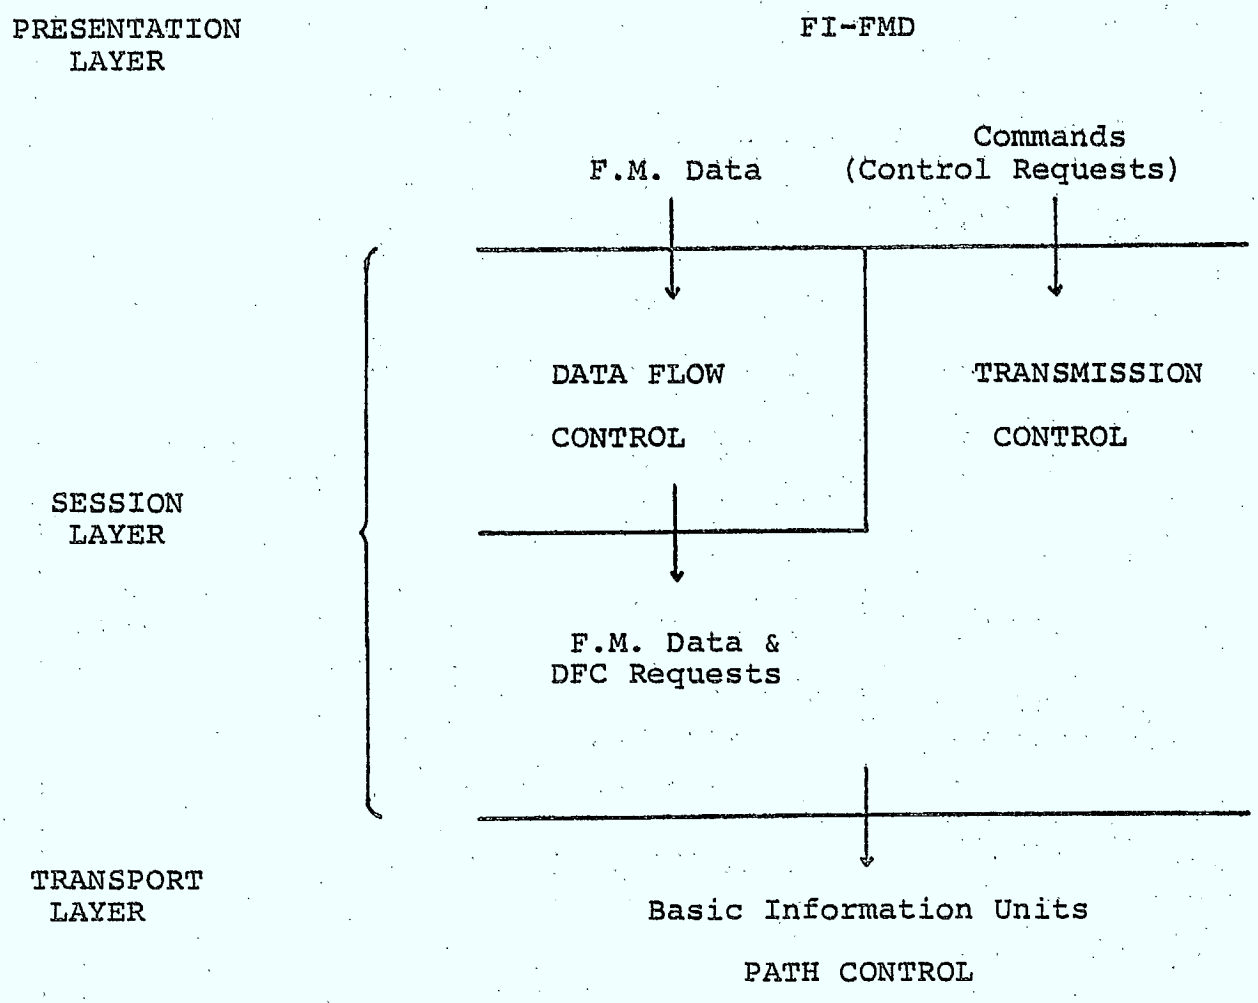

Fig. 2.3.2.1.1-2 DATA FLOW CONTROL STRUCTURE OVERVIEW

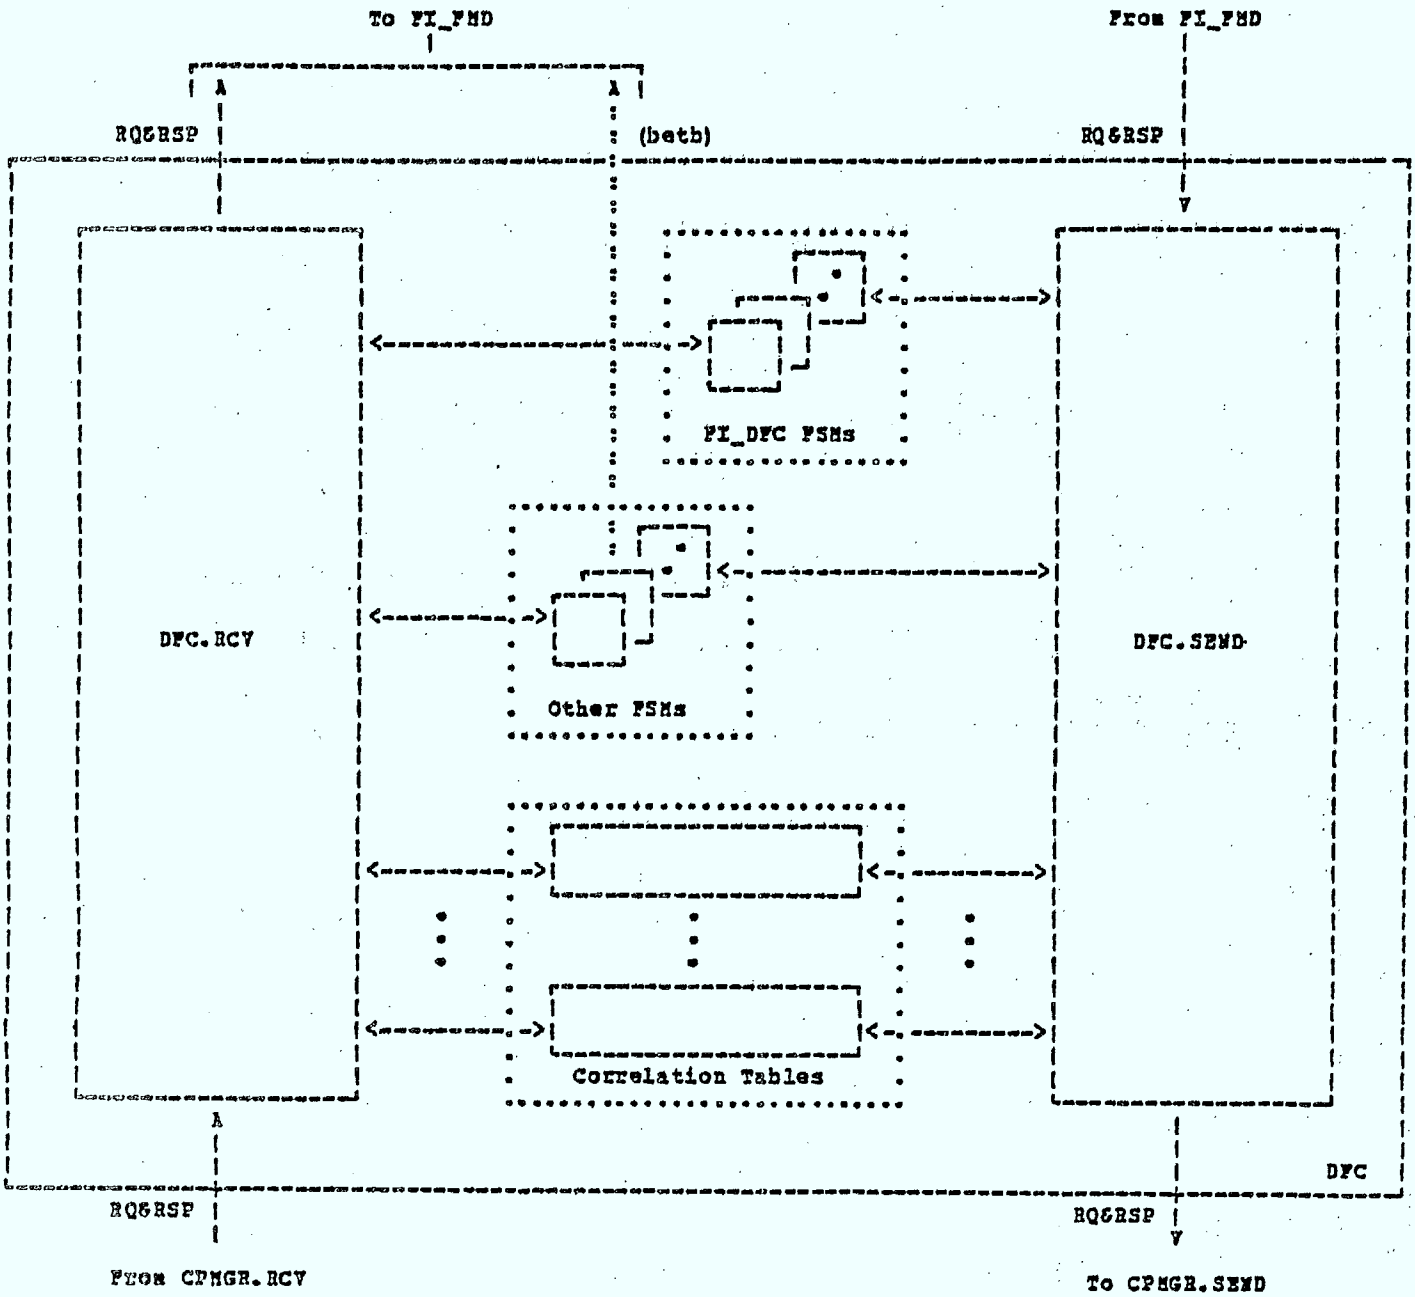

Structure of DFC

- CORRELATION TABLES, are used for correlating responses and requests. Enforcement of proper response sending and receiving, requires the use of these tables. The four tables used by DFC are for

  - expedited flow received requests
  - normal flow received requests
  - expedited flow send requests
  - normal flow send requests

  The receive correlation tables are used to correlate received requests with sent responses on the separate flows. The send correlation tables are used to correlate sent requests with received responses.

DFC controls the following:

- Dequeuing from the Q-CPMGR-RCV queue, when a Q-CPMGR-RCV queue exists, all received requests and responses (that are queued) are received from Transmission Control into the DFC layer by dequeuing them from that queue. This function is performed by the OPEN-CHECK.Q-CPMGR-RCV routine.

- Correlation of requests and responses, DFC correlates responses with their associated requests.

- Normal-flow send/receive modes, which are full-duplex, half-duplex contention, and half-duplex flip-flop specify a particular form of co-ordination between sending and receiving of normal-flow requests and responses. DFC checks that this is done correctly.

- Request and response control modes, immediate and delayed request/response modes are enforced by DFC.

- Chaining, which is enforced and checked to provide a means of sending or receiving a sequence of requests as an error recovery entity.

- Brackets, which are enforced to provide a means of sending or receiving a sequence of chains as a delimited transaction entity.

- Quiesce/shutdown processing, the normal-flow traffic may be suspended using various DFC requests; DFC enforces suspension rules following quiescing or shutdown of the normal flows.

- Request/response header formats, DFC enforces correct header parameter settings for data and commands.

The data flow control of SNA provides functions corresponding to the dialogue management and data delimiting functions of the ISO Session Layer. A detailed description of the functions and protocol sub-sets used is contained in the following pages:

(a) Dialogue Management

The DFC.SEND and DFC.RCV protocol boundary with FI.FMD will support full-duplex or half-duplex mode. One-way transmission is a special case of half-duplex mode in which the turn never changes from the first sender. The send/receive mode is determined during session establishment. It is completely independent of the transmission modes at the link control level.

(i) Full-Duplex (two-way simultaneous)

The primary and secondary half-session DFC.SEND and DFC.RCV protocol boundaries with FI.FMD are full-duplex. The normal-flow request and response flows in each direction are independent. Any correlation between flows is done at a level of control above that supplied by DFC.

(ii) Half-Duplex (two-way alternate)

The turn alternates between the two partners. Each half-session has a half-duplex DFC.SEND and DFC.RCV protocol boundary with FI.FMD. There are two ways of controlling the change of turn

- half-duplex flip-flop
- half-duplex contention

In half-duplex flip-flop mode there is some variation in the techniques depending on whether bracket formats are also being used.

Half-duplex flip-flop (not using bracket formats) - at session activation one half-session is designated first sender and the other first receiver. The sender issues normal-flow requests and the receiver issues responses. When the sender completes its transmission of normal-flow requests, it transfers control of sending to the other half-session by setting a change direction indicator on the last request sent.

Half-duplex flip-flop (using bracket formats) - at session activation one half-session is designated bidder and the other first speaker. Using bracket formats with half-duplex flip-flop mode requires a synchronization between the two half-sessions. When between brackets each half-session is in contention state - either may send. The contention winner is always the first speaker. When not between brackets, the half-sessions use the technique described for half-duplex flip-flop not using brackets. Brackets are further described in (b) Data Delimiting.

In half-duplex contention, at session activation, one half-session is designated the contention winner and the other the contention loser. Initially both winner and loser are in the contention state and either one may independently begin sending normal-flow requests. Normal-flow requests arriving at the loser, if it is sending, are queued. Normal-flow requests arriving at

the winner, if it is sending, may be temporarily queued or may be rejected with an appropriate negative response. Valid normal-flow requests, arriving at a non sending half-session, place the half-session in a receiving state. The contention loser reverts to contention state after sending or receiving the last request of a chain. Chains are further described in (b) Data Delimiting. On reverting to the contention state a contention loser, or a contention winner that queues received normal-flow requests, may dequeue any requests (and responses). Contention can be avoided by the use of a change direction indicator.

(b) Data Delimiting

Data flow control allows data to be grouped in several ways. They include:

(i) Basic information unit, which is the logical unit of user data exchanged between the half-sessions. It includes the request header which is generated by transmission control using parameters supplied by data flow control. It is equivalent to the session service data unit defined by the ISO Model.

(ii) Chaining, which provides a means to send (and receive) a sequence of requests as one entity, and to manage the recovery of the chain as a unit if necessary. A chain consists of a sequence of one or more requests with the following properties:

- the RUs belong to the same flow

- the RUs are all requests

- the first RU in the chain is marked with a 'begin chain' indicator

- the last RU in the chain is marked with an 'end of chain' indicator

- all other RUs in the chain are marked with a 'middle of chain' indicator

The proper chaining of requests is enforced for each half-session by DFC.SEND, and the checking of received requests for proper chaining is provided for each half-session by DFC.RCV.

Each response and each expedited request is a chain containing only one RU which is marked as both 'begin chain' and 'end of chain'.

Chains are further defined by the responses required for them:

- no-response chain, in which each request is marked no-response

- exception-response chain, in which each request is marked exception-response

- definite-response chain, in which the last request is marked definite-response.

The choice of the request/response mode for each RU in the chain is indicated by setting Form of Response Requested bits in the request/response header. The receiver of a chain only needs to examine these bits in the last request of a chain, or in a request in error. In addition, the Form of Response Requested bits, in the last request in a chain, need be examined only when the half-session activation parameter indicates that both definite response and exception-response chains may be received. If only one type of chain response will be received, then the setting of the Form of Response Requested bits on the last response in the chain may be assumed by the receiver without checking.

One of the data flow control functions is to police the adherence to the chain regulations and the alert the higher levels to any violations of these regulations. This includes checking to ensure

- proper sequence of 'begin chain', 'middle of chain' and 'end of chain'

- only the last element of a chain asks for a definite response

- only one definite response received for each definite response chain

SNA permits multiple codes for the definite-response and exception-response. This means that it is possible to distinguish between different types of definite-responses and different types of exception-responses. This would be implementation dependent.

To simplify implementation and to better manage error recovery situations, every half-session issues request and responses according to defined control mode options. The request control modes used on the normal flows are enforced by data flow control. The following request control modes are defined:

- immediate request mode, in which all request chains are sent under a single constraint, that no request may be sent on the flow by a given half-session when a previously sent definite-response request is still outstanding on that flow.

- delayed request mode, in which there are no constraints on the sending of chains.

The response control modes used on the normal flows are enforced by data flow control. The following response control modes are defined:

- immediate response mode, in which responses are sent in the order the requests are received, i.e. requests are processed and responses issued on a first-in, first-out basis. When a response to a particular request is received, it means that all requests in the same flow sent before the request that has been responded to, have been processed by the receiver and their responses, if any, have been sent.

- delayed response mode, in which responses may be sent in any order.

A chain, as defined within SNA, is equivalent to the quarantine unit specified in the ISO Model. The quarantine unit cancellation as defined by the ISO Model is equivalent to the CANCEL command in data flow control. This command is a single RU chain that is sent along the normal flow and terminates the previous partially sent chain without informing the FI.FMD (corresponding to the ISO Presentation Layer).

(iii) Bracket, which is a series of normal-flow request chains and their responses, exchanged in either or both directions between two half-sessions. The chains within a bracket comprise a unit of work. The bracket is delimited by the FI.FMDs but data flow control enforces the bracket exchange rules. These rules allow half-sessions to contend for activating a bracket, and assist in resolving the race condition that can result from that contention. The use of brackets is defined during session establishment.

A bracket is delimited by the use of a 'begin bracket' indicator in the first request of the first chain and an 'end bracket' indicator in the first request of the last chain.

If brackets are used in a session then at session establishment one of the half-sessions is specified as first speaker and the other as bidder. The first speaker has the freedom to begin a bracket without requesting permission to do so. The bidder must request and receive permission from the first speaker to begin a bracket.

One of the following bracket termination rules is specified for the session during session establishment.

- conditional termination, in which bracket termination is controlled by the form of response requested (definite, exception or no-response) for the chain containing 'end of bracket'.

  If the chain requests a definite response the bracket is not terminated until a positive response is processed. A negative response to the last request causes the bracket to be continued. A negative response to any but the last request in the chain allows the option of terminating or continuing the bracket. The sender of the chain may end the bracket by sending

CANCEL with 'end of bracket', or by ending the chain with a request specifying exception response or no-response. Alternatively the sender of the chain may continue the bracket by sending CANCEL without 'end of bracket', or by ending the chain with a request specifying definite response.

If the chain requests exception response or no-response, the bracket is terminated unconditionally when the last request of the chain that has 'end of bracket' in its first request is processed.

- unconditional termination, in which a bracket is terminated unconditionally when the last request of the chain that has 'end of bracket' in its first request is processed, regardless of the form of response requested.

No more than one 'begin bracket' and one 'end of bracket' can be outstanding from a half-session.

A bracket, as defined by SNA, is equivalent to the session interaction unit specified in the ISO Model.

Some of the data flow control options are selectable at session establishment. Specific combinations of these selectable options are known as function management (FM) profiles. These are further defined by the FM usage parameter settings. See figure 2.3.2.1.1-3 for a diagrammatic representation of the location of the FM profile and usage within the half-session.

2.3.2.1.2 Transmission Control

The transmission control elements provide three protocol machines for each locally supported half-session:

- connection point manager, controls sequence number checking, pacing, enciphering/deciphering, routing and other support functions relating to the half-session flows.

- function interpreter for session control, provides session-specific support for activating and deactivating half-sessions, and for starting, clearing and resynchronizing session-related data flows.

- function interpreter for network control, provides functions such as identifying sub-areas to which routing has been interrupted.

These protocol machines are interconnected as shown in figure 2.3.2.1.2-1. The functions provided are examined in more detail in the following pages:

Fig. 2.3.2.1.1-3 FUNCTION MANAGEMENT PROFILE AND USAGE

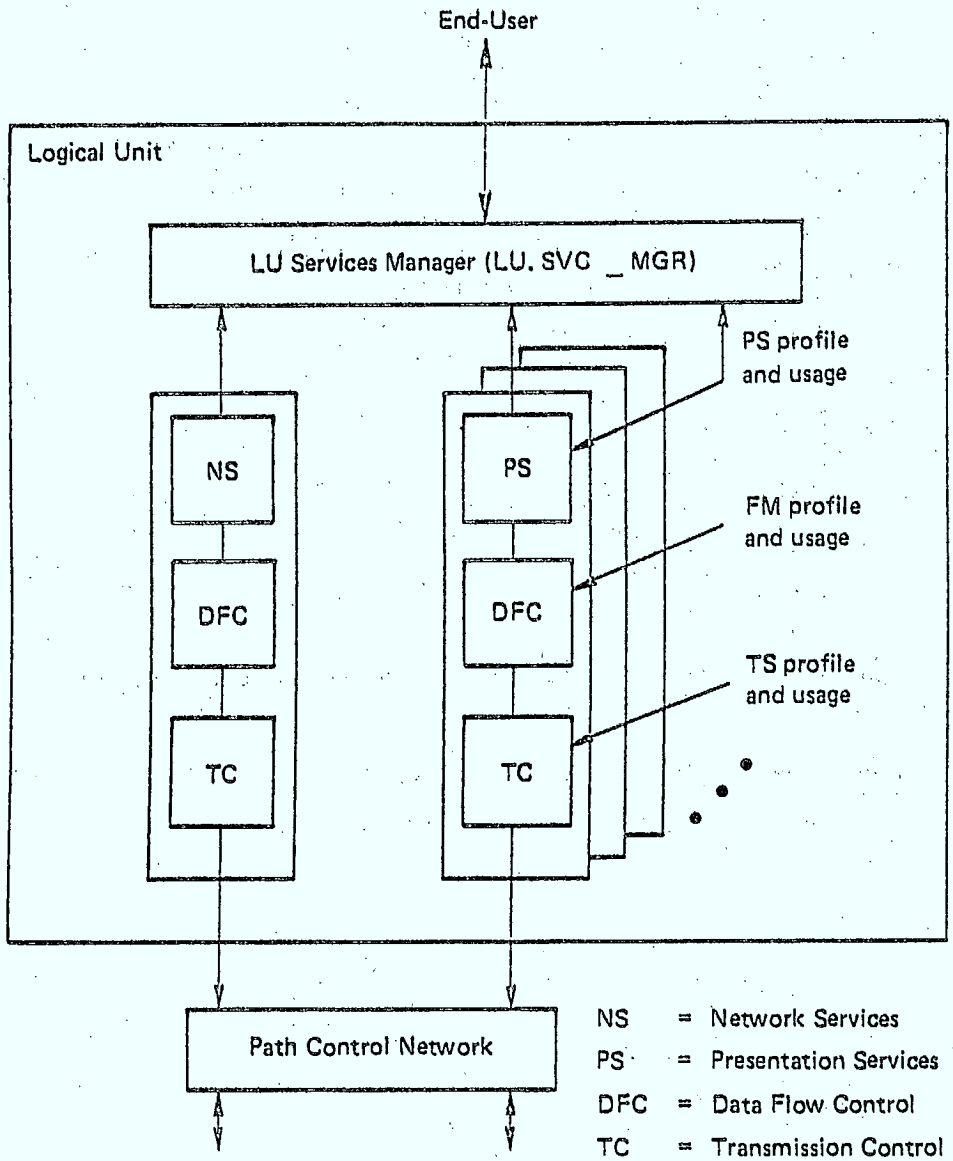

Half-Session Component	Selectable Functions (not a complete list)	BIND Fields Used
Data Flow Control (DFC)	◦ Request and Response Modes • Half-Duplex or Full-Duplex Send-Receive Mode ◦ Brackets and Chaining Rules • Data Flow Control Techniques and Requests	FM Profile and Usage

Fig. 2.3.2.1.2-1 TRANSMISSION CONTROL STRUCTURE OVERVIEW

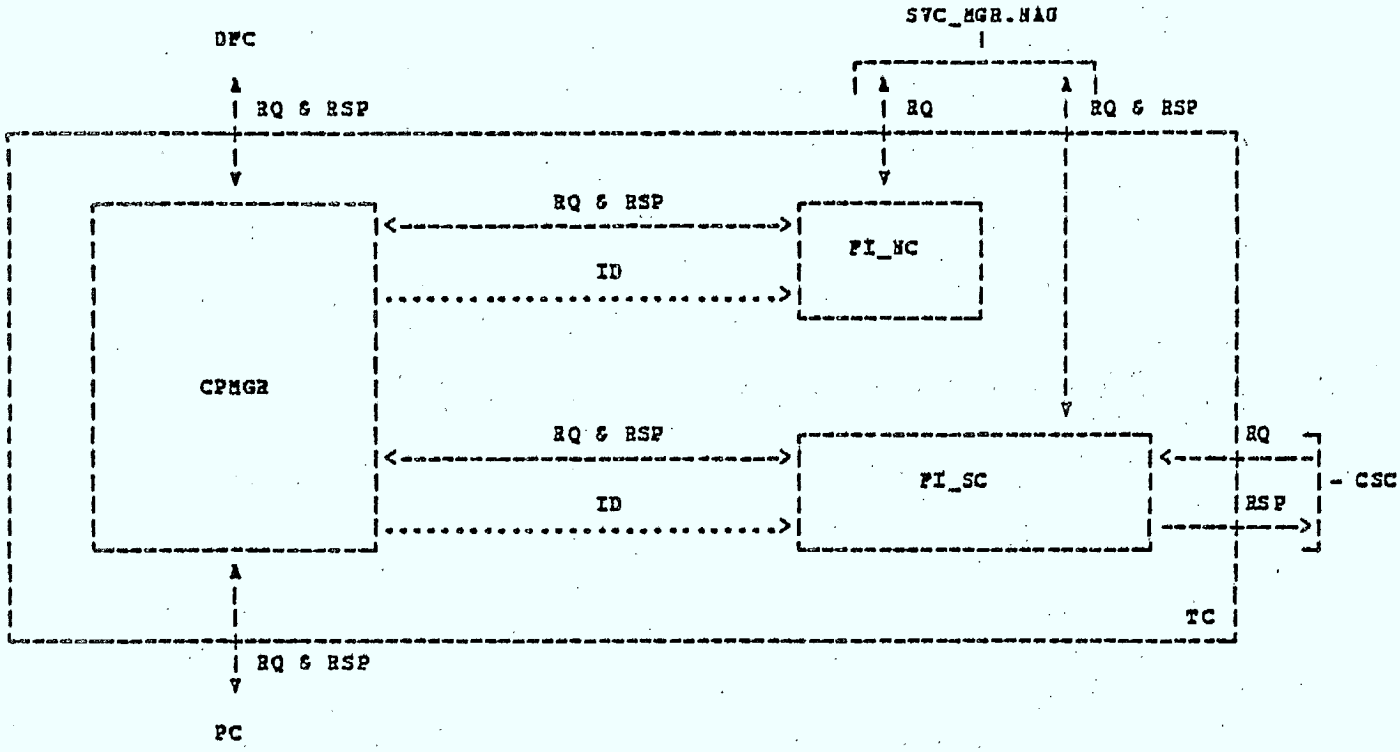

Structure of a TC element

(a) Session Control

Each function interpreter for session control supports functions related to session and data traffic activation, deactivation and recovery. All session control requests and responses are sent via the expedited flow.

(i) Session Establishment

It is the function interpreter for session control that manages the session establishment functions. This includes provision for defining the FM and TC profiles and usages, and also includes details of the presentation profile and usage to be passed to the FI.FMD. This function corresponds to the session establishment described by the ISO Model Session Layer.

(ii) Context Management

The high level protocols to be used are identified during session establishment by the presentation services profile. This field contains a format indicator and an LU type designation, which together determine the format and meaning of the presentation services usage field. Predefined LU types architecturally determine the following for the session:

- the mandatory and optional values allowed for the FM and TC (Session) profiles and usages.

- the usage of various controls, FM headers, parameters, etc.

- the presentation services options.

The presentation services usage field supplements the information specified by the session profiles and usages by identifying additional FM options.

Negotiation of high level protocols may be performed during session establishment to modify the session parameters in a limited way.

(iii) Session Recovery

Either session partner can be responsible for recovering from errors and unsuccessful transmissions. Session establishment parameters indicate which partner has the responsibility, and in some cases this responsibility may be shared.

Most half-sessions that have no session recovery responsibility begin recovery by sending, from session control, a request recovery command. The half-session data flow control could request shut down if it detects an error.

The session control of the half-session with recovery responsibility may then initiate clearing of all data traffic, cancelling and purging the current chain, or send a data RU (without data) that specifies end of chain.

Resynchronization after data traffic has been cleared is a function of the connection point manager and is detailed in (b).

If the session cannot be resumed it may be possible to start a new session where the previous session ended.

(iv) Session Release

The session release functions, provided by the function interpreter for session control, allow for normal end of session and temporary end of session. Abnormal termination commands are initiated in the higher layers and are supported by the session functions.

(b) Connection Point Manager

The connection point manager is responsible for the data exchange functions defined in the ISO Model Session Layer. SNA controls the flow of data by the use of pacing and sequence checking.

(i) Sequence Numbering

Sequence numbers, if used, are assigned by data flow control to each normal-flow request, and are checked in the receiving half-session by the connection point manager. For the expedited flow, an identifier is assigned to each request sent. Data flow control assigns the identifiers to data flow control requests, and the connection point manager assigns the identifiers for other requests.

If sequence numbers are not used on the normal flow, then identifiers will be used here as well.

Resynchronization during recovery, relies on the sequence numbers or identifiers to determine a point at which transmission may re-start, whether in the same session or a new session.

(ii) Pacing

Pacing allows a connection point manager to control the rate at which it receives requests on the normal flow. If the pacing option is to be used, it applies only to the normal flow and does not affect the expedited flow. It is normally used when the sending connection point manager is capable of sending requests faster than the receiving connection point manager can process them.

The sending connection point manager sends a limited number of requests and then waits until the receiving connection point manager indicates that it is again ready to receive before sending another group of requests.

(iii) Request/Response Modes on Expedited Flow

The connection point manager polices the control modes used for expedited flow requests and responses. Normally the expedited flow requests use immediate request mode - no request may be sent on the flow when a previous definite-response request is still outstanding; and delayed response mode - responses may be sent in any order.

(iv) Cryptography

Session level cryptography is an option provided by SNA. If it is selected, the connection point manager enciphers all data RUs being sent and deciphers all data RUs received. The function may be mandatory, in which all data RUs are enciphered, or selective, in which only those data RUs indicated are enciphered.

(c) Session Identification

Each half-session denoted by the half-session identifier, HSID, has the structure shown in figure 2.3.2.1.2-2. A session is uniquely identified by the pair of network addresses for the NAUs engaged in the session. This pair of addresses is called the session ID. The SID can also be used to reference all of the various elements that are part of a particular session.

The transmission control profile specifies which transmission sub-system facilities will be used while the session remains active. The information specified by the transmission control profile may be supplemented by that from the transmission control usage field, which is used to specify pacing parameters and maximum RU sizes on the normal flow. See figure 2.3.2.1.2-3 for a diagrammatic representation of the location of the transmission control profile and usage within the half-session.

Fig. 2.3.2.1.2-2 SESSION IDENTIFICATION

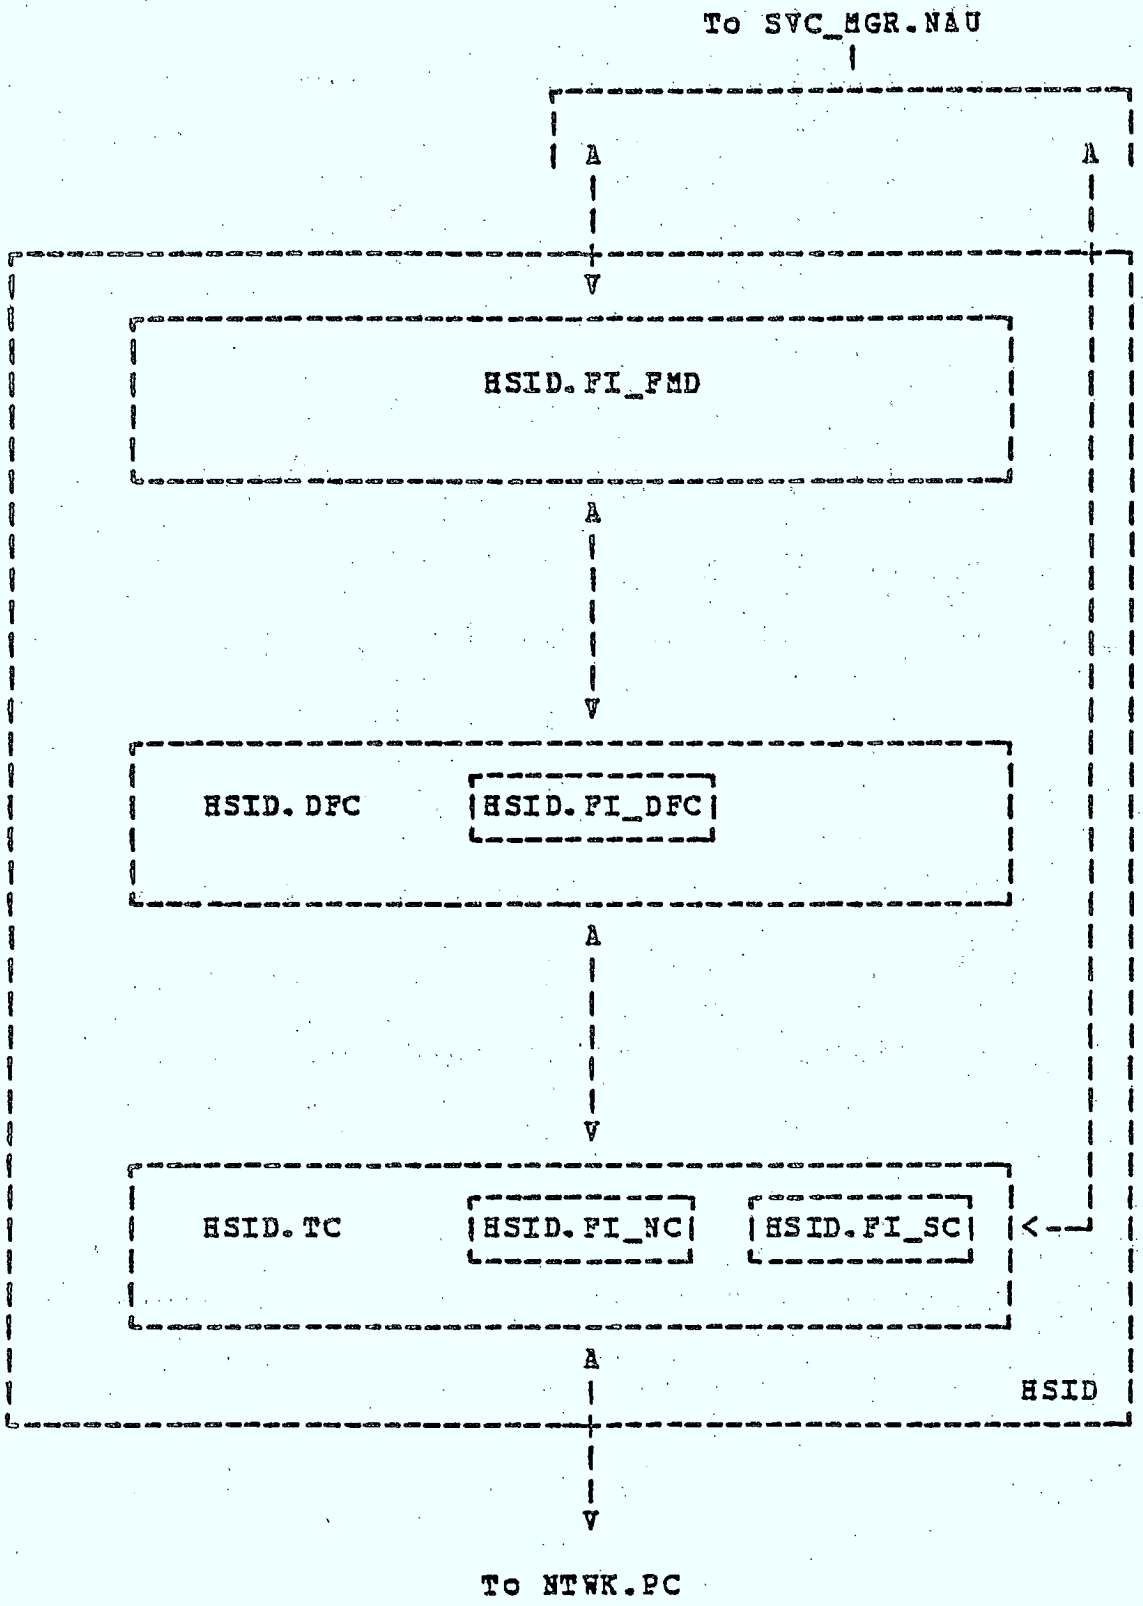

Half-Session Structure

Fig. 2.3.2.1.2-3 TRANSMISSION CONTROL PROFILE AND USAGE

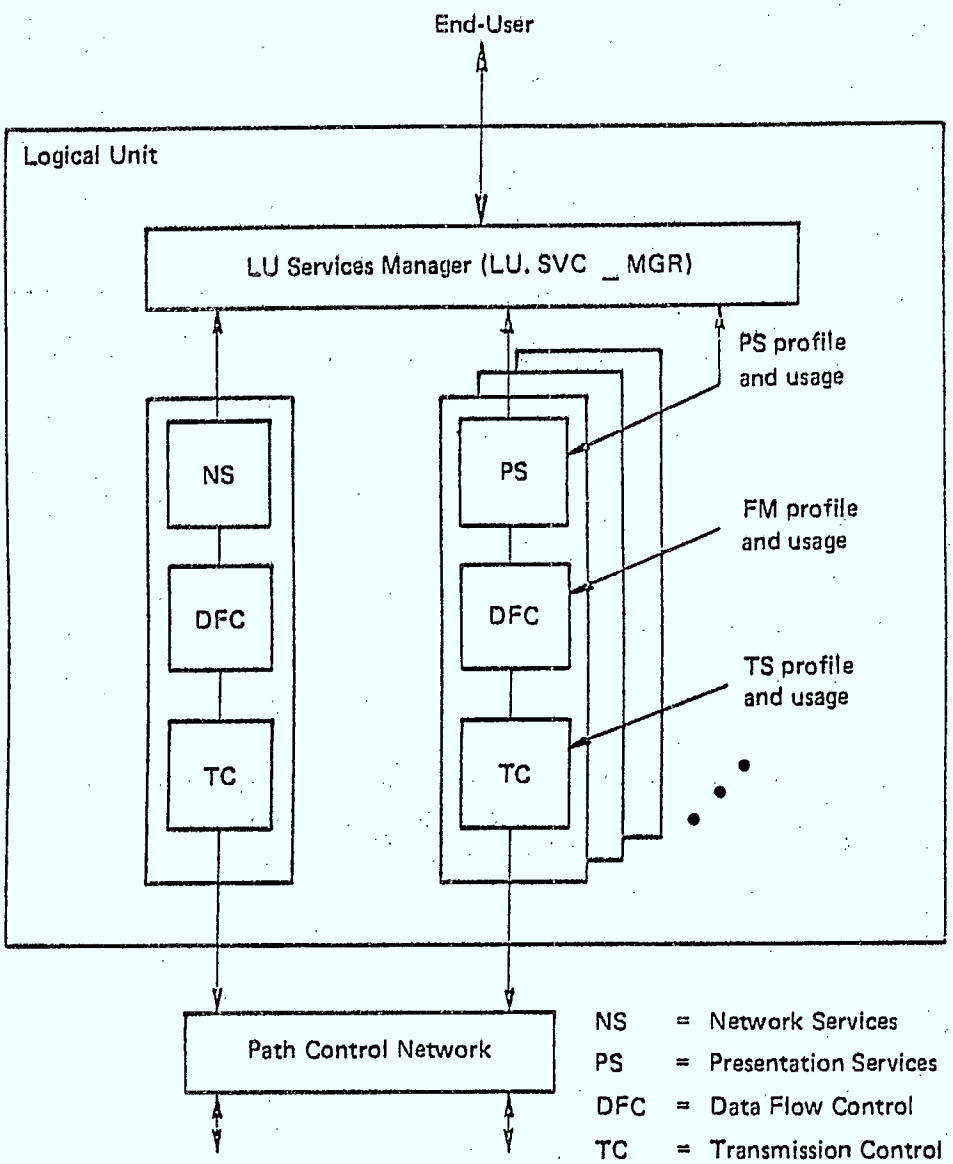

Half-Session Component	Selectable Functions (not a complete list)	BIND Fields Used
Transmission Control (TC)	• RU Size • Pacing • Session Control Requests • Cryptography Options	TS Profile and Usage

2.3.2.2 Functions at the Presentation Layer Boundary

Both data flow control and transmission control interface with the FI.FMD's (corresponding to the ISO Presentation Layer).

2.3.2.2.1 Data Flow Control

Data flow control only handles FM data and DFC requests. The functions performed that correspond to those described in the ISO Model include:

(a) Data for Transformation

The FM data, input by the sending FI.FMD is passed, by the receiving data flow control, to the receiving FI.FMD as complete chains after header information has been removed.

(b) Data Grouping Information

The FM data within a chain is only passed to the receiving FI.FMD when it is all available. No information regarding the chain is passed to the FI.FMD.

Data flow control informs the FI.FMD of bracket information as no more than one bracket may be used at a time.

Other data groupings are not defined within SNA.

(c) Data Transfer Mode Indication

All data passed to data flow control from the Presentation Layer is to be sent via normal flow.

2.3.2.2.2 Transmission Control

The interface with the FI.FMD's is handled by the function interpreters for session control and network control. The functions performed that correspond with those described in the ISO Model include:

(a) Session Establishment Requests/Responses

The session establishment requests and responses are passed from session control to the FI.FMD's.

(b) Control Data

The control information passed to the Presentation Layer consists of data from the RU received, together with the sequence number that has been allocated to the RU and any sense data generated. An indicator will also be passed to the FI.FMD's to show whether this is a request or a response.

(c) Identifiers

Session identification information is passed to the FI.FMD's to identify which session the particular RU belongs to.

(d) Session Failure/Recovery Information

Sense data, identifying the type of error that has occurred, is passed to the FI.FMD's if the error cannot be rectified within transmission control.

(e) Session Release Information

Session release is performed within session control and any information concerning the session release is passed to the FI.FMD's.

2.3.3 DEC - DNA Functions and Protocol

Functions within the Session Layer of the ISO Model are provided within the Logical Link Layer of DNA by the Network Services Protocol (NSP). NSP provides an interprocess communication mechanism among the nodes of a DECnet network.

2.3.3.1 Functions Within the Layer

The functions provided by NSP that are defined by the ISO Model as belonging to the Session Layer are provided by the logical link service within NSP. They are as follows:

- session establishment
- context management
- session identification
- session recovery
- session release
- data exchange
- data delimiting
- dialogue management

These functions are described in more detail in the following sections.

2.3.3.1.1 Session Establishment

The session establishment request is received from the higher layers. The session establishment routine requires that the source NSP sends a connect initiate message to the proposed partner. This is a control message which contains information to determine the initial session options to be used. The session establishment response is in the same format.

2.3.3.1.2 Context Management

Each node within the DNA network is currently defined as a processor, and therefore the identification of the high level protocol is assumed. Negotiation of initial options may take place during session establishment. The destination process must be defined during session establishment. Authentication of process access is carried out within the receiving NSP. Dynamic specification of some session options is performed.

2.3.3.1.3 Session Identification

NSP identifies a session or logical link in a node by a logical link address which is unique within that node. A particular link will have two addresses associated with it, one for each node in which the link terminates.

An NSP implementation must maintain the following information for each logical link that terminates in the node.

- local link address, The logical link address by which the implementation identifies a particular logical link - corresponding to the session identifier as described in the ISO Model.

- remote node identity, The identity of the node in which the other end of the logical link terminates - corresponding to the global identifier as described in the ISO Model.

- remote link address, The logical link address by which the NSP implementation in the remote node identifies a particular logical link.

2.3.3.1.4 Session Recovery

To ensure the integrity and cohesiveness of the information being exchanged, and to allow a receiving NSP to discard received segments, each NSP module employs a segment acknowledgement scheme. This scheme keeps track of the data segments sent (numbered by NSP) and ascertains whether or not retransmission is necessary. Resynchronization of the session will therefore be at the last acknowledged segment.

2.3.3.1.5 Session Release

At any time during a dialogue exchange, NSP will allow either partly to abort or terminate the conversation. When NSP has been properly notified to do so, it will disassemble or destroy the logical link connection. The two functions involved are:

- disconnect, is synchronous with previous data transmit requests. All data from the dialogue process must be sent and acknowledged before the disconnect message is sent. Equivalent to the non-destructive session termination defined by the ISO Model.

- abort, allows the operations resulting from a disconnect to be performed immediately without waiting for previous data transmit requests to complete. Equivalent to the session abort defined by the ISO Model.

2.3.3.1.6 Data Exchange

Data exchange is managed by NSP using flow control and segment number checking.

The receiver of a flow of data may stop the flow of normal data segments from the transmitter by sending a control message. Any segments received after this may be discarded by sending a negative acknowledgement. The flow may be started again by sending another control message. When dialogue messages are being received, the receiver may send a control message requesting one or more dialogue messages to be sent. This function may also be performed by requesting a number of segments to be sent, which may comprise part of one or many dialogue messages. Message and segment counts are maintained at both the transmitter and the receiver.

2.3.3.1.7 Data Delimiting

NSP groups user data as follows:

- segment, is the logical unit of user data exchanged via the network. The segment size is determined during session establishment. It is equivalent to the session service data unit defined by the ISO Model.

- dialogue message, is a unit of information that has meaning to the end user. This may be one or more segments. The dialogue message will be reconstructed by the receiving NSP prior to delivery. Start and end of message flags are used to delimit the dialogue message. It is equivalent to the quarantine unit defined by the ISO Model. Dialogue message cancellation does not occur within the NSP. The receiving dialogue process allocates a buffer for each dialogue message it wants to receive. Segments received (at the end of a dialogue message) for which there is no buffering space are positively acknowledged and then discarded. When the last segment of such a dialogue message is discarded, the receiving dialogue process (presentation-entity) is informed that the receive is complete but that data was lost.

2.3.3.1.8 Dialogue Management

Dialogue management within NSP is performed on a flow of data from a transmitting NSP to a receiving NSP. If data is also flowing in the other direction then the control of the interaction between the two flows will be controlled at a higher layer. The dialogue management is duplicated for each of the normal and interrupt flow sub-channels. It is described in terms of the receiver and transmitter operation.

(a) Receiver Operation

Segments received on the sub-channel are processed in sequence (by segment number). A segment out of order may be rejected or held. Each segment must be acknowledged positively or negatively within a specified time period. The positive acknowledgement of a segment is assumed to positively acknowledge any previous segments. The acknowledgement may be sent as a separate message, or included in a message travelling in the opposite direction.

(b) Transmitter Operation

The transmitter assigns unique sequential segment numbers to segments to be transmitted. The segment number must be associated with the particular segment until a positive acknowledgement is received. If the data is to be retransmitted it must have the same segment number as on the previous transmission. The transmitter must be able to process received acknowledgements out of order by assuming the highest number positive acknowledgement to acknowledge any lower number segments still active.

2.3.3.2 Functions and the Presentation Layer Boundary

At the boundary of the data access protocol (DAP) and the network services protocol (NSP), NSP passes received information to DAP. This includes:

- session establishment requests/responses
- control data
- identifiers
- session failure/recovery information
- session release information
- data for transformation
- data grouping information
- data transfer mode indication

These are discussed in more detail in the following sections.

2.3.3.2.1 Session Establishment Requests/Responses

The session establishment messages are delivered to the partner dialogue process after passing through the lower layers. The parameters exchanged between the dialogue processes and NSP include the following:

- destination name and process identification;
- source process identification;
- link identifiers (one for each end);
- buffer addresses;
- access control information;
- segment interface information.

2.3.3.2.2 Control Data

Control information for the dialogue process functions is passed to the receiving dialogue process after passing through the lower layers. The parameters exchanged between the dialogue processes and NSP include the following:

- buffer addresses
- segment and message information
- indication of success or failure (NSP to dialogue process)
- sense data on failure
- interrupt data (control commands)

2.3.3.2.3 Identifiers

The NSP provides the dialogue process with the local link identifier, the remote link identifier, the destination node name, the destination process identification and the local process identification. Link identifiers correspond to the session identifiers described in the ISO Model. The destination node name corresponds to the global identifier described in the ISO Model.

2.3.3.2.4 Session Failure/Recovery Information

NSP informs the dialogue process of the success or failure of each segment or message. A failure initiates retransmission of the failed segment. If the session has failed then a new link has to be established before retransmission.

2.3.3.2.5 Session Release Information

Disconnect and abort commands are passed from the dialogue layer to NSP for transmission. The reason for the disconnect or abort will be supplied by NSP to the dialogue layer.

2.3.3.2.6 Data for Transformation

The sending dialogue process passes the data to NSP by identifying the location of the data (buffer address). Other control information is also passed to NSP including identifiers etc. The received data messages are passed from NSP to the dialogue process. NSP will also pass on to each dialogue process positive or negative acknowledgements for each message.

2.3.2.2.7 Data Grouping Information

Message delimiters are passed from NSP to the dialogue process, together with an indication as to whether that message was transmitted successfully or not.

2.3.3.2.8 Data Transfer Mode Indication

There are two data streams on a logical link. One stream is used for interrupts and link service messages - equivalent to the expedited flow defined by the ISO Model, and the other contains segments of data messages - equivalent to the normal flow defined by the ISO Model.

NSP decides which stream is to be used by checking a message identifier supplied by the dialogue process.

On the interrupt and link service sub-channel all messages are single segments. On the data sub-channel message segmentation can occur. NSP segment acknowledgement is performed independently for each sub-channel, and NSP informs the dialogue process of the success or failure of each message transmission on each flow.

2.4 Transport Layer

Any proposed transport service must be able to support a wide variety of communication systems such as leased lines, packet switching services, circuit switching services and various local area networks. In addition it must support numerous types of applications incuding Teletex, Facsimile, Videotex, Directory Information, Plant Surveillance, and Message Handling facilities. To complicate the mandate even further, it must do this without becoming unduly complicated; in particular it must be applicable to small hardware units of limited capability, such as microprocessor controlled terminals.

The transport service will provide similar services to those provided by the Network Layer across a particular network (i.e. Network Layer Protocol). Network interconnection may be achieved easily only if networks provide the same elements of service. The transport service provides for a network connection across combinations of public and private networks. In some cases, where network services are identical, this type of connection extention can be realized at the Network Layer. Transport functions, then, enhance the quality of the network service (e.g. cost reduction by means of multiplexing, reliability by means of error recovery procedures, re-establishment of virtual circuits after failure, etc.). If all functions of the transport layer vanish then the transport layer service is then identical to the network layer service not only in nature but also in quality.

The characteristics of the network service is currently a contentious issue between ISO and CCITT. The two OSI models contain similar definitions of the services provided at the Transport/Session boundary, but differ significantly in the way this service is to be provided. The CCITT model assumes a consistent set of services at the Network Layer, similar to the services of the X.25 protocol. These services will be available over packet and non-packet networks, by means of an end-to-end X.25-like protocol in the latter case. The transport service may also have the function of providing connections over two or more networks not connected at the network level. The transport service functions are proposed to be individually selectable as required and may be completely absent if user requirements are met by the underlying network service.

The ISO model, on the other hand, assumes an end-to-end network connection with rudimentary services, significantly less than those of X.25. It also assumes (although this has been proposed for further study) that networks may differ in the services they offer. A transport protocol of considerable complexity (as yet unproposed) will be needed to provide the services required by the Session Layer. It is also unclear (also under further study) how an end-to-end connection across two dissimilar networks can be achieved, and how interconnectability is achieved between transport entities which support different sets of functions.

An end-to-end transport protocol which assumes a minimal network layer service can provide the required service. The weakness of the end-to-end protocol approach is that it requires all systems to implement an additional protocol above the network layer and introduces substantial software and communications overhead. The motivation for developing an end-to-end transport protocol stems from a concern that the three lower layers (i.e. physical, link, and network) cannot be structured in a manner that allows a consistent service to be provided, particularly when networks of different types are used. X.25, although normally associated with PDN access, is also applicable to applications operating over circuit switched networks. In fact, if applications are to internetwork over combinations of local, packet, circuit, etc., networks, then at some layer, a common level of service must be provided. In achieving this consistency at the network layer, it is important to note that the protocol may be based on X.25, or if only a single connection is required, then a "mini-protocol" directly above HDLC may be sufficient for providing virtual circuit service.

There exists valid views on both sides of the Transport definition debate, and the decision to use ISO's reference model as a basis of comparison for the following analysis, was for consistency only.

It has been used as a base for the ensuing discussion which compares and contrasts transport functions proposed by ISO, and the transport-like functions inherent in IBM's SNA, and Digital's DECNET.

EXECUTIVE SUMMARY

The following section details the transport functions of the ISO reference model, and SNA. It was shown that DECNET Phase II has no transport type functions since there is no equivalent to virtual circuits (routes) in DECNET. In DECNET Phase II there is a one-to-one mapping between sessions (logical links) and network connections. In OSI and SNA there may exist a many-to-one or one-to-many relationship between transport and network connections. Therefore, there is need for extra addressing, multiplexing logic, flow control, etc.

In particular, the direct correspondences between the ISO model and SNA are as follows:

- addressing, Both SNA and the ISO model must support address mapping functions at this layer. The mapping/switching functions have not been defined by ISO, and the mapping functions for SNA virtual routes are not yet available.

- multiplexing, ISO proposes both upward and downward multiplexing at the Transport Layer. SNA supports upward multiplexing only at the Path Control Layer. ISO also proposes multiplexing at the Session Layer. The practical utility of this is unclear, and the proposal may be associated with misunderstanding of X.25 functionalities. SNA assumes a one-to-one mapping of session and transport connections, which is in agreement with CCITT's reference model.

- segmenting and blocking, Both architectures support segmenting and blocking at the Transport Layer.

- end-to-end sequence control, error detection, error recovery. Both architectures support these virtual circuit functions.

2.4.1 ISO Model Description

2.4.1.1 Functions within the Layer

The transport layer is responsible for the provision of all functions which bridge the gap between the services provided by the Network Layer and the services needed by the Session Layer. The Transport layer functions include:

- mapping transport-addresses onto network-addresses;
- multiplexing (end-to-end) transport-connections onto network-connections;
- establishment and termination of transport-connections;
- end-to-end sequence control on individual connections;
- end-to-end error detection and any necessary monitoring of the quality of service;
- end-to-end error recovery;
- end-to-end segmenting and blocking;
- end-to-end flow control on individual connections;
- supervisory functions.

Detailed discussions of each of these functions follow.

2.4.1.1.1 Addressing

The Transport Layer defines and maintains transport addresses which higher layer entities use to communicate with other entities. Transport entities are identified by network addresses provided by the network layer. (Refer to Figure 2.4.1.1.1-1). Transport entities may support more than one session entity. Therefore several transport addresses may be associated with a single network address within the same transport entity. The corresponding mapping or switching functions have not been defined by ISO but must reside within the transport entity.

2.4.1.1.2 Connection Multiplexing

Transport connections are mapped onto network connections and the correspondence need not be one-to-one. Two types of multiplexing are possible: upward multiplexing (where one network connection supports more than one transport connection) and downward multiplexing (where more than one network connection supports one transport connection).

The benefits of upward multiplexing are first, that the network service can be used more efficiently and economically (in that network charges will be reduced) and secondly, it is possible to supply multiple transport connections where only one network connection exists.

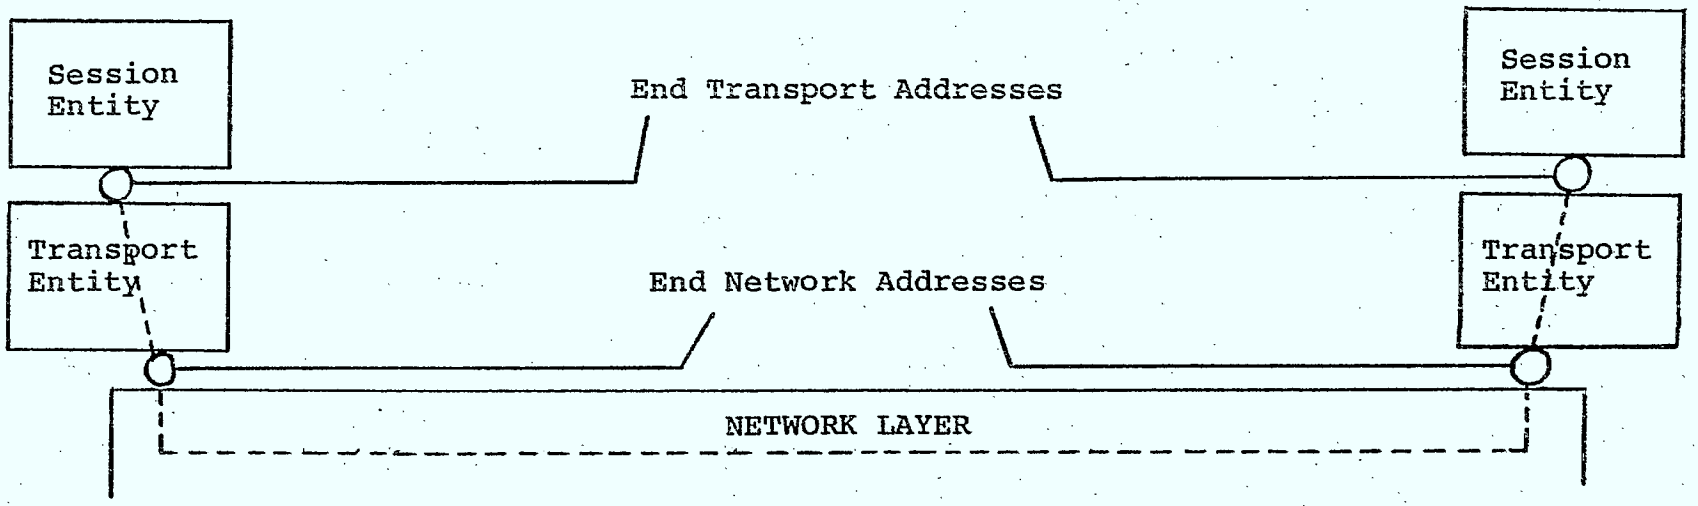

Figure 2.4.1.1.1-1 Association of transport-entities

Similarly, downward multiplexing has an associated set of benefits. First, reliability can be improved where multiple network connections exist. Secondly, the required grade of service can be made available by using multiple network connections. Thirdly, some cost benefits may be realized by the utilization of multiple low cost network connections with low grades of performance.

Multiplexing benefits do not come free however. Mapping functions involve a number of associated functions which need not be present when connection mapping is done on a one-to-one basis. The functions associated with upward multiplexing are:

(i) Single stream flow control

When the available bandwidth of a network connection is shared it is necessary to flow control each individual stream.

(ii) Individual flow identification

In order to ensure that data from the various multiplexed flows are not mixed, it is necessary that an identification of the individual flows be provided.

Similarly, downward multiplexing is associated with:

(i) Scheduling the utilizaton of multiple network connections.

(ii) Resequencing of data units associated with a transport connection that is multiplexed onto several network connections, even though each network connection guarantees sequence of delivery.

It is clear, because of the complexity of the above functions, that a cost effectiveness analysis be made in each particular implementation, to determine whether connection multiplexing is an important requirement.

2.4.1.1.2.1 Multiplexing - A Contentious Issue

The need for multiplexing in the lower three layers of the reference models is clear from the above discussion. Users of European packet switching networks have identified the need for multiplexing above virtual circuits. In general, packet switching tariffs are based on data volume transmitted and are largely independent of connection duration. However, several European networks have duration charges which dictate that, for low volume traffic applications, the most economical solution requires multiplexing above virtual circuits.

North American packet networks have tariffs based on volume only and there is a cost penalty for multiplexing above virtual circuits. For example, flow control is provided through the virtual circuit service but, when several independent streams are multiplexed above a single virtual circuit, independent flow control for each stream must be provided through the exchange of control information contained within data packets.

The CCITT model proposes multiplexing at the transport layer, and at the network layer; that is, multiple sessions can be multiplexed over a virtual circuit, and multiple virtual circuits over a data link. This implies a one-to-one mapping of sessions to transport connections (i.e. no multiplexing at the Session Layer).

The ISO model, on the other hand, proposes multiplexing at the Session Layer. The utility of this is unclear, and it is believed by some that the proposal is based on the false impression that in the ISO model X.25 may be a transport layer protocol.

2.4.1.1.3 Establishment and Termination

The phases of operation within the Transport Layer are the same as those in most other layers:

- establishment phase
- data phase
- termination phase

The transition from one phase to another has not yet been defined by ISO, but would be specified within the Transport Layer protocol.

2.4.1.1.4 Data Transfer

The following functions are associated with the Data Transfer phase of operation:

(i) End-to-end segmenting and blocking

The maximum size of transport service data units is not limited by the architecture of the Transport Layer but is related to the flow control functions and network service data unit size. Blocking of transport service data units during transfer may be desirable for cost optimization.

(ii) End-to-end sequence control

Sequence control may be necessary on individual connections if not supplied by the network connections or if downward multiplexing is used.

(iii) End-to-end error detection

Error detection in the Transport Layer uses error notificaton from the Network Layer. If additional error detection capabilities are necessary to provide the required quality of transport service, the Transport Layer may utilize peer-to-peer mechanisms such as end-to-end checksums, end-to-end acknowledgements, etc.

(iv) End-to-end error recovery

In order to achieve the quality of service guaranteed by the class of service selected, various error recovery functions may be performed within the Transport Layer. These functions include retransmission, use of alternative network-connections, etc.

(v) Flow Control

Flow control functions associated with each transport connection within the Transport Layer operate at the interface between the Transport Layer and the Session Layer, and between transport entities (as related to the peer-to-peer protocol).

(vi) Transport Layer Delivery Confirmation

This feature is now defined in X.25 (1980 version), and may be migrated to the Network Layer. ISO, however, specifies this as a Transport Layer function.

2.4.1.1.5 Layer Management

Management functions required by the Transport Layer have been defined by ISO, but must be handled within the layer and in conjunction with transport protocol.

2.4.1.2 ISO's Items for Further Study

Several issues related to the specification of the Transport Layer are currently unresolved and under further study. These issues are essential to the definition of even a basic Transport Service.

2.4.1.2.1 Scope of the Transport Service

The question here is whether the Transport Layer is expected to provide a universal and reliable process-to-process communication service or merely a system-to-system communication service. The scope of addressing provided by the Transport Service is also undefined.

2.4.1.2.2 Transport Protocols Assumptions Base

As mentioned in the introduction to this section, ISO's transport protocols, unlike CCITT's transport protocols, may not be based on the assumption that the network services which are seen at each end of a network-connection are necessarily provided by networks of the same nature. However this requires further study.

2.4.1.2.3 Expedited-Transport-Service-Data-Unit Transfer and Purge

It is felt that these two services may be redundant. An investigation has been undertaken to determine whether an expedited-transport-service-data-unit provides services other than purge.

2.4.1.2.4 Termination

The provision of an orderly termination service is under consideration. This function would guarantee data delivery when no unrecoverable error occurs during the termination phase. The termination phase already proposed by the ISO model does not guarantee data delivery.

2.4.1.2.5 Flow Control

It may be desirable that the transport service guarantees that the amount of data that it accepts from the sender can really be accepted by the receiving transport service user.

2.4.1.2.6 Transport Connection Identifiers

Is there a need for a transport connection identifier to distinguish between transport connections at the same transport address? Is there a need for this identifier to be known by the correspondent?

2.4.2 IBM-SNA Functions and Protocol

ISO's level 4, or the Transport Layer, provides control from user node to user node across a network. Some functions of ISO's Transport Layer are included in the Path control layer of SNA. The Path control layer, however, spans both the ISO Transport and Network layers, and is, in fact, more closely related to Network Layer functions (e.g. routing). This architectural organization seems to support CCITT'S view that the Transport Layer is in fact very sparse and serves only to enhance the quality of service [2]. In SNA, the composite of path, link, and physical control layers is termed the transport sub-system, analogous to ISO's transport service (layers 1 to 4 inclusive).

2.4.2.1 Transport Functions within Path Control

With the recent announcement of virtual routes in SNA, Path control now provides functions similar to those proposed in ISO's Transport Layer. The Path control functions include:

- address transformation
- multiplexing (upward multiplexing only)
- end-to-end sequencing on virtual routes
- end-to-end error detection and recovery
- end-to-end segmenting and blocking
- end-to-end flow control on virtual routes

Detailed discussions of each of these functions follow.

2.4.2.1.1 Addressing

SNA assigns network addresses to each SSCP, PU, and LU that communicates via the common transmission network. Network addresses have a sub-area and an element address component. Each PU type 4 and PU type 5 is assigned a specific sub-area address. All other PUs, LUs, and links in a domain are in the sub-areas of one or another of the Type 4 or 5 PU nodes, in which they reside or to which they are attached. The element address distinguishes a particular PU, LU, or link within a sub-area.

Each node in the network is characterized by the PU type it contains. PU4 and PU5 nodes - IBM 3204 and 3705 Communications Controllers, the System/370, the 303X, and 43X1 procesors - provide full network address routing, as well as both local and global flow control capabilities. PU1 and PU2 nodes - terminals and cluster controllers - are free of network address awareness and of the responsibility of network routing and global flow control. It is the responsibility of the boundary function within Path control to perform address translation for PU1 and PU2 nodes.

2.4.2.1.2 Session Multiplexing

Upward multiplexing, where one network connection (called explicit route in SNA) supports more than one transport connection (virtual route in SNA) is supported in SNA. Downward multiplexing does not exist at the Path Control Level.

Figure 2.4.2.1.2-1 shows the relationship of half sessions, virtual routes, and explicit routes. A transmission group is a group of one or more concurrently operating links between two adjacent sub-areas, and an explicit route (ER) is an ordered set of transmission groups. An ER is denoted by the sub-area addresses at the two ends and an explicit route number. A virtual route (VR) is a full duplex connection between two sub-area nodes. It is denoted by the sub-area addresses on the two ends of the VR, a virtual route number, and a transmission priority. Each VR must use an underlying ER, and there may be one or more VRs using the same ER (upward multiplexing). Several sessions may use the same VR simultaneously (upward multiplexing at the DFC/TC layer).

The VR and ER constructs provide flexibility in managing routes. VRs provide an end-to-end control, whereas ERs provide the physical representation of a route. Sessions using the same ER can be grouped into different VRs, thereby permitting different flow control and priority options to be exercised on each VR.

2.4.2.1.3 Segmenting and Blocking

Another function of Path Control which corresponds to the ISO Transport Layer is a repackaging job. This is called segmenting if the BIU is to be broken up into smaller packages (i.e. multiple PIU), or blocking, if a number of PIUs are to be combined in a single package. Segmenting and blocking are illustrated in figure 2.4.2.1.3-1.

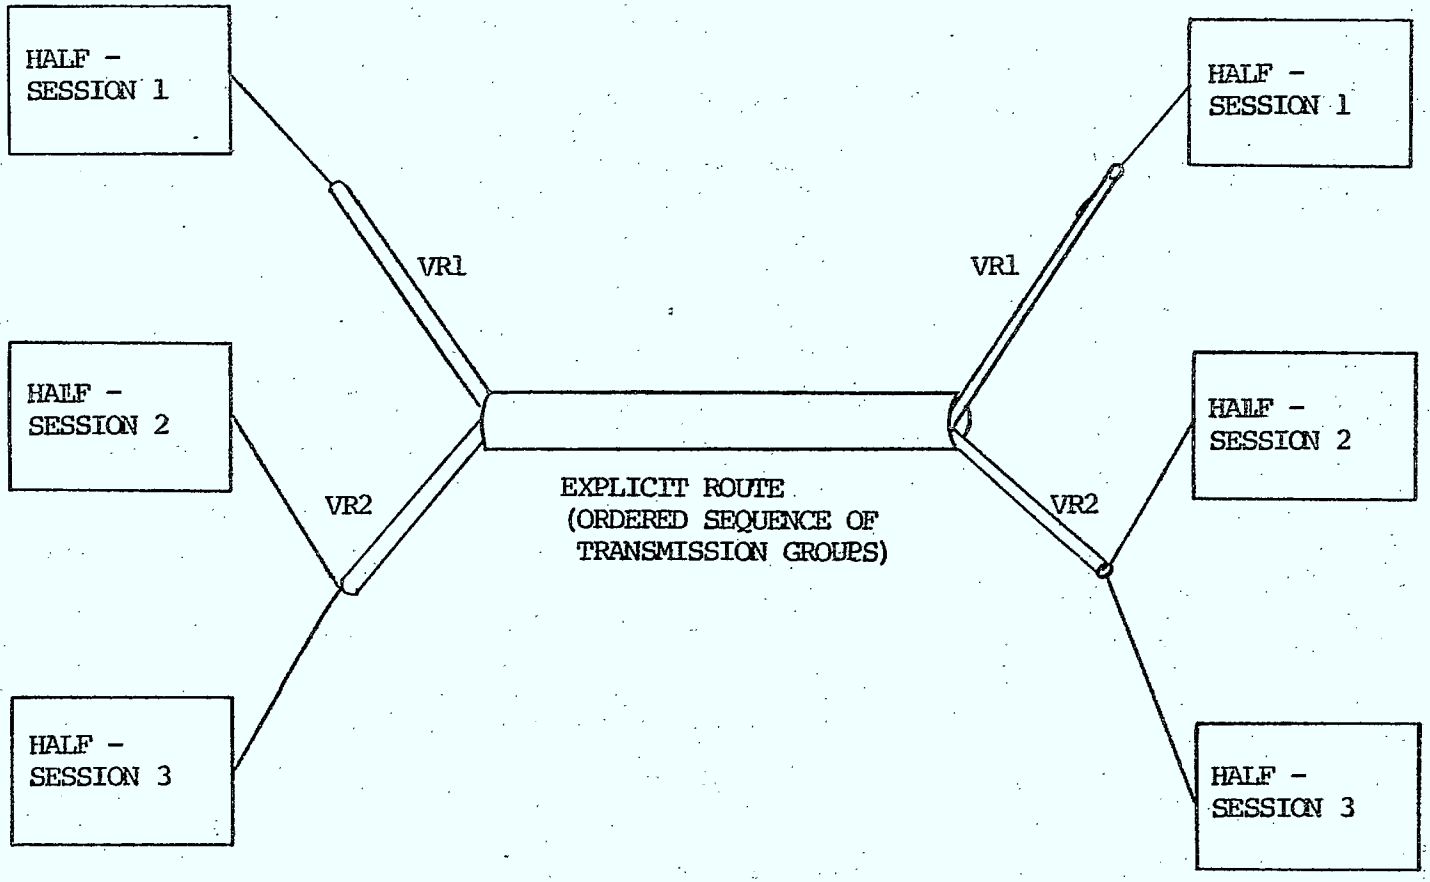

FIGURE 2.4.2.1.2-1 HALF SESSIONS, VIRTUAL ROUTES, AND AN EXPLICIT ROUTE

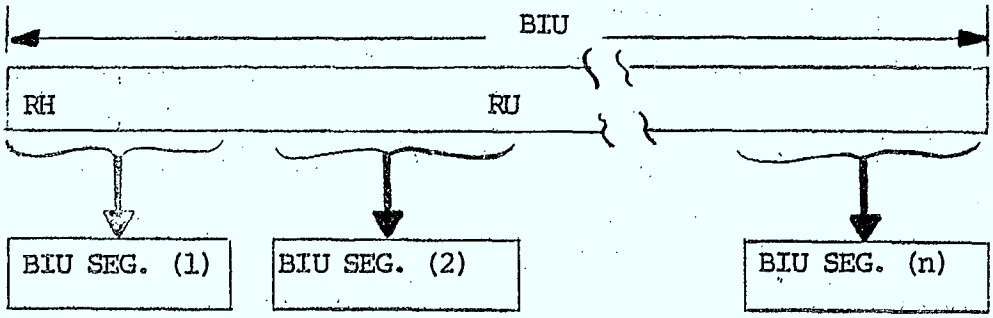

(a) SEGMENTING

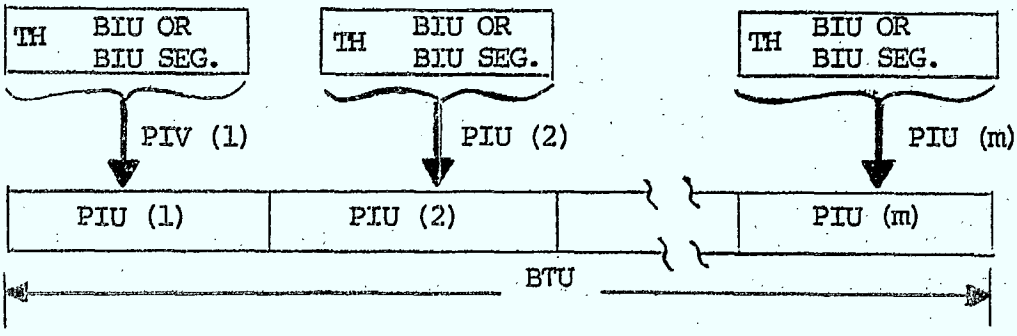

(b) BLOCKING

SEG. - SEGMENT

FIGURE 2.4.2.1.3-1 SEGMENTING AND BLOCKING

Response time is a major reason for segmenting. Figure 2.4.2.1.3-2 shows the potential for shortened transmission delay with segmenting. Also segmenting of large messages can aid in reducing the amount of retransmission on the occasion of a line failure.

The information in a block consists of a string of path information units (PIUs). This string is termed the basic transmission unit (BTU). It is the unit of information passed to the next lower layer. At each node, path control must examine the block, because it may be that some messages are to be forwarded to different destinations. When blocking, PIUs are grouped within a single set of data link control headers, thus saving some header overhead on the links. On the other hand, too much blocking forces longer queues to a specific data link, and it implies that a delay exists until a sufficient number of BIUs are assembled to fill a block. Although blocking is an available option, it must be used with discretion to avoid effects on response time, and is therefore not usually implemented as an SNA Path Control option.

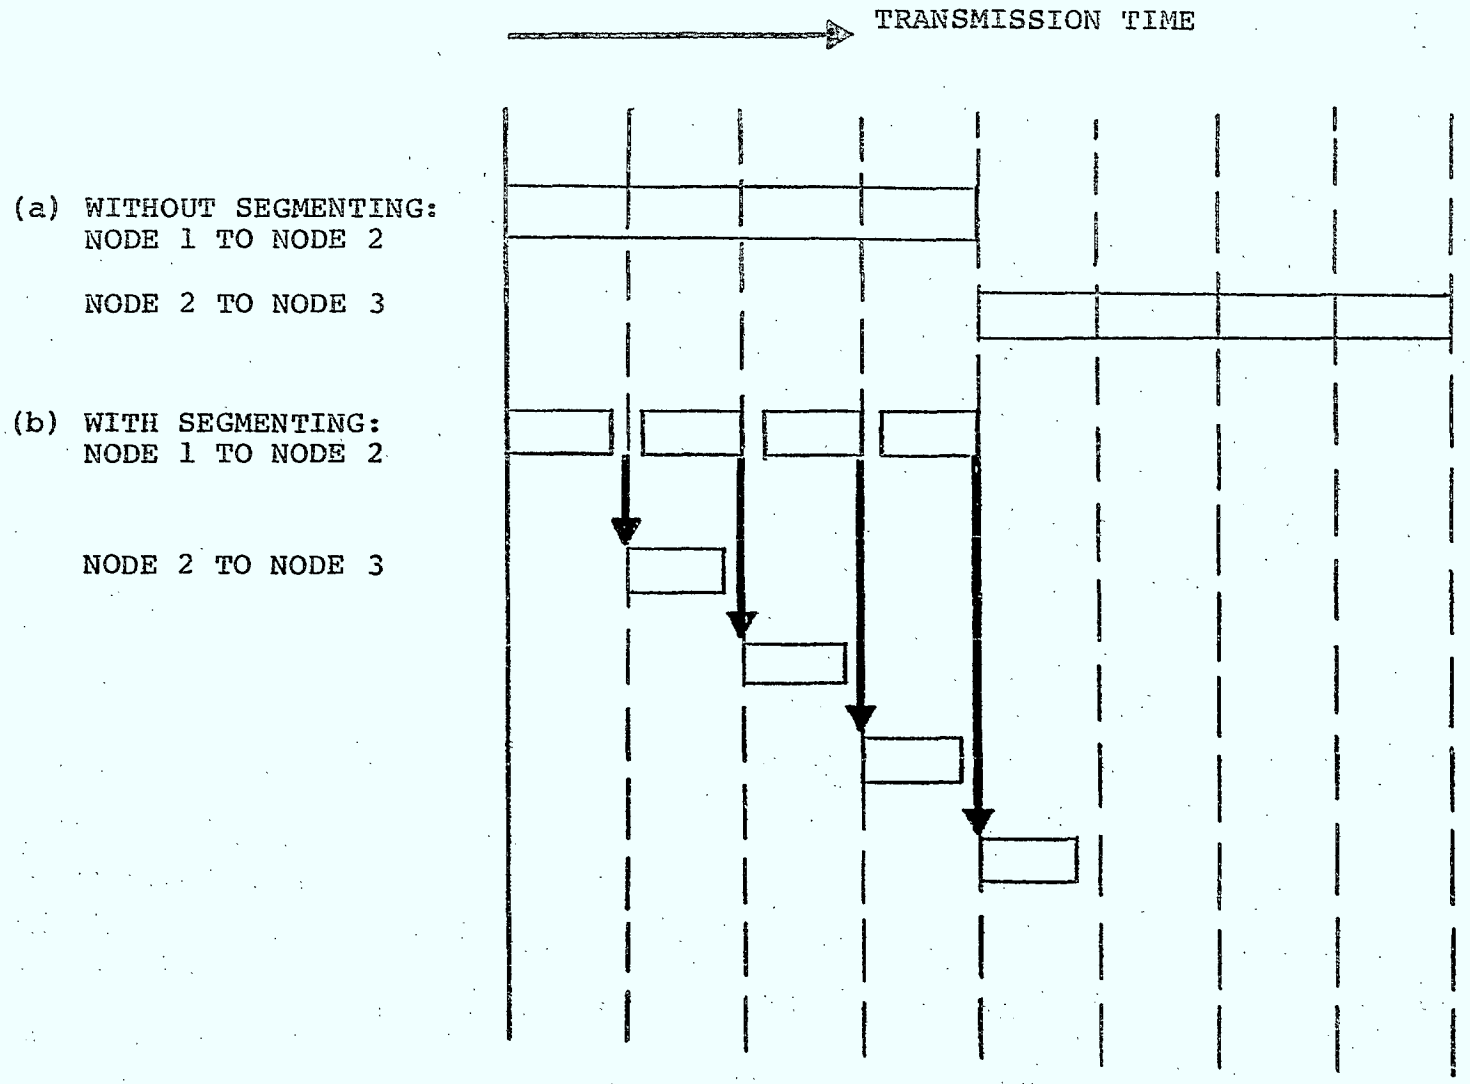

FIGURE 2.4.2.1.3-2 POTENTIAL FOR SHORTENED DELAY IN INTERMEDIATE NODE, WITH SEGMENTING

2.4.2.1.4 Flow Control

The VR end-to-end flow control is illustrated in figure 2.4.2.1.4-1. Nodes s and t are the two endpoints of a virtual route. The window size is the maximum number of PIUs that the sender can transmit in a group, or "window", following permission to send the window. Node s seeks permission to send the next window by setting a pacing request bit in the transmission header of the first PIU of a window destined for Node t. A pacing response from Node t permits Node s to send another window of PIUs after completing the current window. SNA offers an adaptive windowing scheme to adapt to network traffic and available buffer space.

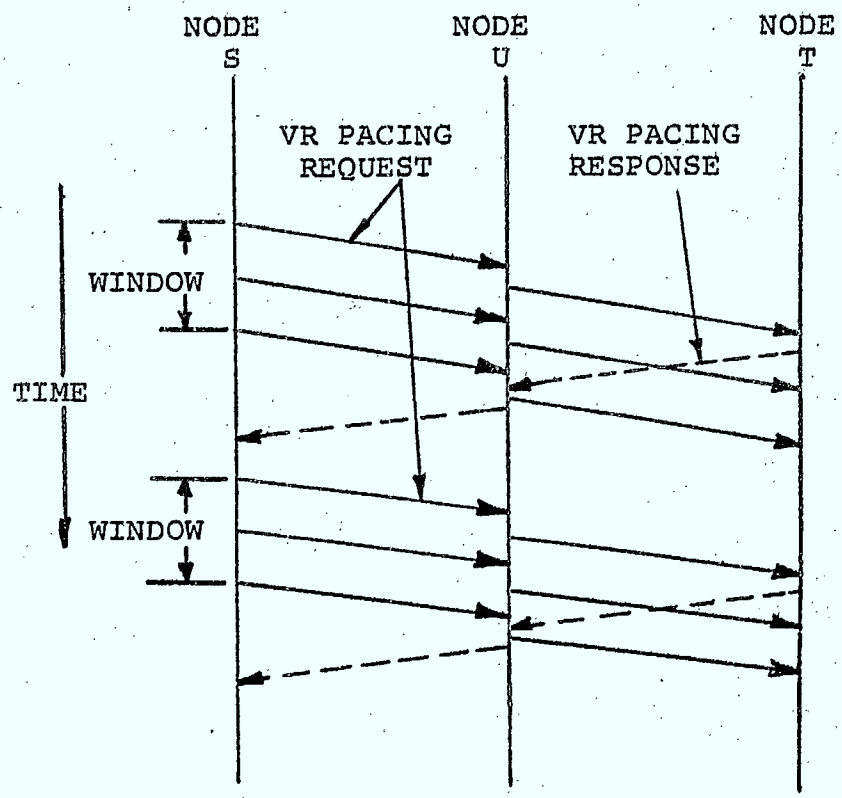

Figure 2.4.2.1.4-1 END-TO-END FLOW CONTROL

2.4.3 DEC-DNA Functions and Protocol

The logical link, as described in Section 2.3.3 is the functional equivalent of a Session connection in the OSI and SNA models. Since DECNET Phase II has no functional or logical equivalent to OSI's virtual circuits, or SNA's virtual routes, (i.e. no multiplexing below the logical link level), there exists a one-to-one mapping of session connections to network connections. Since the functionalities associated with the logical link were described in Section 2.3.3, they will not be reiterated here.

DECNET Phase III, however, proposes the introduction of adaptive path routing, which may imply significant redefinition of the Logical Link Layer services (a functional specification is not yet available). It has also been announced that DECNET Phase III will support an SNA interface product, an X.25 interface product, multipoint links, network command terminals (terminal communication with remote systems), and Network Management Facilities.

3.0 REQUIRED FUNCTIONALITIES AND ARCHITECTURAL SUPPORT

3.1 Introduction

This section examines the Canadian industry requirements identified in Section 1.4.1 and defines the functions that must be provided by a standard network architecture to meet these requirements. The location of the functions within the ISO model layering structure will be identified.

Those functions, needed to meet the Canadian industry requirements, that are identified as being within the ISO Model Presentation, Session or Transport Layers are discussed in Sections 3.3, 3.4 and 3.5 respectively. These sections discuss the extent to which the functional requirements are met by the ISO Model functions, the IBM-SNA functions and protocol, the DEC-DNA functions and protocol.

3.2 Functional Meaning of Requirements

In this section the Canadian industry requirements are defined in terms of the functions that must be provided by a standard network architecture to accommodate those requirements. The requirements are categorized as follows:

- extent of communications
- data traffic types
- end user services
- associated requirements

These are discussed in the following sections.

3.2.1 Extent of Communications

The requirements discussed in this section are:

- data communications with the United States
- data communications with Europe
- satellite data communications
- interfacing telex to data communications networks
- digitized voice transmission via data communications networks
- support for very large networks
- internetworking

The discussion of associated functions is presented in the following sub-sections.

3.2.1.1 Data Communications with the United States

The important networking function that must be provided to meet the requirement for trans-border communications is interfacing public networks in Canada with those in the United States. This will be met by the implementation of common standards at the bottom three layers of the ISO Model and provision of common end-to-end functions at the Transport Layer. The Transport Layer functions that must be provided are those identified in the following sections.

3.2.1.2 Data Communications with Europe

The functions that must be provided to meet the requirement for data communications with Europe are similar to those for data communications with the United States. The only differences being that the interfacing must be between public networks in Canada and those in Europe.

3.2.1.3 Satellite Data Communications

To meet the Canadian industry requirement to support satellite data communications, the end-to-end functions and protocol at each layer must be able to accommodate relatively long transmission times.

3.2.1.4 Interface Telex to Data Communications Network

To provide the required interface to Telex, the data communications network architecture must provide the following functions:

- support for slow transmission speeds, flow control, must be available at all layers of the architecture;

- presentation services must support all teletype terminals, a function of the Presentation Layer;

- asynchronous terminal operation must be supported by the Session Layer. This involves start-stop session functions.

3.2.1.5 Digitized Voice Transmission

The functions needed to support digitized voice transmission in data communications networks are:

- provision of large available bandwidth for each logical voice connection;

- provision of a logical connection that is free of error checking and retransmission routines;

- provision of high priority data routes.

These are visible to all the higher layers.

3.2.1.6 Support for Very Large Networks

The particular functions needed to meet the requirement to support very large networks include:

- addressing routines capable of uniquely defining every destination node/process in the network must be provided by some layer. The scope of addressing is currently being studied by ISO;

- the capability to support large volumes of traffic within the network itself (i.e. nodal buffering).

3.2.1.7 Internetworking

The particular functions needed to meet the requirement to support internetworking include:

- addressing standards that are either common between dissimilar networks, or that can be mapped at the network to network interface, must be provided (i.e. X.75, X.121);

- higher level functions to be provided across tandem networks must be common to the two communicating networks, i.e. common functions must be provided at each of the higher layers.

3.2.2 Data Traffic Types

The requirements discussed in this section are:

- communication between dissimilar processors;
- communication between terminals and processors;
- transaction transmission;
- bulk data transmission;
- communication between terminals.

The discussion of associated functions is contained in the following sections.

3.2.2.1 Communication Between Dissimilar Processors

The particular functions needed to meet the requirement to support data traffic between dissimilar processors include:

- provision of large available bandwidth between processors at all layers;

- support for multiple logical connections between processors, to allow multiple sessions, should be provided by the Transport Layer.

3.2.2.2 Communication Between Terminals and Processors

The particular functions needed to meet the requirement to support data traffic between terminals and processors include:

- virtual terminal service must be provided at the Presentation Layer;

- presentation options must be provided to support each terminal type to be used. This is a Presentation Layer function;

- transmission of small quantities of data to allow for fast response within conversations. This is a Session Layer function;

- flow control must be provided by the Session Layer if a one to one mapping does not exist between session connections and transport connections.

3.2.2.3 Transaction Transmission

The particular functions needed to meet the requirement to support transaction transmission data traffic include:

- the ability to start a session, send a small amount of data (the transaction), and stop the session in one physical transmission, should be provided by the higher layers for single transaction transmissions;

- delimiting of transactions should be provided by the Session Layer for multiple transaction transmissions.

3.2.2.4 Bulk Data Transmission

The particular functions needed to meet the requirement to support bulk data transmission data traffic include:

- the ability to transmit large volumes of data in one direction efficiently, within a session, should be provided by the Session Layer;

- the ability to use a low transmission priority should be provided by the Transport Layer.

- flow control must be provided by all layers.

3.2.2.5 Communication Between Terminals

The particular functions needed to meet the requirement to support data traffic between terminals include:

- virtual terminal service must be provided at the Presentation Layer.

3.2.3 End User Services

The requirements discussed in this section are:

- virtual terminal service;
- virtual file service;
- job transfer and manipulation service;
- other end user services.

The discussion of association functions is contained in the following sections.

3.2.3.1 Virtual Terminal Service

The particular functions needed to meet the requirement for a virtual terminal service include:

- the ability to identify and select a terminal profile to support each type of terminal to be connected;

- the ability to control presentation of data at each type of terminal;

- support for interactive processing, message switching and asynchronous terminals;

- support for word processing systems;

- support for specialized graphics systems.

These functions hide the detailed individual terminal characteristics from the communicating devices and are provided by the Presentation Layer.

3.2.3.2 Virtual File Service

The particular functions needed to meet the requirement for a virtual file service include:

- present internal file structures in standard formats to support file inquiry and update;

- present file commands in standard formats to support file inquiry/update and bulk data transmission;

- present the external file structures in standard formats to support bulk data transmission.

These functions are all provided at the Presentation Layer. Other required functions would be provided at the Application Layer.

The functions needed to meet the requirement for a data entry/collection service will be the same as for bulk data transmission if collected data is stored locally and transmitted to a remote site at intervals, and the same as for file inquiry/update if collected data is transmitted as soon as it is entered. The functions needed to meet the requirement for remote printing are the same as for bulk data shipping. The functions needed to meet the requirement for on-line programming will be the same as for bulk data transmission if the file is held locally and processed remotely, and the same as for file inquiry/update if the file is held and processed remotely.

3.2.3.3 Job Transfer and Manipulation Service

The particular functions needed to meet the requirements for a job transfer and manipulation service include:

- functions defined for bulk data transmission within the file transfer service to support job transfer;

- conversion of JCL to a standard format if required by the particular implementation;

- conversion of control commands to a standard format to support resource sharing and remote device control.

These functions are provided by the Presentation Layer. Other required functions would be provided at the Application Layer.

3.2.3.4 Other End User Services

The services discussed in this section are:

- electronic mail;

- provision of network services to outside users;

- network management functions.

The discussion of associated functions is contained in the following sections.

3.2.3.4.1 Electronic Mail

The electronic mail end user service is an application that uses the same virtual file service functions as those required to support inquiry/update file access. Other functions are required that would be provided by the Application Layer.

3.2.3.4.2 Provision of Network Services to Outside Users

This service is required by companies that sell their networking capabilities and resources. The particular functions needed to meet this requirement include:

- accountability of all network transmission and network resource utilization to the individual end users. This has implications at all layers.

Other functions are required that would be provided by the Application Layer.

3.2.3.4.3 Network Management Functions

This service includes such functions as network configuration management, network testing, provision of network operator services. These functions will be provided by the Application Layer.

3.2.4 Associated Requirements

The architectural requirements discussed in this section are:

- assured delivery;
- sequential delivery;
- availability and reliability;
- ability to interface with process control type computers;
- support for existing terminals and protocols;
- support for low speed and microprocessor controlled terminals;
- provision of device independence for applications;
- standardization of automatic teller terminals;
- provision of a low error rate service;
- accessibility to public databases;
- provision of an optional high security service;
- provision of an optional high priority service;
- ability to provide a fast response;
- accountability of network and processing resources.

The discussion of associated functions is contained in the following sections. Some non-architectural requirements specified by Canadian industry are:

- standards must be compatible with those of the United States;
- open data communications border with the United States.

3.2.4.1 Assured Delivery

The particular functions needed to meet the requirement for assured delivery include:

- the ability to acknowledge individual transmissions must be provided by the Transport Layer or Network Layer.

3.2.4.2 Sequential Delivery

The particular functions needed to meet the requirement for sequential delivery include:

- sequence numbering and sequence checking of all transmissions must be provided by the Transport Layer or Network Layer.

3.2.4.3 Availability and Reliability

The particular functions needed to meet the requirement to provide the identified level of availability and reliability include:

- error checking and error recovery procedures at all levels of the architecture.

3.2.4.4 Ability to Interface with Process Control Type Computers

The particular functions needed to meet the requirement to interface with process control type computers include:

- a defined set of presentation options suitable for process control type computers must be provided by the Presentation Layer.

3.2.4.5 Support for Existing Terminals and Protocols

The particular functions needed to meet the requirement to support existing terminals and protocols include:

- a defined set of presentation options suitable for each terminal type must be provided by the Presentation Layer;

- a particular node in the network will act as the interface to a particular terminal, this node will provide the presentation transformation required and will support the terminal protocol outside the defined network.

3.2.4.6 Support for Low Speed and Microprocessor Controlled Terminals

The particular functions needed to meet the requirement to support low speed and microprocessor controlled terminal include:

- a defined set of presentation options suitable for each type of terminal must be provided by the Presentation Layer.

3.2.4.7 Provision of Device Independence for Applications

The particular functions needed to meet the requirement to provide device independence for applications include:

- screening of device characteristics from the applications must be provided by the Presentation Layer.

3.2.4.8 Standardization of Automatic Teller Terminals

The particular functions needed to meet the requirement to standardize automatic teller terminals include:

- a defined set of presentation options suitable for automatic teller terminals must be provided by the Presentation Layer.

3.2.4.9 Provision of a Low Error Rate Service

The particular functions needed to meet the requirement to provide a low error rate service for applications include:

- error checking and error recovery procedures must be provided at all levels of the architecture.

3.2.4.10 Allow Access to Public Databases

The particular functions needed to meet the the requirement to access public data bases include:

- the provision of a common set of end-to-end functions in the public data base networks and private networks;

- the provision of common standards at the higher layers in both the public database networks and private networks, or the provision of a defined set of common functions within those layers.

3.2.4.11 Provision of an Optional High Security Service

The particular functions needed to meet the requirements to provide an optional high security service include:

- a class of service facility must be provided by all layers, which will include a high security option.

3.2.4.12 Provision of an Optional High Priority Service

The particular functions needed to meet the requirement to provide an optional high priority service include:

- a class of service facility must be provided by all layers, which will include a high priority option.

3.2.4.13 Ability to Provide a Fast Response

The particular functions needed to meet the requirement to be able to provide fast response include:

- a class of service facility must be provided by all layers, which will include a high priority option;

- the ability to indicate to the receiving application that a fast response is required must be provided by the Application Layer;

- the provision of an expedited flow facility.

3.2.4.14 Accountability of Network and Processing Resources

The particular functions needed to meet the requirement to provide accountability of network and processing resources must be provided by the Application Layer.

3.2.4.15 Standard Compatible with Those of the United States

This Canadian industry requirement shows the relative importance of North American standards, over international standards. The implication here is that a representative of Canadian Industry, for example the Department of Communications, be an active participant in the OSI standardization process.

3.2.4.16 Open Data Communications Border with the United States

This Canadian industry requirement implies that no restrictions be applied to data, interconnections, or network services.

3.3 Presentation Layer

This section compares the required Presentation Layer functions, identified by the survey, with the functions defined by the ISO Model and those provided by IBM SNA and DEC DNA. The functions are categorized as follows:

- functions within the layer
- end user services

They are discussed in the following sections.

3.3.1 Functions Within the Layer

The functions specifically identified as being needed to meet the Canadian industry requirements include:

- selection of presentation options, must take place to identify different terminal types, processor types, etc. This is defined by the ISO Model as taking place at the beginning of a session and also during a session if required (see sections 2.2.1.1.2,3). Similar functions are defined by SNA (see sections 2.2.2.1.1 (b) and 2.2.2.1.2 (a), (b)), and DNA (see sections 2.2.3.1.2,3);

- transformation of data, must take place to screen device characteristics from the applications. This is defined by the ISO Model (section 2.2.1.1.4) and by SNA (section 2.2.2.1.2 (c)) as being performed by either or both communicating partners. A more restricted version is defined by DNA (section 2.2.3.1.4);

- support for required class of service, must be provided to meet individual application requirements. The Presentation Layer functions involved with class of service include:

  i) security, by the provision of special function options such as encryption and decryption. These are defined within the ISO Model (section 2.2.1.1.6). SNA defines compression and compaction functions (section 2.2.2.1.2 (e)) at the FI.FMD layer, but encryption and decryption are performed at the DFC/TC layer (section 2.3.2.1.2 (b)(iv)). Special function options are not provided by DNA;

  ii) optional error checking, must be provided to support both applications requiring no error checking and retransmission, such as packetized voice transmission, and applications requiring error checking and retransmission at the Presentation Layer. Error checking and retransmission at the Presentation Layer is not defined by the ISO Model, but is defined by SNA (section 2.2.2.1.2 introduction) and DNA.

- accountability, of all network and processing resources must be provided for some applications. It is not defined by the ISO Model but is provided for by network services (section 2.2.2.1.1) in SNA. The function is not defined for DNA;

- protocol time independance, must be provided to support long transmission delays, slow transmission speeds etc. The ISO Model does not define the protocol. SNA FI.FMD protocol is not time dependant, DNA data access protocol may be made time dependant if required.

Other Presentation Layer functions implicitly needed to meet the Canadian industry requirements include:

- session establishment, termination and recovery, are required because of the implied requirements to establish sessions. These functions are defined by the ISO Model (sections 2.2.1.1.1,8,10), by SNA (sections 2.2.2.1.1 (a), (d), (f)), and by DNA (sections 2.2.3.1.1,6,8);

- control access to data formats, implied by the requirement to negotiate presentation options. This function is defined by the ISO Model (section 2.2.1.1.5) and by SNA (section 2.2.2.1.2 (d)), but is not defined by DNA;

- expedited flow, implied by the requirement to use control mechanisms outside of normal data transmission. This is defined by the ISO Model (section 2.2.1.1.9), by SNA (section 2.2.2.1.1 (e)) and by DNA (section 2.2.3.1.7).

3.3.2 End User Services

The end user services provided by the Presentation Layer are categorized as follows:

- virtual terminal service
- virtual file service
- job transfer and manipulation service

They are discussed in the following sections.

3.3.2.1 Virtual Terminal Service

The functions specifically identified as being needed to meet the Canadian industry virtual terminal service requirements include:

- selection of terminal class, must take place to support connection of all the terminal types required including existing terminals, data entry terminals, word processing terminals, graphics terminals, microprocessor controlled terminals etc. This is defined by the ISO Model (section 2.2.1.2.1 (a)) and by SNA (section 2.2.2.2.1 (a));

- negotiation of profile, must take place to identify mode of terminal operation to support interactive, asynchronous, low speed etc. This is defined by the ISO Model (section 2.2.1.2.1 (b)) and by SNA (section 2.2.2.2.1 (b)).

Other Presentation Layer functions implicitly needed to meet the Canadian industry virtual terminal service requirements include:

- data and command transfer, implied by the negotiation of a presentation profile. This is defined by the ISO Model (section 2.2.1.2.1 (c)) and by SNA (section 2.2.2.2.1 (c));

- forms management and control of operation, implied by the need to control printers, word processors, graphics terminals etc. This is defined by the ISO Model (section 2.2.1.2.1 (d), (e)) and by SNA (2.2.2.2.1 (d), (e)).

DNA does not define a virtual terminal service.

3.3.2.2 Virtual File Service

The functions specifically identified as being needed to meet the Canadian industry virtual file service requirements include:

- communication of file structure attributes, must take place to identify different file types and record structures. This is defined by the ISO Model (section 2.2.1.2.2 (a), (b)), by SNA (section 2.2.2.2.2 (a), (b)), and by DNA (section 2.2.3.2.1 (a), (b));

- communication of file commands and data, in standard format must take place to perform file access and manipulation operations. This is defined by the ISO Model (section 2.2.1.2.2 (c), (d)), by SNA (section 2.2.2.2.2 (c), (d)), and by DNA (section 2.2.3.2.1 (c), (d)).

3.3.2.3 Job Transfer and Manipulation Service

The functions specifically identified as being needed to meet the Canadian industry job transfer and manipulation service requirements include:

- data formatting and transfer, as for the virtual file service, are required to transmit the job stream in a standard format. This is defined by the ISO Model (section 2.2.1.2.3 (c)), by SNA (section 2.2.2.2.3 (c)), and by DNA (section 2.2.3.2.2).

- control of record structures and devices, must take place to allow resource sharing and remote device control. This is defined by the ISO Model (section 2.2.1.2.3 (a)) and by SNA (section 2.2.2.2.3 (a)), but not by DNA.

Other Presentation Layer functions implicitly needed to meet the Canadian industry job transfer and manipulation service requirements include:

- command formatting, is implied if the service is to allow communication between devices using different command formats. This is defined by the ISO Model (section 2.2.1.2.3 (b)). SNA allows certain job management commands to be formatted in a standard way (section 2.2.2.2.3 (b)). DNA does not define the function.

3.4 Session Layer

This section compares the required Session Layer functions, identified by the survey, with the functions defined by the ISO Model and those provided by IBM SNA and DEC DNA. The functions specifically identified as being needed to meet the Canadian industry requirements include:

- session recovery, must be available to increase network reliability. This is defined by the ISO Model (section 2.3.1.1.4), by SNA (section 2.3.2.1.2 (a)(iii)), and by DNA (section 2.3.3.1.4);

- flow control must be provided to support varying transmission speeds, delays, multiplexing etc. This is mentioned and is under further study by ISO. It is defined by SNA (section 2.3.2.1.1 introduction and section 2.3.2.1.2 (b)(ii)), and by DNA (section 2.3.3.1.6);

- support for required class of service, must be provided to meet individual application requirements. The Session Layer functions involved with class of service include:

i) session mapping, must be provided to support multiplexing. This is defined by the ISO Model (section 2.3.1.1.6), but is not fully supported in SNA and DNA;

ii) optional error checking and retransmission, must be provided to support applications not requiring error checking and retransmission, such as a packetized voice as well as those that do require error checking and retransmission. Error checking and retransmission is not defined by the ISO Model, but is defined by SNA (section 2.3.2.1.1 introduction, section 2.3.2.1.2 (a)(iii)), and by DNA (section 2.3.3.1.4).

- data delimiting, must be provided to improve transmission efficiency. Identified requirements are for the session service data unit, the quarantine unit and quarantine unit cancellation, and the session interaction unit. These are defined by the ISO Model (section 2.3.1.1.7) and by SNA (section 2.3.2.1.1 (b)), but DNA defines only session service data unit and the quarantine unit equivalents (section 2.3.3.1.7);

- accountability, of all network and processing resources must be provided for some applications. It is not defined by the ISO Model, but is provided for by network services (section 2.2.2.1.1) in SNA. The function is not defined for DNA;

- protocol time independance, must be provided to support long transmission delays, slow transmission speeds, etc. The ISO Model does not define the protocol. SNA DFC/TC protocol is not time dependant. DNA network services protocol has an implementation dependant time dependance.

Other Session Layer functions implicitly needed to meet the Canadian industry requirements include:

- session establishment, termination and identification, are required because of the implied requirement to establish sessions. These functions are defined by the ISO Model (sections 2.3.1.1.1,3,5), by SNA (sections 2.3.1.2.1.2 (a) (i),(iv);(c)), and by DNA (sections 2.3.3.1.1,3,5);

- context management, is required because of the Presentation Layer requirement to select presentation options. This is defined by the ISO Model (section 2.3.1.1.2), by SNA (section 2.3.2.1.2 (a)(ii)), and by DNA (section 2.3.3.1.2);

- dialogue management, implied by the requirement for different types of data traffic and dialogues. This is defined by the ISO Model (section 2.3.1.1.8) and by SNA (section 2.3.2.1.1 (a)). DNA does not define two-way dialogue controls.

3.5 Transport Layer

This section compares the required Transport Layer functions, identified by the survey, with the functions defined by the ISO Model and those provided by IBM SNA. As the DNA Transport Layer is a null set because of the current one to one mapping of logical links onto network connections, it cannot be compared to the Transport Layer requirements. The functions specifically identified as being needed to meet the Canadian industry requirements include:

- addressing, must support large networks and multiple networks. This is defined by the ISO Model (section 2.4.1.1.1), and by SNA (section 2.4.2.1.1);

- connection multiplexing, must be provided. This is defined by the ISO Model (section 2.4.1.1.2). Only upward multiplexing is currently defined by SNA (section 2.4.2.1.2);

- sequencing, must be provided to satisfy particular applications. This is defined by the ISO Model (section 2.4.1.1.4 (ii)). It is also provided by SNA, but not currently documented;

- support for required class of service, must be provided to meet individual applications requirements. The Transport Layer functions involved with class of service include:

i) optional error detection and recovery, must be provided to support applications not requiring error detection and recovery, such as packetized voice transmission, as well as those that do require error detection and recovery. This is defined by the ISO Model (section 2.4.1.1.4 (iii),(iv)). It is also provided by SNA but not currently documented;

ii) selectable transmission priority, must be provided to support both high and low priority transmission requirements identified by the survey. This is currently under study by ISO. SNA defines a priority class, but it does not have the same meaning as it is used for route allocation.

- segmenting and blocking, must be provided to support asynchronous transmissions and bulk transmissions. This is defined by the ISO Model (section 2.4.1.1.4(i)), and by SNA (section 2.4.2.1.3);

- flow control must be provided to accommodate multiplexing. This is defined by the ISO Model (section 2.4.1.1.4(v)), and by SNA (section 2.4.2.1.4);

- delivery confirmation, must be provided for some particular applications. This is defined by the ISO Model (section 2.4.1.1.4(vi)), but not within SNA;

- accountability, of all network and processing resources must be provided for some applications. It is not defined by the ISO Model, but is provided for by network services (section 2.2.2.1.1) in SNA.

REFERENCES

1. Reference Model of Open Systems Interconnection (Version 4, June 1979), ISO/TC97/SC16/H227.

2. International Telegraph and Telephone Consultative Committee (CCITT) Study Group VII. Structure for and use of a reference model for Public Data Network applications, September 1979.

3. R.J. Cypser, Communications Architectures for Distributed Systems, 1978.

4. F.P. Corr and D.H. Neal, SNA and emerging international standards, IBM Systems Journal, Volume 18, No. 2, 1979.

5. J.P. Gray and T.B. McNeill, SNA multiple-system networking, IBM Systems Journal, Volume 18, No. 2, 1979.

6. V. Ahuja, Routing and Flow control in Systems Network Architecture, IBM Systems Journal, Volume 18, No. 2, 1979.

7. J.P. Gray, Services provided to users of SNA Networks, IEEE 6th Data Communications Symposium, Nov. 1979.

8. SNA - Logical Unit Types, IBM publication GC20-1868.

9. SNA - Introduction to Sessions between logical units, IBM publication, GC20-1869.

10. SNA - Format and Protocol Reference Manual - Architectural Logic, IBM publication SC30-3112.

11. SNA - Information paper for SC16, June 1978.

12. J.P. Gray, T.B. McNeill, SNA-4 features, September 1978.

13. G. Schultz, B. Sundstrom, A Brief look at SNA, Sept. 1979.

14. R. Bird, D. Rose, G. Schultz, SNA Specification, Oct. 1979.

15. C. West, J. Rubin, SNA Validation, Oct. 1979.

16. SNA Information Paper for SC16, June 1978.

17. Interface - SNA protocols, Nov. 1979.

18. Network Interface Adapter, IBM Publication GA11-8632-0.

19. Datapac/Transpac Program Description and Operation Manual, IBM Publication, SH19-6052.

20. Introduction to Advanced Communications Function, IBM Publication GC30-3033-1.

21. V. Ahuja, J.E. Merkel, Recent Developments in Systems Network Architecture and Reference Model for Open Systems Interconnection, to be delivered at Computer Networks Conference, June 1980.

22. Z. Cvijan, SNA/X.25 Interface Possibilities, SHARE 54, March 1980.

23. Z. Cvijan, SHARE 53.5 Trip Report, Nov. 1979.

24. Z. Cvijan, Packet Data Networks and SNA, SHARE 53.5, Nov. 1979.

25. A.M. Rybezynski, J.D. Palframan, A Common X.25 Interface to Public Data Networks, SHARE 54, March 1980.

26. I.M. Cunningham, Transport Service Standardization Issues In Open Systems Interconnection, Pacific Telecommunication Conference 80, Jan. 1980.

27. Data Access Protocol, DEC Publication AA-D601A-TC.

28. Network Services Protocol, DEC Publication, AA-D600A-TC.

APPENDIX A

COVERING LETTER

Systemhouse Ltd.

February 1, 1980

Dear Sir/Ms:

Systemhouse Ltd. has recently undertaken a study of higher level protocols for open system interconnection for the Department of Communications. Open Systems Interconnection (OSI) refers to standardized procedures for the exchange of information among terminal devices, computers, people, networks, processes, etc., that are "open" to one another for this purpose by virtue of their mutual use of these procedures.

"Openness" does not imply any particular systems implementation, technology or interconnection means, but rather refers to the mutual recognition and support of the standardized information exchange procedures.

There are several manufacturers currently implementing OSI architectures to meet the ever-increasing requirements of distributed processing. Without external impetus, the computer/communications architectures being developed and made available by hardware manufacturers will be, by and large, unique to that manufacturer's equipment. This is due not only to the desirable result of competition fostering original development, but to a conscious desire to create a unique product as a marketing consideration.

The user and developer of information systems, on the other hand, has a real and identifiable requirement to build applications that span interconnected systems. The existing and announced manufacturer's architectures are not generally applicable to dissimilar system interconnection, and custom development of the systems software is prohibitive, particularly where there is no standard to which to adhere.

The overall result is no doubt a dampening effect on the development of state-of-the-art and cost effective software, in the areas of both systems and application software. In a nation like Canada, so dependent on communications as a means of business, and so intent on information processing as an industry sector, such an impediment is of great concern.

.../2

99 Bank Street, 3rd Floor, Ottawa, Canada K1P 6B9 (613) 238-9734 Telex 053-4309

In order for the Department of Communications to accurately represent Canadian communication needs in the formulation of international standards we intend, as part of this study, to survey the Canadian marketplace in order to ascertain the current status quo of distributed systems. Distributed systems range anywhere from the connection of intelligent peripherals to a central processing unit, to a network of mainframes which perform something in excess of remote job entry functions; that is to say, some sort of task to task communication.

The attached survey is intended to accomplish a threefold objective:

(i) to ascertain the commitment to manufacturer provided architectures and the effect of non-standardization on forward systems planning.

(ii) to ascertain the projected development of custom architectures/protocols to accomplish dissimilar system interconnection.

(iii) to ascertain the Canadian requirements for distributed systems.

Your submission is strictly confidential, and there is no requirement to identify your company with your return. Our intention is to formulate a conceptual overview of national distributed processing, not to deal in the specifics of your particular applications. Your confidentiality will be respected by means of a non-disclosure agreement, if so desired. Your reply would be appreciated by the end of February, 1980.

Sincerely,

Patrick Power
Project Manager

/lm

QUESTIONNAIRE

1. Does your organization have the requirement to build applications that span interconnected systems? If so in diagrammatic form (conceptual diagram) please show:

   (i) planned or aspired processing networks.

   (ii) type of communication between the components of the above network(s); i.e. batch, inquiry, data entry, output printing, etc.

   (iii) constraints or incumbents imposed by manufacturer hardware or software packages. Briefly explain.

2. Are your requirements particularly stringent or otherwise noteworthy? Please elaborate.

3. To satisfy 1, have you, or are you planning the implementation of an available manufacturer network architecture (e.g. IBM's SNA, DIGITAL's DECNET, etc.)? If so, what network architecture, and briefly, why?

4. To satisfy 1, have you, or are you planning the interconnection of multiple manufacturer networks? If so, did you, or do you plan to:

   (i) Make all networks emulate some master architecture (e.g. SNA)?

   (ii) Not use higher level protocols but provide compatibilities at the link level.

   (iii) Have parallel (or independent) systems that exchange bulk data on a periodic basis.

5. To satisfy 1, do you interface directly to some other value added network, or provide access to your own network? What network services do you use or provide?

6. Do your applications have international requirements?

7. Are you aware of the current non-standardization of higher level protocols for system interconnection (above level 3 of CCITT's X.25 standard)? Are you actively involved in the design of such higher level protocols?

8. Do you feel that the existing manufacturer architectures/protocols will meet your networking requirements? Do you find the incompatibility (non-standardization) a problem? How is this affecting your forward systems planning?

9. What type of organization do you represent?

--A study of higher level protocols for open system interconnection: final report.