INTERNATIONAL TELECOMMUNICATION UNION

CCITT

X.213

(09/92)

THE INTERNATIONAL
TELEGRAPH AND TELEPHONE
CONSULTATIVE COMMITTEE

DATA COMMUNICATION NETWORKS

INFORMATION TECHNOLOGY – NETWORK SERVICE DEFINITION FOR OPEN SYSTEMS INTERCONNECTION

Recommendation X.213

Foreword

ITU (International Telecommunication Union) is the United Nations Specialized Agency in the field of telecommunications. The CCITT (the International Telegraph and Telephone Consultative Committee) is a permanent organ of the ITU. Some 166 member countries, 68 telecom operating entities, 163 scientific and industrial organizations and 39 international organizations participate in CCITT which is the body which sets world telecommunications standards (Recommendations).

The approval of Recommendations by the members of CCITT is covered by the procedure laid down in CCITT Resolution No. 2 (Melbourne, 1988). In addition, the Plenary Assembly of CCITT, which meets every four years, approves Recommendations submitted to it and establishes the study programme for the following period.

In some areas of information technology which fall within CCITT's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC. The text of CCITT Recommendation X.213 was approved on 10 September 1992. The identical text is also published as ISO/IEC International Standard 8348.

CCITT NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized private operating agency.

© ITU 1993

All rights reserved. No part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the ITU.

Contents

Page

SECTION 1 – GENERAL 1

1 Scope 1

2 Normative references 1

2.1 Paired Recommendations | International Standards equivalent in technical content 2

2.2 Additional references 2

3 Definitions 2

3.1 Basic reference model definitions 2

3.2 Service conventions definitions 3

3.3 Network Service definitions 3

3.4 Network addressing definitions 3

3.5 Network layer architecture definitions 4

4 Abbreviations 4

5 Conventions 5

5.1 General conventions 5

5.2 Parameters 5

5.3 NC endpoint identification convention 6

6 Overview and general characteristics 6

7 Types and classes of Network Service 6

SECTION 2 – DEFINITION OF THE CONNECTION-MODE SERVICE 7

8 Features of the connection-mode Network Service 7

9 Model of the connection-mode Network Service 7

9.1 Model of the connection-mode Network Layer Service 7

9.2 Model of a Network Connection 7

10 Quality of the connection-mode Network Service 11

10.1 Determination of QOS 11

10.2 Definition of QOS-parameters 12

11 Sequence of primitives 16

11.1 Relation of primitives at the two NC end points 16

11.2 Sequence of primitives at one NC endpoint 16

12 Connection establishment phase 19

12.1 Function 19

12.2 Types of primitives and parameters 19

12.3 Sequence of primitives 27

13 Connection release phase 29

13.1 Function 29

13.2 Types of primitive and parameters 30

13.3 Sequence of primitives when releasing an established NC 31

13.4 Sequence of primitives in an NS user rejection of an NC establishment attempt 32

13.5 Sequence of primitives in an NS provider rejection of an NC establishment attempt 33

14 Data transfer phase 33
14.1 Data transfer 33
14.2 Receipt confirmation service 35
14.3 Expedited data transfer service 36
14.4 Reset service 37

SECTION 3 – DEFINITION OF THE CONNECTIONLESS-MODE SERVICE 41

15 Features of the connectionless-mode Network Service 41

16 Model of the connectionless-mode Network Service 41
16.1 Model of the connectionless-mode Network Layer Service 41
16.2 Model of a network connectionless-mode transmission 41

17 Quality of the connectionless-mode Network Service 42
17.1 Determination of QOS 42
17.2 Definition of network connectionless-mode QOS-parameters 43
17.3 Route selection considerations 44

18 Sequence of primitives 44

19 Data transfer 44
19.1 Function 44
19.2 Types of primitives and parameters 45
19.3 Sequence of primitives 46

Annex A Network Layer Addressing 47
A.1 General 47
A.2 Scope 47
A.3 Concepts and terminology 47
A.4 Principles for creating the OSI Network addressing scheme 50
A.5 Network address definition 51
A.6 Character based DSP allocation 57
A.7 Reference publication formats 58
A.8 Network entity titles 58

Annex B Rationales for the material in Annex A 59
B.1 IDI formats (see A.5.2.1.2) 59
B.2 Reservation of AFI values 00-09 (see Table A.2) 59
B.3 Derivation of the preferred encodings (see A.5.3) 60

Annex C Facilities for conveying service characteristics in the connectionless-mode Network Service 61
C.1 Introduction 61
C.2 Function 61
C.3 Types of primitives and parameters 62
C.4 Service characteristics 62
C.5 Types of primitives and parameters 63

Introduction

This Recommendation | International Standard is one of a set of Recommendations and International Standards produced to facilitate the interconnection of computer systems. It is related to other Recommendations and International Standards in the set as defined by the Reference Model of Open Systems Interconnection (OSI). The OSI Reference Model (CCITT Rec. X.200 | ISO 7498) subdivides the area of standardization for interconnection into a series of layers of specification, each of a manageable size.

This Recommendation | International Standard defines the Service provided by the Network Layer to the Transport Layer at the boundary between the Network and Transport Layers of the Reference Model. It provides for the designers of Transport protocols a definition of the Network Service existing to support the Transport protocol and for the designers of Network protocols a definition of the services to be made available through the action of the Network protocol over the underlying service. This relationship is illustrated in the figure below.

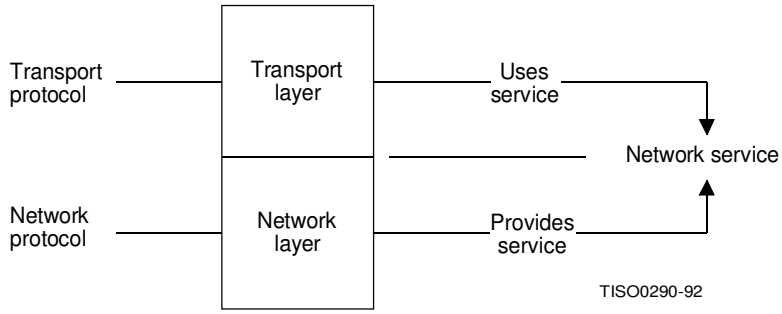

Relationship of the Network Service to OSI Network and Transport protocols

The use of the word "Network" to name the "Network" Layer of the OSI Reference Model should be distinguished from the use of the word "network" to denote a communications network as conventionally understood. To facilitate this distinction, the term "subnetwork" is used for a collection of physical equipment, commonly called a "network" (see CCITT Rec. X.200 | ISO 7498). Subnetworks may be either public networks or privately supplied networks. In the case of public networks, their properties may be determined by separate Recommendations such as Recommendation X.21 for a circuit-switched network or Recommendation X.25 for a packet-switched network.

Throughout the set of OSI Recommendations and International Standards the term "Service" refers to the abstract capability provided by one layer of the OSI Reference Model to the layer above it. Thus, the Network Service defined in this Recommendation | International Standard is a conceptual architectural Service, independent of administrative divisions.

NOTE – It is important to distinguish the specialized use of the term "Service" within the set of OSI Recommendations and International Standards from its use elsewhere to describe the provision of a service by an organization (such as the provision of a service, as defined in other Recommendations, by an Administration).

Any particular subnetwork may or may not support the OSI Network Service. The OSI Network Service may be provided by a combination of one or more subnetworks and optional additional functions between or outside these subnetworks.

INTERNATIONAL STANDARD

CCITT RECOMMENDATION

INFORMATION TECHNOLOGY – OPEN SYSTEMS INTERCONNECTION – NETWORK SERVICE DEFINITION

SECTION 1 – GENERAL

1 Scope

This Recommendation | International Standard defines the OSI Network Service in terms of

a) the primitive actions and events of the Service;

b) the parameters associated with each primitive action and event, and the form which they take;

c) the interrelationship between, and the valid sequences of, these actions and events.

The principal objectives of this Recommendation | International Standard are

1) to specify the characteristics of a conceptual Network Service and thus, supplement the Reference Model in guiding the development of Network Layer protocols;

2) to encourage convergence of the capabilities offered by providers of subnetworks;

3) to provide a basis for the individual enhancement of existing heterogeneous subnetworks to a common subnetwork-independent Network Service to enable them to be concatenated for the purpose of providing global communication. (Such concatenation may involve optional additional functions which are not defined in this Recommendation | International Standard.) A definition of the quality of service is an important element of this Recommendation | International Standard;

4) to provide a basis for the development and implementation of subnetwork-independent Transport Layer protocols decoupled from the variability of underlying public and private subnetworks and their specific interface requirements.

This Recommendation | International Standard does not specify individual implementations or products nor does it constrain the implementation of entities and interfaces within a system.

There is no conformance of equipment to this Recommendation | International Standard. Instead, conformance is achieved through implementation of conforming OSI Network protocols which fulfill the Network Service defined in this Recommendation | International Standard.

2 Normative references

The following CCITT Recommendations and International Standards contain provisions which, through reference in this text, constitute provisions of this Recommendation | International Standard. At the time of publication, the editions indicated were valid. All Recommendations and International Standards are subject to revision, and parties to agreements based on this Recommendation | International Standard are encouraged to investigate the possibility of applying the most recent edition of the CCITT Recommendations and International Standards listed below. Members of the IEC and ISO maintain registers of currently valid International Standards. The CCITT Secretariat maintains a list of the currently valid CCITT Recommendations.

2.1 Paired Recommendations | International Standards equivalent in technical content

– CCITT Recommendation X.200 (1988), *Reference model of Open Systems Interconnection for CCITT Applications*.

ISO 7498:1984, *Information processing systems – Open Systems Interconnection – Basic Reference Model*.

– CCITT Recommendation X.210 (1988), *Open Systems Interconnection layer service definition conventions*.

ISO/TR 8509:1987, *Information processing systems – Open Systems Interconnection – Service conventions*.

– CCITT Recommendation X.224 (1988), *Transport protocol specification for Open Systems Interconnection for CCITT applications*.

ISO/IEC 8073:1992, *Information technology – Telecommunications and information exchange between systems – Open Systems Interconnection – Protocol for providing the connection-mode transport services*.

2.2 Additional references

– CCITT Recommendation E.163 (1998), *Numbering plan for the international telephone service*.

– CCITT Recommendation E.164 (1991), *The numbering plan for the ISDN era.*

– CCITT Recommendation F.69 (1988), *Plan for telex destination codes*.

– CCITT Recommendation X.121 (1992), *International numbering plan for public data networks*.

– CCITT Recommendation X.300, (1988) *General principles for interworking between public networks and between public networks and other networks for the provision of data transmission services*.

– ISO/IEC 646: 1991, *Information technology – ISO 7-bit coded character set for information interchange*.

– ISO 2375: 1985, *Data processing – Procedure for registration of escape sequences*.

– ISO 3166: 1988, *Codes for the representation of names of countries*.

– ISO 6523: 1984, *Data interchange – Structures for the identification of organizations*.

– ISO 7498:1984/Add. 1:1987, *Information processing systems – Open Systems Interconnection – Basic Reference Model – Addendum 1*: *Connectionless-mode transmission*.

– ISO 8648:1988, *Information processing systems – Open Systems Interconnection – Internal organization of the Network Layer*.

3 Definitions

For the purposes of this Recommendation | International Standard, the following definitions apply.

3.1 Basic reference model definitions

This Recommendation | International Standard is based on the concepts developed in the Basic Reference Model for Open Systems Interconnection and makes use of the following terms defined in CCITT Rec. X.200 | ISO 7498:

a) expedited Network-Service-data-unit;

b) Network-address,

c) Network Connection;

d) Network-entity,

e) Network-protocol control information,

f) Network-protocol data unit,

g) Network Layer.

h) Network-relay.

i) Network-routing,

j) Network Service;

k) Network-Service-access-point;

l) Network-Service-access-point-address;

m) Network-Service-data-unit;

n) OSI environment,

o) subnetwork;

p) title.

3.2 Service conventions definitions

This Recommendation | International Standard also makes use of the following terms defined in CCITT Rec. X.210 | ISO/TR 8509, as they apply to the Network Layer:

a) confirm;

b) indication;

c) Network Service user;

d) Network Service provider;

e) primitive;

f) request;

g) response.

3.3 Network Service definitions

For the purpose of this Recommendation | International Standard, the following definitions also apply:

3.3.1 calling NS user: An NS user that initiates an NC establishment request.

3.3.2 called NS user: An NS user with whom a calling NS user wishes to establish an NC.

NOTE – Calling NS users and called NS users are defined with respect to a single NC. An NS user can be both a calling and a called NS user simultaneously.

3.3.3 generic address: An address which identifies a set of NSAPs rather than a single specific NSAP.

3.3.4 Network Connection: An association established by a Network Layer between two NS users for the transfer of data, which provides explicit identification of a set of Network data transmissions and agreement concerning the services to be provided by the set.

NOTE – This definition clarifies that given in CCITT Rec. X.200 | ISO 7498.

3.3.5 Network connection-mode Data Transmission: The transfer of an NSDU from a source NSAP to a destination NSAP within the context of an NC that has previously been established

3.3.6 Network connectionless-mode Data Transmission: The transmission of an NSDU from a source NSAP to a destination NSAP outside the context of an NC and without any requirement to maintain any logical relationship among multiple invocations.

3.4 Network addressing definitions

Annex A, describing network addressing makes use of the following terms as defined below:

3.4.1 DTE address: Information used to identify a point of attachment to a public data network.

3.4.2 subnetwork point of attachment: A point at which a real end system, interworking unit, or real subnetwork is attached to a real subnetwork, and a conceptual point at which a subnetwork service is offered within an end or intermediate system.

3.4.3 subnetwork point of attachment address: Information used in the context of a particular real subnetwork to identify a subnetwork point of attachment; or information used in the context of a particular subnetwork to identify the conceptual point within an end or intermediate system at which the subnetwork service is offered. This term is used interchangeably with the (equivalent) shortened form *subnetwork address*.

3.4.4 network protocol address information: Information encoded in a Network protocol data unit to carry the semantics of a Network service access point address. (This is known as an "address signal" or as the "coding of an address signal" in the public network environment.)

3.4.5 naming domain: A context within which a name allocated by a naming authority is unambiguous. Where the name is an address, the context within which the name is allocated is called an *addressing domain*.

3.4.6 global network addressing domain: An addressing domain consisting of all of the Network service access point addresses in the OSI environment.

3.4.7 network addressing domain: A subset of the global network addressing domain consisting of all of the Network service access point addresses allocated by one or more addressing authorities.

3.4.8 naming authority: That which allocated names from a specified naming domain, and which ensures that names so allocated are unambiguous. Where the naming authority allocates addresses, it is called an *addressing authority*.

3.4.9 network addressing authority: An addressing authority that assigns and administers Network service access point addresses within one or more network addressing domains.

3.4.10 abstract syntax: A notation which enables data types to be defined, and values of those types specified, without determining the way in which they will be represented (encoded) for transfer by protocols.

3.5 Network layer architecture definitions

This Recommendation | International Standard makes use of the following terms defined in CCITT Rec. X.300 and ISO 8648.

a) subnetwork;

b) real subnetwork;

c) subnetwork service;

d) real end system;

e) interworking unit;

f) intermediate system;

g) relay entity.

4 Abbreviations

AFI	Authority and format identifier
CC	Country code
COR	Confirmation of receipt
DCC	Data country code
DSP	Domain specific part
ENSDU	Expedited Network-Service-data-unit
ICD	International code designator
IDI	Initial domain identifier

IDP	Initial domain part
ISDN	Integrated services digital network
N	Network
NC	Network Connection
NL	Network Layer
NPAI	Network protocol addressing information
NPDU	Network protocol data unit
NS	Network Service
NSAP	Network-Service-access-point
NSDU	Network-Service-data-unit
OSI	Open Systems Interconnection
PSTN	Public switched telephone network
PTT	Postal, telephone and telegraph
QOS	Quality of Service
RPF	Reference publication format
SNPA	Subnetwork point of attachment

5 Conventions

5.1 General conventions

This Service Definition uses the descriptive conventions given by CCITT Rec. X.210 | ISO/TR 8509.

The layer service model, service primitives, and time-sequence diagrams taken from those conventions are entirely abstract descriptions; they do not represent a specification for implementation.

5.2 Parameters

Service primitives, used to represent service-user/service-provider interactions (see CCITT Rec. X.210 | ISO/TR 8509), convey parameters which indicate information available in the user/provider interaction.

The parameters which apply to each group of Network Service primitives are set out in tables in clauses 12 to 14 and 19. Each "X" in the tables indicates that the primitive labeling the column in which it falls may carry the parameter labeling the row in which it falls.

Some entries are further qualified by items in brackets. These may be

a) an indication that the parameter is conditional in some way:

   (C) indicates that the parameter is not present on the primitive for every NC; the parameter definition describes the conditions under which the parameter is present or absent;

b) a parameter specific constraint:

   (=) indicates that the value supplied in an indication or confirm primitive is always identical to that supplied in the corresponding request or response primitive occurring at the peer NSAP;

c) an indication that some note applies to the entry:

   (Note x) indicates that the referenced note contains additional information pertaining to the parameter and its use.

In any particular interface, not all parameters need be explicitly stated. Some may be implicitly associated with the NSAP at which the primitive is issued.

5.3 NC endpoint identification convention

If an NS user needs to distinguish among several NCs at the same NSAP, then a local NC endpoint identification mechanism must be provided. All primitives issued at such an NSAP would be required to use this mechanism to identify NCs. Such an implicit identification is not described as a parameter of the service primitives in this Service Definition.

NOTE – The implicit NC endpoint identification must not be confused with the address parameters of the N-CONNECT primitives (see 12.2).

6 Overview and general characteristics

The Network Service provides for the transparent transfer of data (i.e., NS-user-data) between NS users. It makes invisible to these NS users the way in which supporting communications resources are utilized to achieve this transfer.

In particular, the Network Service provides for the following:

a) *Independence of underlying transmission media* – The Network Service relieves NS users from all concerns regarding how various subnetworks are used to provide the Network Service. The Network Service hides from the NS user differences in the transfer of data over heterogeneous subnetworks, other than quality of service.

b) *End-to-end transfer* – The Network Service provides for transfer of NS-user-data between NS users in end systems. All routing and relaying functions are performed by the NS provider including the case where several similar or dissimilar transmission resources are used in tandem or in parallel.

c) *Transparency of transferred information* – The Network Service provides for the transparent transfer of octet-aligned NS-user-data and/or control information. It does not restrict the content, format or coding of the information, nor does it ever need to interpret its structure or meaning.

d) *Quality of service selection*–- The Network Service makes available to NS users a means to request and to agree to the quality of service for the transfer of NS-user-data. Quality of service is specified by means of QOS-parameters representing characteristics such as throughput, transit delay, accuracy, and reliability.

e) *NS-user-addressing* – The Network Service utilizes a system of addressing (NSAP addressing) which allows NS users to refer unambiguously to one another.

7 Types and classes of Network Service

There are two types of Network Service:

a) a connection-mode service (defined in clause 2); and

b) a connectionless-mode service (defined in clause 3).

For a given instance of communication, the mode of service provided to both NS users is the same (i.e. connection-mode or connectionless-mode). Choice of provision of the connectionless-mode Network Service or the connection-mode Network Service is made in accordance with ISO 7498/Add.1.

When referring to this Service Definition, an NS user or NS provider shall state which types(s) of service it expects to use or provide.

There are no distinct classes of Network Service defined.

However, for the connection-mode service, two Network Layer Services, Receipt Confirmation and Expedited Data Transfer, are NS provider-options.

A service which is an NS provider-option is one which an NS provider can choose either to provide or not to provide for a particular NC. In circumstances where the NS provider chooses not to provide a provider-option service, it will not be available in the Network Service. If the provider-option Receipt Confirmation or Expedited Data Transfer is provided, it shall be as defined in 14.1 to 14.3.

SECTION 2 – DEFINITION OF THE CONNECTION-MODE SERVICE

8 Features of the connection-mode Network Service

The connection-mode Network Service offers the following features to an NS user:

a) the means to establish an NC with another NS user for the purpose of transferring NS-user-data in the form of NSDUs. More than one NC may exist between the same pair of NS users;

b) the establishment of an agreement between the two NS users and the NS provider for a certain QOS associated with each NC;

c) the means of transferring NSDUs in sequence on an NC. The transfer of NSDUs, which consist of an integer number of octets, is transparent, in that the boundaries of NSDUs and the contents of NSDUs are preserved unchanged by the Network Service, and there are no constraints on the NSDU content imposed by the Network Service;

d) the means by which the receiving NS user may flow control the rate at which the sending NS user may send NSDUs;

e) in some circumstances, the means of transferring separate expedited NSDUs in sequence (see clause 7). Expedited NSDUs are limited in length and their transmission is subject to a different flow control from normal data across the NSAP;

f) the means by which the NC can be returned to a defined state and the activities of the two NS users synchronized by use of a reset service;

g) in some circumstances, the means for the NS user to confirm the receipt of an NSDU (see clause 7);

h) the unconditional, and therefore possibly destructive, release of an NC by either of the NS users or by the NS provider.

9 Model of the connection-mode Network Service

9.1 Model of the connection-mode Network Layer Service

This Service Definition uses the abstract model for a layer service defined in clause 4 of CCITT Rec. X.210 | ISO/TR 8509. The model defines the interactions between the NS users and the NS provider which take place at the two NSAPs. Information is passed between the NS user and the NS provider by service primitives, which may convey parameters.

9.2 Model of a Network Connection

Between the two endpoints of an NC, there exists a flow control function which relates the behaviour of the NS user at one end receiving NS-user-data to the ability of the NS user at the other end to send NS-user-data. As a means of specifying this flow control feature and its relationship with other capabilities provided by the Network Service, the queue model of an NC, described in the following subclauses, is used.

This queue model of an NC is discussed only to aid in the understanding of the end-to-end service features perceived by users of the Network Service. It is not intended to serve as a substitute for a precise, formal description of the Network Service, nor as a complete specification of all allowable sequences of NS primitives, (Allowable primitive sequences are specified in clause 11 – also, see Note below.) In addition, this model does not attempt to describe all the functions or operations of Network Layer entities (including relay entities) which are used to provide the Network Service. No attempt to specify or constrain Network Service implementations is implied.

In interpreting this service definition, statements in clauses 12 to 14 concerning the properties of individual primitives have precedence over the general statements in this clause.

NOTE – In addition to the interaction between service primitives described by this model, there may be constraints applied locally on the ability to invoke primitives, as well as service procedures defining particular sequencing constraints on some primitives.

9.2.1 Queue model concepts

The queue model represents the operation of an NC in the abstract by a pair of queues linking the two NSAPs. There is one queue for each direction of information flow (see Figure 1).

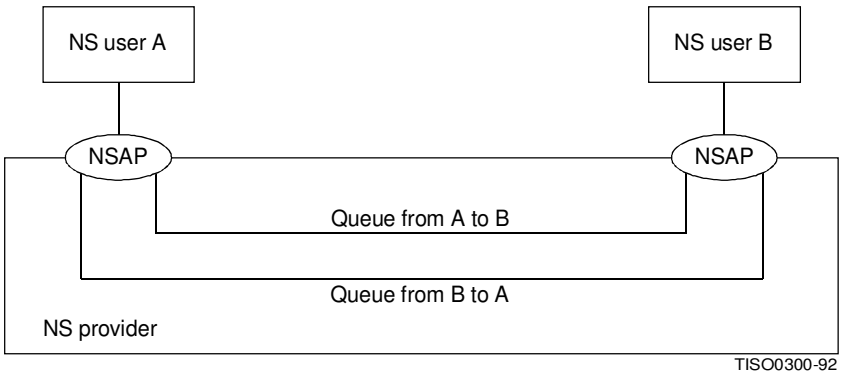

Figure 1 – Queue model of a Network Connection

Each queue represents a flow control function in one direction of transfer. The ability of an NS user to add objects to a queue will be determined by the behaviour of the NS user removing objects from that queue and the state of the queue. Objects are entered or removed from the queue, either as the result of interactions at the two NSAPs, or as the result of NS provider initiatives.

The pair of queues is considered to be available for each potential NC.

The objects which may be placed in a queue as a result of interactions at an NSAP (see clauses 12 to 14) are

a) connect objects (associated with N-CONNECT primitives and all of their parameters);

b) octets of normal NS-user-data (associated with an N-DATA primitive);

c) indications of end-of-NSDU (associated with completion of an N-DATA primitive);

d) expedited NSDUs (associated with N-EXPEDITED-DATA primitives and all their parameters);

e) data acknowledgment objects (associated with N-DATA-ACKNOWLEDGE primitives);

f) reset objects (associated with N-RESET primitives and their parameters);

g) disconnect objects (associated with N-DISCONNECT primitives and all their parameters).

NOTE – The description of flow control (see 9.2.3) requires a less abstract description than that used for describing sequences of primitives in clauses 11 to 14. While primitives are defined to be indivisible, for purposes of this queue model, information associated with N-DATA primitives is conceptually subdivided into a sequence of octets of NS-user-data followed by an end-of-NSDU indication. This does not imply any particular subdivision in any real interface.

The objects which may be placed in a queue as a result of NS provider initiatives (see clauses 12 to 14) are

1) reset objects (associated with N-RESET primitives and all their parameters);

2) synchronization mark objects (see 9.2.4);

3) disconnect objects (associated with N-DISCONNECT primitives and all their parameters).

The queues are defined to have the following general properties:

i) a queue is empty until a connect object has been entered and can be returned to this state, with loss of its contents, by the NS provider (see 9.2.4 and 9.2.5);

ii) objects may be entered into a queue as a result of the actions of the source NS user, subject to control by the NS provider; objects may also be entered into a queue by the NS provider;

iii) objects are removed from the queue under the control of the receiving NS user;

iv) objects are normally removed under the control of the NS user in the same order that they were entered (however see 9.2.3);

v) a queue has a limited capacity, but this capacity is not necessarily either fixed or determinable.

9.2.2 NC establishment

A pair of queues is associated with an NC between two NSAPs when the NS provider receives an N-CONNECT request primitive at one of the NSAPs and a connect object is entered into one of the queues. From the standpoint of one of the NS users of the NC, the queues remain associated with the NC until a disconnect object (associated with an N-DISCONNECT primitive) is either entered or removed from a queue at that NSAP.

If NS user A denotes the NS user who initiates NC establishment (resulting in a connect object being entered into the queue from NS user A to NS user B), then no object other than a disconnect object may be entered into the queue from A to B until after the connect object associated with the N-CONNECT confirm has been removed. In the queue from NS user B to NS user A, objects can be entered only after a connect object associated with an N-CONNECT response from NS user B has been entered; it is possible for a disconnect object to be placed in the queue from B to A instead of a connect object to release the NC.

The properties exhibited by the queues while the NC exists represent the agreements reached among the NS users and the NS provider during the NC establishment procedure concerning quality of service and the use of the receipt and expedited data transfer services.

9.2.3 Data transfer operations

Flow control on the NC is represented in this queue model by the management of the queue capacity, allowing objects of certain types to be added to the queues. The conditions affecting entry of reset and disconnect objects are described in item b) below and in 9.2.4 and 9.2.5. The flow control relationship between the other types of objects is summarized by Table 1.

Table 1 – Flow control relationships between queue model objects

The addition of object x may prevent further addition of object y	Octets of NS-user-data or end-of-NSDU	Expedited NSDU	Data acknowledgment
Octets of normal NS-user-data or end-of-NSDU	Yes	Yes	No
Expedited NSDU	No	Yes	No
Data acknowledgment	No	No	No

Once in the queue, the NS provider may manipulate pairs of adjacent objects, resulting in

a) *Change of order* – the order of any pair of objects may be reversed, if and only if, the following object is of a type defined to be able to advance ahead of the preceding object. No object is defined to be able to advance ahead of another object of the same type.

b) *Deletion* – any object may be deleted if, and only if, the following object is defined to be destructive with respect to the preceding object. If necessary, the last object in the queue will be deleted to allow a destructive object to be entered. Destructive objects may therefore always be added to the queue. Disconnect objects are defined to be destructive with respect to all other objects. Reset objects are defined to be destructive with respect to all other objects except connect and disconnect objects.

The relationships between objects which may be manipulated as described in a) and b) above are summarized in Table 2.

Table 2 — Ordering relationships between queue model objects

Following object x is defined / with respect to preceding object y	Connect	Octets of normal NS user-data	End-of-NSDU	Expedited NSDU	Data acknowledgment	Reset	Synchronization mark	Disconnect
Connect	N/A	—	N/A	—	—	—	N/A	DES
Octets of normal NS user-data	N/A	—	—	AA	AA	DES	N/A	DES
End-of-NSDU	N/A	—	N/A	AA	AA	DES	N/A	DES
Expedited NSDU	N/A	—	—	—	AA	DES	N/A	DES
Data acknowledgment	N/A	—	—	AA	—	DES	N/A	DES
Reset	N/A	—	N/A	—	—	DES	—	DES
Synchronization mark	N/A	—	—	—	—	DES	N/A	DES
Disconnect	N/A	N/A	N/A	N/A	N/A	N/A	N/A	DES

AA Indicates that object x is defined to be able in advance ahead of the preceding object y.

DES Indicates that object x is defined to be destructive with respect to the preceding object y.

— Indicates that object x is neither destructive with respect to object y nor able to advance ahead of object y.

N/A Indicates that object x will not occur in a position succeeding object y in a valid state of a queue.

Whether the NS provider performs actions resulting in change of order and deletion or not will depend upon the behaviour of the NS users and the agreed QOS for the NC. In general, if an NS user does not cause objects to be removed from a queue, the NS provider shall, after some unspecified period of time, perform all permitted actions of types a) and b).

9.2.4 Reset operations

The invocation of a reset procedure is represented in the two queues as follows:

a) Invocation of a reset procedure by the NS provider is represented by the introduction into each queue of a reset object followed by a synchronization mark object.

b) A reset procedure invoked by an NS user is represented by the addition of a reset object to one queue. In this case, the NS provider will insert a reset object followed by a synchronization mark object into the other queue.

The completion of a reset procedure by the issuance of an N-RESET response by an NS user results in a reset object being placed in the queue from the responding NS user.

A synchronization mark object cannot be removed from a queue by an NS user; a queue appears empty to an NS user when a synchronization mark object is the next object in it. Unless destroyed by a disconnect object, a synchronization mark object remains in the queue until the next object following it in the queue is a reset object. Both the synchronization mark object and the following reset object are then deleted by the NS provider.

NOTE – Associated with the invocation of a reset procedure are restrictions on the issuance of certain other types of primitives. These restrictions will result in restrictions on the entry of certain object types into the queue until the reset procedure is complete.

9.2.5 NC release

The insertion into a queue of a disconnect object, which may occur at any time, represents the initiation of an NC release procedure. The release procedure may be destructive with respect to other objects in the two queues and eventually results in the emptying of the queues and the disassociation of the queues with the NC.

The insertion of a disconnect object may also represent the rejection of an NC establishment attempt or the failure to complete NC establishment. In such cases, if a connect object representing an N-CONNECT request primitive is deleted by a disconnect object, then the disconnect object is also deleted. The disconnect object is not deleted when it deletes any other object, including the case where it deletes a connect object representing an N-CONNECT response.

10 Quality of the connection-mode Network Service

The term quality of service (QOS) refers to certain characteristics of an NC as observed between the NC endpoints. QOS describes aspects of an NC which are attributable solely to the NS provider; it can only be properly determined in the absence of NS user behaviour (which is beyond the control of the NS provider) which specifically constrains or impairs the performance of the Network Service.

A value of QOS applies to an entire NC. When determined or measured at both ends of an NC, the QOS observed by the NS users at the two ends of the NC is the same. This is true even in the case of an NC spanning several subnetworks where each subnetwork offers different services.

10.1 Determination of QOS

QOS is described in terms of QOS-parameters. The definition of each of these QOS-parameters specifies the way in which the QOS-parameter's value is measured or determined, making reference where appropriate, to primitive events of the NS.

NOTES

1 It is important to distinguish the use of the term "QOS-parameters" from the more general term "parameters" as defined in 5.2 and used throughout this Service Definition. A "QOS-parameter" refers to a specific aspect or component of the QOS for an NC. As described below, a particular QOS-parameter may or may not be related to a parameter defined as part of a Network Service primitive.

2 For purposes of accuracy and/or convenience, the definition and measurement formula for some QOS-parameters includes a component attributable to the NS user(s). In such cases, to evaluate the QOS attributable solely to the NS provider, this NS user-dependent component must be factored out.

3 The definition of NS QOS-parameters in terms which provide a means for measurement should not be understood to imply that QOS monitoring or that verification of stated QOS value is, or must be, performed by the NS provider or by the NS users.

It is in terms of the NS QOS-parameters that information about QOS is exchanged among the NS provider and NS users.

Information about the QOS requirements of the NS users may be used by the NS provider for purposes such as protocol selection, route determination, and allocation of resources. Information about the QOS available from the NS provider may be used by NS users for purposes such as selecting QOS enhancement mechanisms and determining the QOS values provided to NS users at higher layers.

The NS QOS-parameters can be divided into two categories as follows:

1) those whose values are "conveyed" between peer NS users by means of the NS provider during the Establishment phase of an NC. As part of this conveyance, a three-party "negotiation" among the NS users and the NS provider for the purpose of agreeing upon a particular QOS-parameter value may take place; and

2) those whose values are not "conveyed" or "negotiated" among the NS users and the NS provider. For these QOS-parameters, however, information about the values which is useful to the NS provider and each NS user may be made known by local means.

The NS QOS-parameters are defined in 10.2.1 to 10.2.12 below.

The set of NS QOS-parameters that belong to the first category, and the procedures and constraints that apply to conveying and negotiating those QOS-parameters, are specified in 12.2.7. Once the NC is established, and throughout the lifetime of the NC, the agreed values for these QOS-parameters are not "renegotiated" at any point, and there is no guarantee that the originally negotiated values will be maintained. The NS user should also be aware that, once an NC is established, changes in QOS on the NC are not explicitly signaled in the NS.

For QOS-parameters in the second category, the values for a particular NC are not negotiated, nor are they directly conveyed from NS user to NS user. As a local matter, however, there may be means by which the values of one or more of these QOS-parameters are known and utilized by the NS provider and each NS user. Despite the local nature of particular NS user/NS provider interactions which may occur for the purposes of exchanging QOS-parameter information, the characteristics of an NC which the QOS-parameters describe are applicable and can be observed on a complete NC, end-to-end basis. Thus, in order to give a full characterization of the properties of NCs, the definitions of the entire set of QOS-parameters which apply to the NS, including those classified in category 2, are included in this Service Definition. Other aspects related to category 2 parameters, such as the circumstances of their availability and use, as well as other QOS issues, such as the relationship to OSI management, and multi-layer QOS relationships, are the subjects of other OSI QOS-related specifications.

NOTE 4 – For non-negotiated QOS-parameters associated with the Data Transfer phase of an NC, when specified, a value of such a QOS-parameter applies to both directions of transfer on the NC.

10.2 Definition of QOS-parameters

QOS-parameters can be classified as

a) QOS-parameters which express Network Service performance, as shown in Table 3.

b) QOS-parameters that express other Network Service characteristics, as shown in Table 4.

NOTE – Some QOS-parameters are defined in terms of the issuance of Network Service primitives. Reference to a primitive in 10.2.1 through 10.2.12 refers to the complete execution of that service primitive at the appropriate NSAP.

10.2.1 NC establishment delay

NC establishment delay is the maximum acceptable delay between an N-CONNECT request and the corresponding N-CONNECT confirm primitive.

NOTE – This delay includes a component, attributable to the called NS user, which is the time between the N-CONNECT indication primitive and the N-CONNECT response.

Table 3 – Classification of performance QOS-parameters

Phase	Performance criterion	
	Speed	Accuracy/reliability
NC establishment	NC establishment delay	NC establishment failure probability (misconnection/NC refusal)
Data transfer	Throughput	Residual error rate (corruption, duplication/loss)
		NC resilience
	Transit delay	Transfer failure probability
NC release	NC release delay	NC release failure probability

Table 4 – QOS-parameters not associated with performance

NC protection
NC priority
Maximum acceptable cost

10.2.2 NC establishment failure probability

NC establishment failure probability is the ratio of total NC establishment failures to total NC establishment attempts in a measurement sample.

NC establishment failure is defined to occur when a requested NC is not established within the specified maximum acceptable time period as a result of NS provider behaviour such as misconnection, NC refusal, or excessive delay. NC establishment attempts which fail as a result of NS user behaviour such as error, NC refusal, or excessive delay are excluded in calculating NC establishment failure probability.

10.2.3 Throughput

Throughput is defined, for each direction of transfer, in terms of a sequence of at least two successfully transferred NSDUs presented continuously to the NS provider at the maximum rate the NS provider can continuously sustain, and unconstrained by flow control applied by the receiving NS user.

Given such a sequence of n NSDUs, where n is greater than or equal to 2, the throughput is defined to be the smaller of:

a) the number of NS-user-data octets contained in the last $n - 1$ NSDUs divided by the time between the first and last N-DATA requests in the sequence; and

b) the number of NS-user-data octets contained in the last $n - 1$ NSDUs divided by the time between the first and last N-DATA indications in the sequence.

Successful transfer of the octets in a transmitted NSDU is defined to occur when the octets are delivered to the intended receiving NS user without error, in the proper sequence, and prior to release of the NC by the receiving NS user.

Throughput is specified separately for each direction of transfer. Each throughput specification will specify both the desired "target" value and the minimum acceptable value (i.e., the "lowest quality acceptable") for the NC. (See also 12.2.7.)

10.2.4 Transit delay

Transit delay is the elapsed time between an N-DATA request and the corresponding N-DATA indication. Elapsed time values are calculated only on NSDUs that are successfully transferred.

Successful transfer of an NSDU is defined to occur when the NSDU is transferred from the sending NS user to the intended receiving NS user without error, in the proper sequence, and prior to release of the NC by the receiving NS user.

Specification of transit delay will define a pair of values: the desired "target" value and the maximum acceptable (i.e., the "lowest quality acceptable") value. (See also 12.2.7.) The specified values will be averages and will be based on an NSDU size of 128 octets.

The pair of transit delay values specified for an NC applies to both directions of transfer. That is, the transit delay in each direction is expected to be no worse than that specified.

The transit delay for an individual NSDU may be increased if the receiving NS user exercises flow control. Such occurrences are excluded in calculating both average and maximum transit delay values.

10.2.5 Residual error rate

Residual error rate is the ratio of total incorrect, lost, and duplicate NSDUs to total NSDUs transferred across the NS boundary during a measurement period. The relationship among these quantities is defined, for a particular NS user pair, as shown in Figure 2.

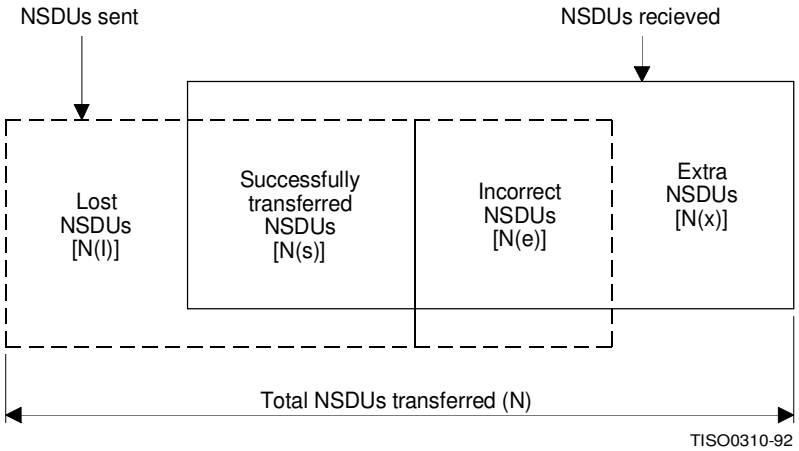

Figure 2 – Components of residual error rate

10.2.6 Transfer failure probability

Transfer failure probability is the ratio of total transfer failures to total transfer samples observed during a performance measurement.

A transfer sample is a discrete observation of NS provider performance in transferring NSDUs between a specified sending and receiving NS user. A transfer sample begins on input of a selected NSDU at the sending NS user boundary, and continues until the outcome of a given number of NSDU transfer requests has been determined. A transfer sample will normally correspond to the duration of an individual NC.

A transfer failure is a transfer sample in which the observed performance is worse than a specified minimum acceptable level. Transfer failures are identified by comparing the measured values for the supported performance parameters with specified transfer failure thresholds. The three supported performance parameters are throughput, transit delay, and residual error rate.

In systems where Network Service QOS is reliably monitored by the NS provider, transfer failure probability can be estimated by the probability of an NS provider invoked N-DISCONNECT during a transfer sample.

10.2.7 NC resilience

NC resilience parameters specify the probability of

a) an NS provider invoked NC release (i.e. issuance of an N-DISCONNECT indication with no prior N-DISCONNECT request); and

b) an NS provider invoked reset (i.e. issuance of an N-RESET indication with no prior N-RESET request);

during a specified time interval on an established NC.

10.2.8 NC release delay

NC release delay is the maximum acceptable delay between an NS user invoked N-DISCONNECT request and the successful release of the NC at the peer NS user. NC release delay is normally specified independently for each NS user. NC release delay does not apply in cases where NC release is invoked by the NS provider.

Issuance of an N-DISCONNECT request by either NS user starts the counting of NC release delay for the other NS user. Successful NC release is signaled to the NS user not initiating the N-DISCONNECT request by an N-DISCONNECT indication.

10.2.9 NC release failure probability

NC release failure probability is the ratio of total NC release requests resulting in release failure to total NC release requests included in a measurement sample. NC release failure probability is normally specified independently for each NS user.

A release failure is defined to occur, for a particular NS user, if that user does not receive an N-DISCONNECT indication within the specified maximum NC release delay of the NS user issuing the N-DISCONNECT request (given that the former NS user has not issued an N-DISCONNECT request).

10.2.10 NC protection

NC protection is the extent to which an NS provider attempts to prevent unauthorized masquerading or monitoring or manipulation of NS-user-data. NC protection for an NC is specified by selecting any combination of the following features:

a) confidentiality of an entire NSDU sequence on the NC;

b) detection of modification, deletion, replay, or insertion of data within the NSDU sequence on an NC;

c) peer entity authentication. The NS user may request that the NS provider should confirm the identity of the remote NSAP such that there is protection against masquerading by T-entities;

d) authentication of the origin of an NSDU such that there is protection against the unauthorized insertion or replay of the NSDU.

10.2.11 NC priority

NC priority specifies independently the relative importance of an NC with respect to the following:

a) priority to gain an NC;

b) priority to keep an NC;

c) priority of data on the NC.

NC priority QOS-parameters a) and b) together define the order in which NCs are to be broken to recover resources if necessary. The NS provider is required to accept new requests for NCs with a high priority type a) if it can, even if NCs with a lower priority type b) have to be released to do so.

NC priority QOS-parameter c) defines the order in which NCs are to have their QOS degraded. The NCs with a high priority type c) are to have their requests serviced within the required QOS first and remaining resources are then used to attempt to satisfy requests on lower priority NCs.

NOTE – The use or abuse of the NC priority QOS-parameters can be controlled by one or more of the following:

– user discipline within a closed group of NS users;

– differential tariffs;

– management facilities within the Network Layer such that requests for NC priority are policed and regulated.

10.2.12 Maximum acceptable cost

The maximum acceptable cost QOS-parameter specifies the maximum acceptable cost for an NC. The cost may be specified in absolute or relative costs units. The cost of an NC is composed of communications and end-system resource costs.

NOTE – The possible actions of the NS provider in the event that the maximum acceptable cost for an NC is exceeded are not specified in this Service Definition.

11 Sequence of primitives

This clause defines the constraints on the sequences in which the primitives defined in clauses 12 to 14 may occur. The constraints determine the order in which primitives occur, but do not fully specify when they may occur. Other constraints, such as flow control of data, will affect the ability of an NS user or an NS provider to issue a primitive at any particular time.

Table 5 is a summary of the NS primitives and their parameters.

11.1 Relation of primitives at the two NC end points

A primitive issued at one NC end point will, in general, have consequences at the other NC end point. The relations of primitives of each type to primitives at the other NC end point are defined in the appropriate clauses 12 to 14; all these relations are summarized in the diagrams in Figure 3.

However, an N-DISCONNECT request or indication primitive may terminate any of the other sequences before completion. An N-RESET request or indication may terminate a data transfer, expedited data transfer, or receipt confirmation sequence before completion.

11.2 Sequence of primitives at one NC endpoint

The possible overall sequences of primitives at an NC endpoint are defined in the state transition diagram, Figure 4. In the diagram:

a) a primitive which is not shown as resulting in a transition (from one state to the same state, or from one state to a different state) is not permitted in that state (however, see 11.1 above concerning the effect of N-DISCONNECT and N-RESET primitives);

b) N-DISCONNECT stands for either the request or the indication form of the primitive in all cases;

c) the labeling of the states NS user invoked reset pending (state 5) and NS provider invoked reset pending (state 6) indicates the party which started the local interaction, and does not necessarily reflect the value of the originator parameter in the associated N-RESET primitive;

d) the Idle state (state 1) reflects the absence of an NC. It is the initial and final state of any sequence, and once it has been re-entered, the NC is released;

e) the use of a state transition diagram to describe the allowable sequences of service primitives does not impose any requirements or constraints on the internal organization of any implementations of the Network Service.

Table 5 – Summary of Network Service primitives and parameters

Phase	Service	Primitive	Parameters
NC establishment	NC establishment	N-CONNECT request	(Called address, calling address, receipt confirmation selection, expedited data selection, QOS-parameter set, NS-user-data)
		N-CONNECT indication	(Called address, calling address, receipt confirmation selection, expedited data selection, QOS-parameter set, NS-user-data)
		N-CONNECT response	(Responding address, receipt confirmation selection, expedited data selection, QOS-parameter set, NS-user-data)
		N-CONNECT confirm	(Responding address, receipt confirmation selection, expedited data selection, QOS-parameter set, NS-user-data)
Data transfer	Data transfer	N-DATA request	(NS-user-data, confirmation request)
		N-DATA indication	(NS-user-data, confirmation request)
	Receipt confirmation (see Note)	N-DATA-ACKNOWLEDGE request•	–
		N-DATA-ACKNOWLEDGE indication	–
	Expedited data transfer (see Note)	N-EXPEDITED-DATA request	(NS-user-data)
		N-EXPEDITED-DATA indication	(NS-user-data)
	Reset	N-RESET request	(Reason)
		N-RESET indication	(Originator, reason)
		N-RESET response	–
		N-RESET confirm	–
NC release	NC release	N-DISCONNECT request	(Reason, NS-user-data, responding address)
		N-DISCONNECT indication	(Originator, reason, NS-user-data, responding address)

NOTE – An NS provider-option service: it need not be provided in every Network Service.

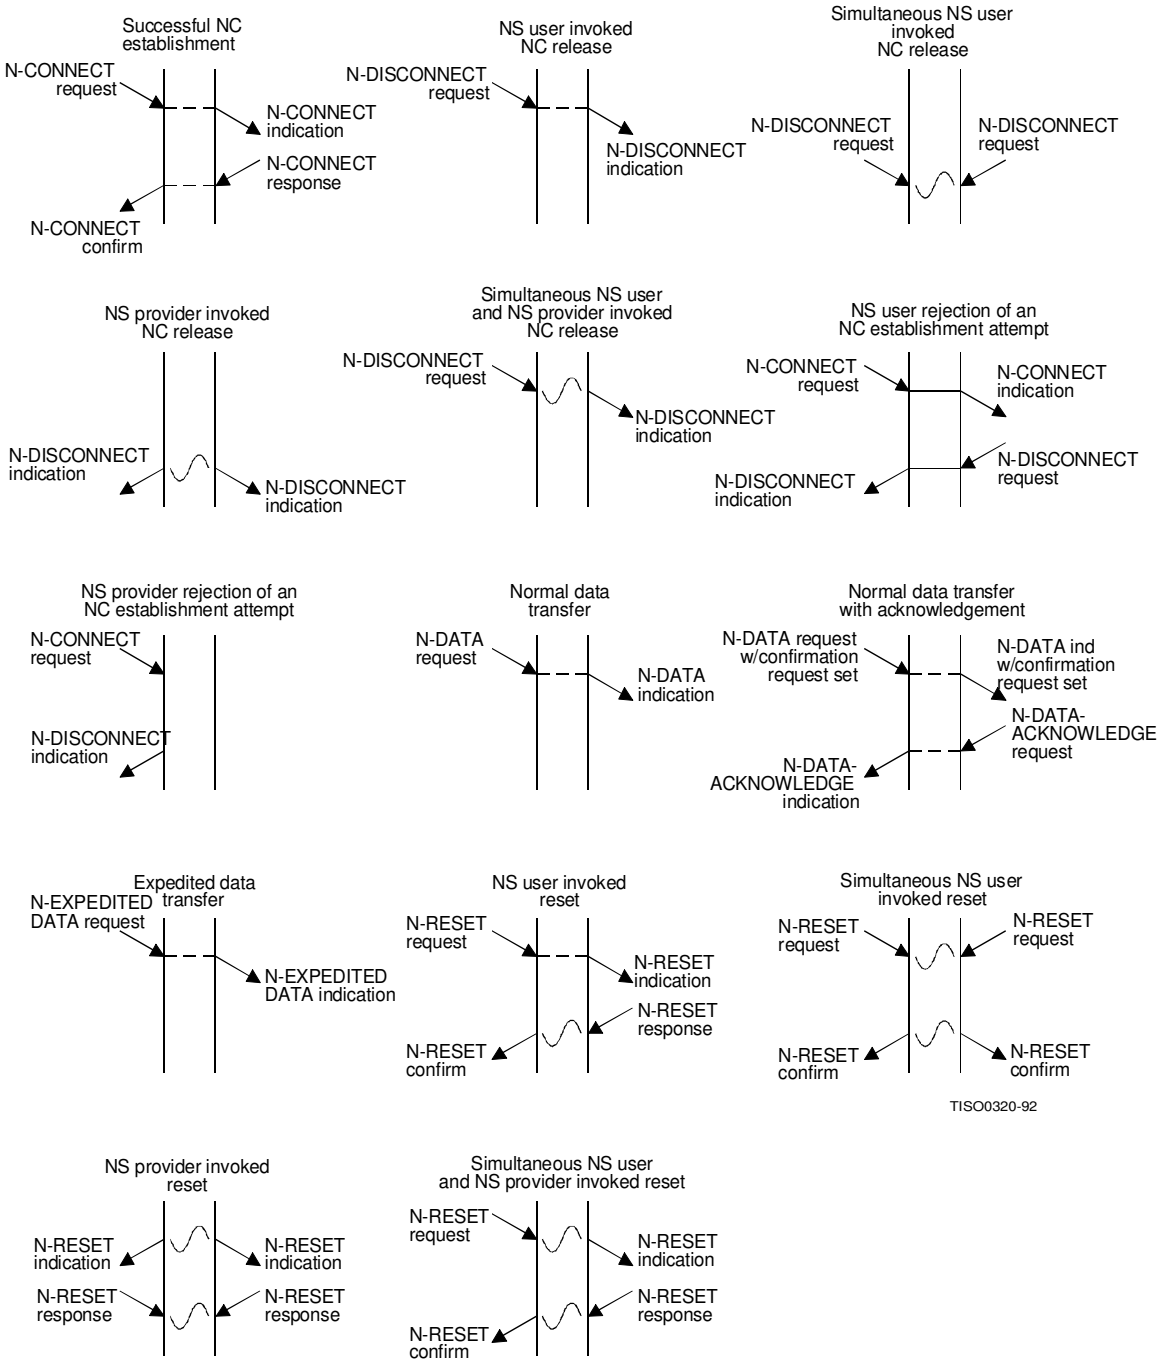

Figure 3 – Summary of Network Service primitive time sequence diagrams

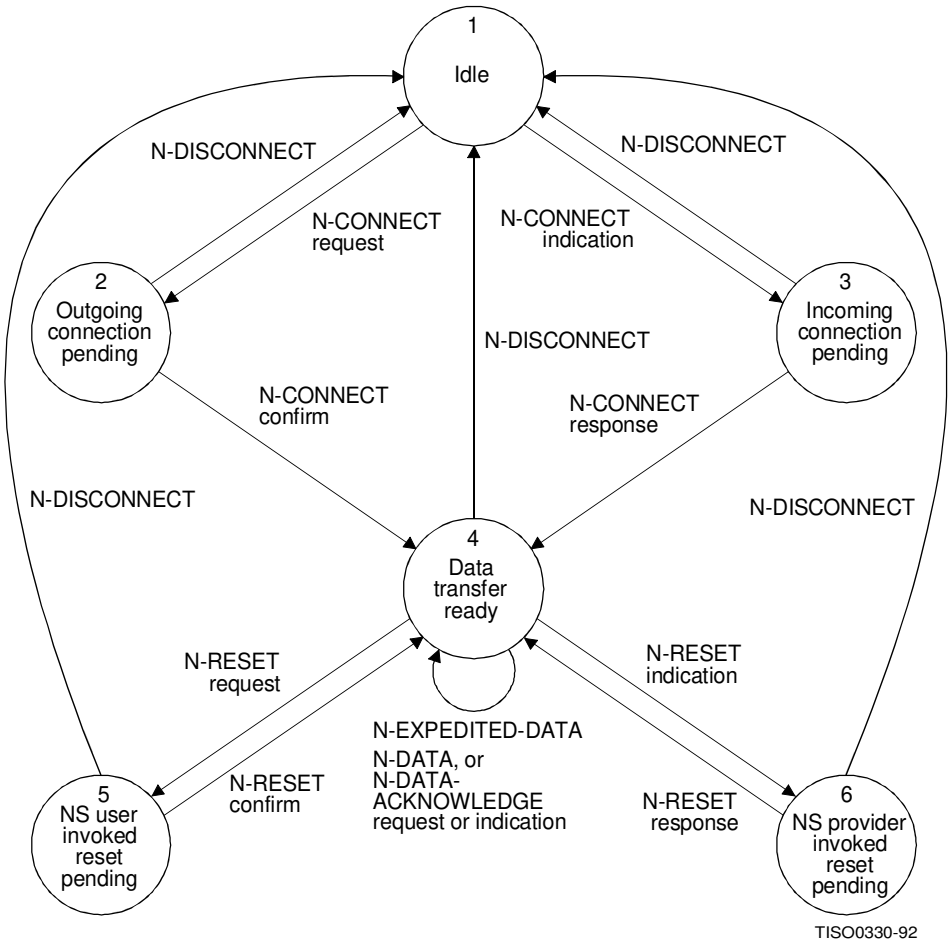

Figure 4 – State transition diagram for sequences of primitives at an NC endpoint

12 Connection establishment phase

12.1 Function

The NC establishment service primitives can be used to establish an NC, provided the NS users exist and are known to the NS provider.

Simultaneous N-CONNECT requests at the two NSAPs are handled independently by the NS provider; they may result in two, one or zero NCs.

12.2 Types of primitives and parameters

Table 6 indicates the types of primitives and the parameters needed for NC establishment.

12.2.1 Addresses

The parameters which take addresses as values (see 12.2.2 to 12.2.4) all refer to NSAP addresses. The NSAP address parameters will accommodate variable length addresses up to a defined maximum. Network Layer Addressing is specified in Annex A.

The values of these addresses as supplied by the NS user are not necessarily checked or authenticated by the NS provider. An NS user receiving these addresses in N-CONNECT indication or confirm primitives can only rely on their validity if the NS user has knowledge that the NS provider guarantees address correctness.

NOTE – Mechanisms operating within the NS provider, such as call redirection or resolution of generic addresses, may result in address parameters in corresponding primitives not being identical in the following cases:

a) the responding address parameter on the N-CONNECT response may not necessarily be the same as the called address parameter on the N-CONNECT indication;

b) the responding address parameter on the N-CONNECT confirm may not necessarily be the same as the called address parameter on the N-CONNECT request.

Table 6 – NC establishment primitives and parameters

Primitive / Parameter	N-CONNECT request	N-CONNECT indication	N-CONNECT response	N-CONNECT confirm
Called address	X	X(=) (see Note)		
Calling address	X (see Note)	X(=)		
Responding address			X (see Note)	X(=)
Receipt confirmation selection	X	X	X	X(=)
Expedited data selection	X	X	X	X(=)
QOS-parameter set	X	X	X	X(=) for ISO/IEC 8348 X(C=) for X.213
NS-user-data	X for ISO/IEC 8348 X(C) for X.213	X for ISO/IEC 8348 X(C=) for X.213	X for ISO/IEC 8348 X(C) for X.213	X(=) for ISO/IEC 8348 X(C=) for X.213
NOTE – This parameter may be implicitly associated with the NSAP at which the primitive is issued.				

12.2.2 Called address parameter

The called address parameter conveys an address identifying the NSAP to which the NC is to be established. Where explicitly supplied, the addresses in corresponding N-CONNECT request and indication primitives are identical.

12.2.3 Calling address parameter

The calling address parameter conveys the address of the NSAP from which the NC has been requested. Where explicitly supplied, the addresses in corresponding N-CONNECT request and indication primitives are identical.

12.2.4 Responding address parameter

The responding address parameter conveys the address of the NSAP to which the NC has been established. Where explicitly supplied, the addresses in corresponding N-CONNECT response and confirm primitives are identical. This parameter always conveys a specific NSAP address and not a generic NSAP address.

12.2.5 Receipt confirmation selection parameter

The receipt confirmation selection parameter indicates the use/availability of the receipt confirmation service on the NC. If the receipt confirmation service is not provided in the Network Service, then it cannot be used on the NC (see clause 7). The value of this parameter is either "use of receipt confirmation" or "no use of receipt confirmation". The values on the various primitives are related such that

a) on the N-CONNECT request, either of the defined values may occur;

b) on the N-CONNECT indication, the value is either equal to the value on the request primitive, or is "no use of receipt confirmation";

c) on the N-CONNECT response, the value is either equal to the value on the indication primitive or is "no use of receipt confirmation";

d) on the N-CONNECT confirm, the value is equal to the value on the response primitive.

Since receipt confirmation need not be provided in the Network Service and since, when it is available, both NS users and the NS provider must agree to its use, there are four possible cases of negotiation of receipt confirmation on an NC:

i) the calling NS user does not request it – it is not used;

ii) the calling NS user requests it but the NS provider does not provide it – it is not used;

iii) the calling NS user requests it and the NS provider agrees to provide it, but the called NS user does not agree to its use – it is not used;

iv) the calling NS user requests it, the NS provider agrees to provide it, and the called NS user agrees to its use – it can be used.

12.2.6 Expedited data selection parameter

The expedited data selection parameter indicates the use/availability of the expedited data transfer service on the NC. If the expedited data transfer service is not available from the NS provider (see clause 7), then it cannot be used on the NC. The value of this parameter is either "use of expedited data" or "no use of expedited data". The values on the various primitives are related such that:

a) on the N-CONNECT request, either of the defined values may occur;

b) on the N-CONNECT indication, the value is either equal to the value on the request primitive, or is "no use of expedited data";

c) on the N-CONNECT response, the value is either equal to the value on the indication primitive or is "no use of expedited data";

d) on the N-CONNECT confirm, the value is equal to the value on the response primitive.

12.2.7 QOS-parameter set

For each QOS-parameter which is conveyed during NC establishment, a set of "subparameters" is defined from among the following possibilities:

i) a "target" value which is the QOS value desired by the calling NS user;

ii) the "lowest quality acceptable" value which is the lowest QOS value agreeable to the calling NS user;

iii) an "available" value which is the QOS value the NS provider is willing to provide; and

iv) a "selected" value which is the QOS value to which the called NS user agrees.

The set of values which can be specified for each subparameter is defined in every Network Service. Each set of values includes the value "unspecified". It may also include a value defined to be a "default" value, which is mutually understood by the NS provider and the NS user between which it is conveyed.

NOTE 1 – "Default" values are defined between a particular NS user and the NS provider. Different "defaults" may exist for different NS users and thus a value which is understood as a "default" at one end of an NC may not be the "default" value at the other end.

In those cases where both the subparameters "target" and "lowest quality acceptable" are specified by the calling NS user, they are boundary parameters defining a range of QOS values to which the calling NS user will agree. Similarly, where both the subparameters "available" and "lowest quality acceptable" are specified by the NS provider, they are boundary parameters defining a range of QOS values which the NS provider is willing to provide. These ranges are

defined to include the values of both of the boundary subparameters, plus any values allowed for these subparameters which lie between the boundary subparameters. In the case where the "target" (or the "available") subparameter has a specified value but the "lowest quality acceptable" value is "unspecified", the rage is defined to consist of the "target" value plus all other values which are allowed for these subparameters and which are lower (in QOS terms) than the "target". If the value for both the "target" and "lowest quality acceptable" is "unspecified", then no range of values is defined.

NOTE 2 – For other value assignments (e.g. "target" is "unspecified" but "lowest quality acceptable" has a specified value), the range is not defined since these assignments are not allowed in the negotiation procedures described in 12.2.7.1 and 12.2.7.4.

12.2.7.1 Throughput

Table 7 indicates the presence of the QOS-subparameters for the throughput QOS-parameters in the N-CONNECT primitives.

The negotiation and conveyance of each of the two throughput QOS-parameters are conducted as follows:

a) In the N-CONNECT request primitive, the calling NS user specifies values for the "target" and "lowest quality "acceptable" (i.e. lowest throughput) subparameters. Permitted value assignments are:

   Case 1: both the "target" and "lowest quality acceptable" are "unspecified";

   Case 2: values other than "unspecified" are specified for both "target" and "lowest quality acceptable";

   Case 3: a value other than "unspecified" is specified for the "target" and the "lowest quality acceptable" is "unspecified".

   NOTE – The case where "target" is "unspecified" and the "lowest quality acceptable" has a value other than "unspecified" is not permitted; logically, this case can be represented by the permitted assignment where an identical value is specified for both the "target" and "lowest quality acceptable" (case 2).

b) If the value assignment of the "target" and "lowest quality acceptable" subparameters are as defined in case 1, then the NS provider determines the highest QOS throughput value which is to be offered on the NC. This value (which may be the "default" value understood by the NS provider and the called NS user) is specified as the "available" subparameter in the N-CONNECT indication while the "lowest quality acceptable" subparameter value is "unspecified". If the requested QOS value assignments are as defined in case 2 or case 3, then, if the NS provider does not agree to provide a QOS in the requested range, the NC establishment attempt is rejected as described in 13.5. If the NS provider does agree to provide a QOS in the requested range, then in the N-CONNECT indication, the "available" subparameter specifies the highest QOS value within the range which the NS provider is willing to provide and the "lowest quality acceptable" subparameter value is identical to that of the "lowest quality acceptable" subparameter in the N-CONNECT request.

c) If the called NS user does not agree to a QOS in the range between the "available" and the "lowest quality acceptable" subparameters of the N-CONNECT indication then the NS user rejects the NC establishment attempt as described in 13.4.

d) If the called NS user does agree to a QOS in the specified range, then the NS user specifies the agreed to value in the "selected" parameter of the N-CONNECT response.

e) In the N-CONNECT confirm, the "selected" subparameter has a value identical to that of "selected" in the N-CONNECT indication.

A summary of the negotiation procedures for the throughput QOS-subparameters is contained in Table 8.

12.2.7.2 Transit delay

NOTE – This is not only applicable to CCITT Rec. X.213: The implementation of the transit delay negotiation requires urgent further study in order to have a harmonized realization in different types of sub-networks. Special attention is requires as regards routing and charging consequences.

Table 9 indicates the presence of the QOS-subparameters for the transit delay QOS-parameter in the N-CONNECT primitives.

Table 7 – Negotiated QOS-subparameters for throughput QOS-parameters

Primitive / Parameter	N-CONNECT request	N-CONNECT indication	N-CONNECT response	N-CONNECT confirm
Throughput 1 “target” (calling to called)	X			
Throughput 1 “lowest quality acceptable” (calling to called)	X	X(=)		
Throughput 2 “target” (called to calling)	X			
Throughput 2 “lowest quality acceptable” (called to calling)	X	X(=)		
Throughput 1 “available” (calling to called)		X		
Throughput 2 “available” (called to calling)		X		
Throughput 1 “selected” (calling to called)			X	X(=)
Throughput 2 “selected” (called to calling)			X	X(=)

The negotiation and conveyance of the transit delay QOS-parameter are conducted as follows:

a) In the N-CONNECT request primitive, the calling NS user specifies values for the “target” and “lowest quality acceptable” (i.e. highest acceptable transit delay) subparameters. Permitted value assignments are:

Case 1: both the “target” and “lowest quality acceptable” are “unspecified”;

Case 2: values other than “unspecified” are specified for both “target” and “lowest quality acceptable”;

Case 3: a value other than “unspecified” is specified for the “target” and the “lowest quality acceptable” is “unspecified”.

NOTE – The case where “target” in “unspecified” and the “lowest quality acceptable” has a value other than “unspecified” is not permitted; logically this case can be represented by the permitted assignment where an identical value is specified for both the “target” and “lowest quality acceptable”.

b) If the value assignments of the “target” and “lowest quality acceptable” subparameters are as defined in case 1, then the NS provider determines the transit delay value to be offered on the NC and specifies it as the “available” subparameter in the N-CONNECT indication.

If the value assignments are as defined in case 2 or case 3, then if the NS provider does not agree to provide a QOS in the requested range, the NC establishment attempt is rejected as described in 13.5. If the NS provider does agree to provide a QOS in the requested range, the “available” subparameter in the N-CONNECT indication specifies the value of QOS which is offered.

c) If the called NS user does not agree to the QOS specified as “available”, the NS user rejects the NC establishment attempt as described in 13.4.

d) If the called NS user does agree to the "available" QOS, then the NS user issues an N-CONNECT response (the N-CONNECT response does not convey any transit delay QOS-subparameters).

e) In the N-CONNECT confirm the "selected" subparameter value is identical to that specified as "available" in the N-CONNECT indication.

A summary of the negotiation procedures for the transit delay QOS-subparameters is contained in Table 10.

12.2.7.3 NC Protection

The values and meaning of the NC Protection QOS parameter are given in 10.2.10. Table 11 indicates the presence of the QOS-subparameter for the NC Protection QOS-parameter in the N-CONNECT primitives.

The negotiation and conveyance of the NC Protection QOS parameter is conducted as follows:

a) In the N-CONNECT request primitive, the calling NS user specifies values for the "Target" and "Lowest Quality Acceptable" subparameters; permitted value assignments are:

Case 1: both the "Target" and "Lowest Quality Acceptable" are "unspecified";

Case 2: values other than "unspecified" are specified for both "Target" and "Lowest Quality Acceptable";

Case 3: a value other than "unspecified" is specified for the "Target" and the "Lowest Quality Acceptable" is "unspecified".

NOTE – The case where "Target" is "unspecified" and the "Lowest Quality Acceptable" has a value other than "unspecified" is not permitted; logically, this case can be represented by the permitted assignment where an identical value is specified for both the "Target" and "Lowest Quality Acceptable" (case 2).

b) If the NS provider does not support a choice of NC Protection levels then the value of the "Target" subparameter is conveyed by the NS provider and passed to the called NS user unchanged as the "Available" subparameter in the N-Connect indication.

c) If the NS provider does support a choice of NC Protection levels, then

1) In case 1

the NS provider determines the QOS value to be offered on the NC and specifies it in the "Available" subparameter in the N-CONNECT indication.

2) In cases 2 and 3

if the NS provider does not agree to provide a QOS in the requested range, then the NC establishment attempt is rejected as described in 13.5. If the NS provider does agree to provide a QOS in the requested range, then in the N-CONNECT indication, the "Available" subparameter specifies the highest QOS value within the range which the NS provider is willing to provide.

d) The value of the "Lowest Quality Acceptable" subparameter in the N-CONNECT indication is identical to that in the N-CONNECT request.

e) If the value of the "Available" subparameter of the N-CONNECT indication is "unspecified" then

1) If the called NS user does not agree to accept establishment of a connection with this unspecified quality, the NS user rejects the NC establishment attempt as described in 13.4.

2) If the called NS user does agree, then the NS user specifies the value "unspecified" in the "Selected" subparameter of the N-CONNECT response.

NOTE – When connection is established with the value of "unspecified" selected, it follows that the QOS provided may be at any level at the discretion of the Network Service provider. Consequently, the called NS user would agree to such a connection only if any level of QOS, even the lowest, is acceptable.

f) If the value of the "Available" subparameter in the N-CONNECT indication is not "unspecified" then

1) If the called NS user does not agree to a QOS in the range identified by the "Available" and "Lowest Quality Acceptable" subparameters of the N-CONNECT indication then the NS user rejects the NC establishment attempt as described in 13.4.

2) If the called NS user does agree to a QOS in the identified range, then the NS user specifies the agreed value in the "Selected" subparameter of the N-CONNECT response.

g) In the N-CONNECT confirm, the "Selected" subparameter has a value identical to that of "Selected" in the N-CONNECT response.

Table 8 — Negotiation of throughput QOS-subparameters

	Calling NS user specifies in N-CONNECT request		NS provider specifies in N-CONNECT indication		Called NS user specifies in N-CONNECT response	NS provider specifies in N-CONNECT confirm	Notes
	"Target"	"Lowest quality acceptable"	"Available"	"Lowest quality acceptable"	"Selected"	"Selected"	
Case 1	"Unspecified"	"Unspecified"	Z	"Unspecified"	A	A	Z may be a "default" value $Z \geq A > 0$
Case 2	X	Y	Z	Y	A	A	X and/or Y may be defined to be the "default" value at the calling NS user end, called NS user end, or both $X \geq Z \geq Y$; $Z \geq A \geq Y$
Case 3	X	"Unspecified"	Z	"Unspecified"	A	A	X may be a "default" value $X \geq Z > 0$; $Z \geq A > 0$

Table 9 – Negotiated QOS-subparameters for transit delay QOS-parameter

Primitive / Parameter	N-CONNECT request	N-CONNECT indication	N-CONNECT response	N-CONNECT confirm
Transit delay "target"	X			
Transit delay "lowest quality acceptable"	X			
Transit delay "available"		X		
Transit delay "selected"				X

12.2.7.4 NC Priority

The values and meaning of NC Priority QOS parameters are specified in 10.2.11. This subclause specifies the conveyance of these parameters and applies to each of the three independent aspects of NC Priority defined in 10.2.11.

Table 12 indicates the presence of the QOS-subparameters for the NC Priority QOS-parameter in the N-CONNECT primitives.

The conveyance of NC Priority QOS-parameter is conducted as follows:

a) In the N-CONNECT request primitive, the calling NS user specifies values for the "Target" and "Lowest Quality Acceptable" subparameters; permitted value assignments are:

Case 1: both the "Target" and "Lowest Quality Acceptable" are "unspecified";

Case 2: values other than "unspecified" are specified for both "Target" and "Lowest Quality Acceptable";

Case 3: a value other than "unspecified" is specified for the "Target" and the "Lowest Quality Acceptable" is "unspecified".

NOTE – The case where "Target" is "unspecified" and the "Lowest Quality Acceptable" has a value other than "unspecified" is not permitted; logically this case can be represented by the permitted assignment where an identical value is specified for both the "Target" and "Lowest Quality Acceptable" (case 2).

b) If the NS provider does not support a choice of NC Priority levels then the value of the "Target" subparameter is conveyed by the NS provider and passed to the called NS user unchanged as the "Available" subparameter in the N-CONNECT indication.

c) If the NS provider does support a choice of NC Priority levels, then

1) In case 1

the NS provider determines the QOS value to be offered on the NC and specifies it in the "Available" subparameter in the N-CONNECT indication.

2) In cases 2 and 3

if the NS provider does not agree to provide a QOS in the requested range, then the NC establishment attempt is rejected as described in 13.5. If the NS provider does agree to provide a QOS in the requested range, then in the N-CONNECT indication, the "Available" subparameter specifies the highest QOS value within the range which the NS provider is willing to provide.

d) The value of the "Lowest Quality Acceptable" subparameter in the N-CONNECT indication is identical to that in the N-CONNECT request.

e) If the value of the "Available" subparameter of the N-CONNECT indication is "unspecified" then

1) If the called NS user does not agree to accept establishment of a connection with this unspecified quality, the NS user rejects the NC establishment attempt as described in 13.4.

2) If the called NS user does agree, then the NS user specifies the value "unspecified" in the "Selected" subparameter of the N-CONNECT response.

NOTE – When connection is established with the value "unspecified" selected, it follows that the QOS provided may be at any level at the discretion of the Network Service provider. Consequently, the called NS user would agree to such a connection only if any level of QOS, even the lowest, is acceptable.

f) If the value of the "Available" subparameter in the N-CONNECT indication is not "unspecified" then

1) If the called NS user does not agree to a QOS in the range identified by the "Available" and "Lowest Quality Acceptable" subparameters of the N-CONNECT indication then the NS user rejects the NC establishment attempt as described in 13.4.

2) If the called NS user does agree to a QOS in the identified range, then the NS user specifies the agreed value in the "Selected" subparameter of the N-CONNECT response.

g) In the N-CONNECT confirm, the "Selected" subparameter has a value identical to that of "Selected" in the N-CONNECT response.

12.2.8 NS-user-data parameter

The NS-user-data parameter allows the transfer of NS-user-data between NS users, without modification by the NS provider. The NS user may send any integer number of octets of NS-user-data between zero and 128 octets inclusive.

NOTE – This note is only applicable to CCITT Rec. X.213: The objective is to make this parameter a mandatory parameter to be supported by all sub-networks in the future. However, a number of existing sub-networks cannot support it now. During the interim period, while these sub-networks exist and are not modified to provide this parameter, it is considered as a provider-option. No negotiation mechanism is needed in the Network Service. Limiting, in some sub-networks, the length of NS-user-data to be provided to a value lower than 128 octets (e.g., 16 to 32 octets) for an interim period would imply fewer changes to existing interfaces and signaling systems and would simplify the introduction of such a service in existing sub-networks.

12.3 Sequence of primitives

The sequence of primitives in a successful NC establishment is defined by the time sequence diagram in Figure 5.

The NC establishment procedure may fail either due to the inability of the NS provider to establish an NC or due to the unwillingness of the called NS user to accept an N-CONNECT indication (for these cases, see NC release service, in 13.4 and 13.5). In addition, the NC establishment attempt may be aborted by the NS provider or either of the NS users at any other time before the issuing of the N-CONNECT confirm.

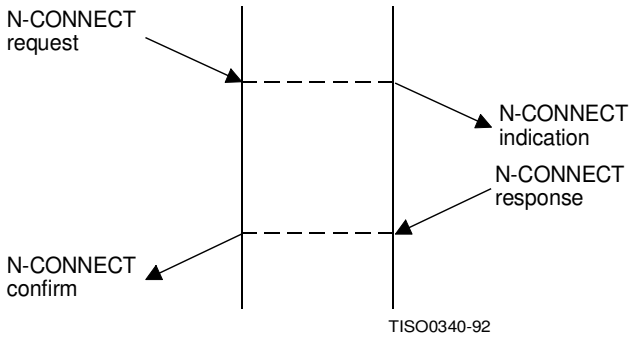

Figure 5 – Sequence of primitives in successful NC establishment

Table 10 — Negotiation of transit delay QOS-subparameters

	Calling NS user specifies in N-CONNECT request		NS provider specifies in N-CONNECT indication	Called NS user specifies in N-CONNECT response	NS provider specifies in N-CONNECT confirm	Notes
	"Target"	"Lowest quality acceptable"	"Available"		"Selected"	
Case 1	"Unspecified"	"Unspecified"			Z	
Case 2	X	Y	Z		Z	X and/or Y may be a "default" value $X \leq Z \leq Y$
Case 3	X	"Unspecified"	Z		Z	X may be a "default" value $X \leq Z \leq \infty$

Table 11 – Negotiated QOS-subparameter for NC Protection QOS-parameters

Parameter \ Primitive	N-CONNECT request	N-CONNECT indication	N-CONNECT response	N-CONNECT confirm
NC Protection "Target"	X			
NC Protection "Lowest Quality Acceptable"	X	X(=)		
NC Protection "Available"		X		
NC Protection "Selected"			X	X(=)

Table 12 – Negotiated QOS-subparameters for NC Priority QOS-parameter

Parameter \ Primitive	NC-CONNECT request	NC-CONNECT indication	NC-CONNECT response	NC-CONNECT confirm
NC Priority "Target"	X			
NC Priority "Lowest Quality Acceptable"	X	X(=)		
NC Priority "Available"		X		
NC Priority "Selected"			X	X(=)

13 Connection release phase

13.1 Function

The NC release service primitives are used to release an NC. The NC release may be performed:

a) by either or both of the NS users to release an established NC;

b) by the NS provider to release an established NC. All failures to maintain an NC are indicated in this way;

c) by the called NS user to reject an N-CONNECT indication;

d) by the NS provider to indicate its inability to establish a requested NC.

NC release is permitted at any time regardless of the current phase of the NC. Once an NC release procedure has been invoked, the NC will be released; a request for NC release cannot be rejected. After NC release has been invoked at one NC endpoint, the NS provider may discard any normal or expedited NS-user-data that has not yet been delivered at the other NC endpoint and may cause any uncompleted sequence of primitives for NC establishment, receipt confirmation, or reset to remain uncompleted.

13.2 Types of primitive and parameters

Table 13 indicates the types of primitives and the parameters needed for NC Release.

Table 13 – NC release primitives and parameters

Primitive / Parameter	N-DISCONNECT request	N-DISCONNECT indication
Originator		X
Reason	X	X
NS-user-data	X for ISO/IEC 8348 X(C) for X.213	X(C=)
Responding address	X(C) (see Note)	X(C=)
NOTE – This parameter may be implicitly associated with the NSAP at which the primitive is issued.		

13.2.1 Originator parameter

The originator parameter indicates the source of the NC release. Its value indicates either the "NS user", "NS provider", or "undefined".

NOTE – The value "undefined" is not permitted when an N-DISCONNECT indication is issued by an NS user or an NS provider in order to reject an NC establishment attempt (see 13.4 and 13.5).

13.2.2 Reason parameter

The reason parameter gives information about the cause of the NC release. The value conveyed in this parameter will be as follows:

a) When the originator parameter indicates an NS provider invoked release, the value is one of:
   1) disconnection-permanent condition;
   2) disconnection-transient condition;
   3) connection rejection-NSAP address unknown (permanent condition);
   4) connection rejection-NSAP unreachable/transient condition;
   5) connection rejection-NSAP unreachable/permanent condition;
   6) connection rejection-QOS not available/permanent condition;
   7) connection rejection-QOS not available/transient condition;
   8) connection rejection-reason unspecified/permanent condition;
   9) connection rejection-reason unspecified/transient condition;

b) When the originator parameter indicates an NS user invoked release, the value is one of:
   1) disconnection-normal condition;
   2) disconnection-abnormal condition;
   3) connection rejection-permanent condition;
   4) connection rejection-transient condition;
   5) connection rejection-QOS not available/transient condition;

6) connection rejection-QOS not available/permanent condition;

7) connection rejection-incompatible information in NS-user-data.

c) When the originator parameter value is “undefined”, then the value of the reason parameter shall also be “undefined”.

13.2.3 NS-user-data parameter

The NS-user-data parameter allows the transfer of NS-user data between NS users, without modification by the NS provider. An NS user invoking NC release may send any integer number of octets of NS-user-data between zero and 128 inclusive. In an N-DISCONNECT indication, this parameter can have a non-zero number of octets of NS-user-data only if the originator parameter has the value of “NS user”.

The NS-user-data sent is lost if NC release is simultaneously invoked by either the NS provider or the intended receiving NS user (see 13.3).

NOTE – This note is only applicable to CCITT Rec. X.213: The objective is to make this parameter a mandatory parameter to be supported by all sub-networks in the future. However, a number of existing sub-networks cannot support it now. During the interim period, while these sub-networks exist and are not modified to provide this parameter, it is considered as a provider-option. No negotiation mechanism is needed in the Network Service.

13.2.4 Responding address parameter

The responding address parameter is present in this primitive only in the case where the primitive is used to indicate rejection of an NC establishment attempt by an NS user (see 13.4). The parameter conveys the address of the NSAP from which the N-DISCONNECT request was issued and, where explicitly supplied, the addresses in the corresponding request and indication primitives are identical. Under certain circumstances (e.g. call redirection, generic addressing, etc.) this address may be different from the “called address” in the corresponding N-CONNECT request primitive.

13.3 Sequence of primitives when releasing an established NC

The sequence of primitives depends on the origin or origins of the NC release action. The sequence may be:

a) invoked by one NS user, with a request from that NS user leading to an indication to the other NS user;

b) invoked by both NS users, with a request from each of the NS users;

c) invoked by the NS provider, with an indication to each of the NS users;

d) invoked independently by one NS user and the NS provider, with a request from the originating NS user and an indication to the other NS user.

The sequences of primitives in these four cases are expressed in the time sequence diagrams in Figures 6 to 9.

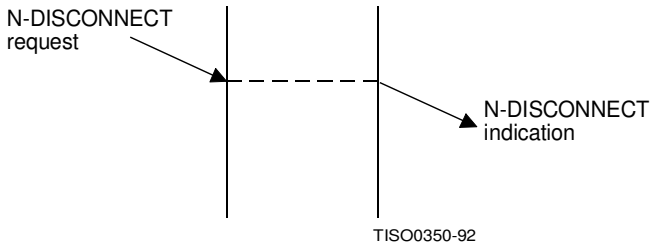

Figure 6 – Sequence of primitives in NS user invoked NC release

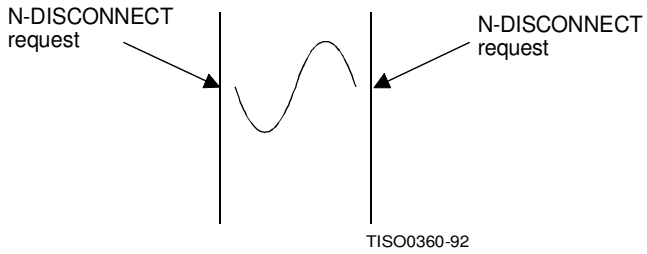

Figure 7 – Sequence of primitives in simultaneous NS user invoked NC release

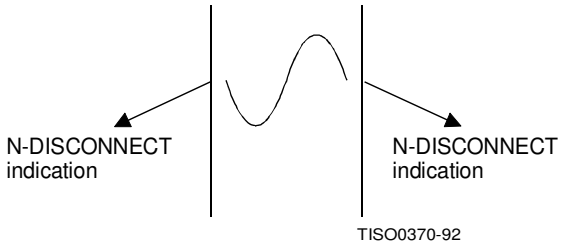

Figure 8 – Sequence of primitives in NS provider invoked NC release

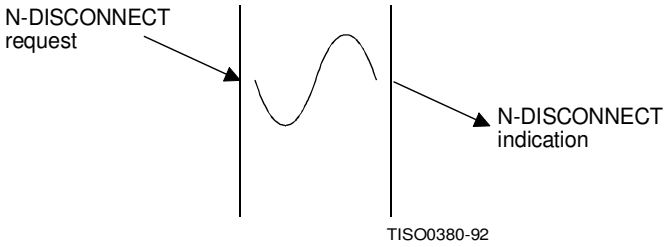

Figure 9 – Sequence of primitives in simultaneous NS user and NS provider invoked NC release

13.4 Sequence of primitives in an NS user rejection of an NC establishment attempt

An NS user may reject an NC establishment attempt by issuing an N-DISCONNECT request. The originator parameter in the N-DISCONNECT primitives will indicate NS user invoked NC release. The sequence of events is defined in the time sequence diagram in Figure 10.

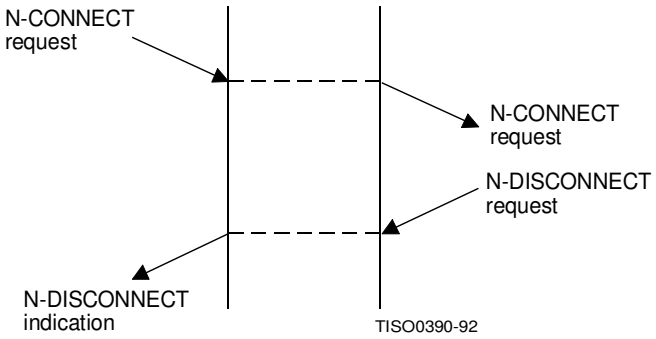

Figure 10 – Sequence of primitives in NS user rejection of an NC establishment attempt

13.5 Sequence of primitives in an NS provider rejection of an NC establishment attempt

If the NS provider is unable to establish an NC, it indicates this to the requester by issuing an N-DISCONNECT indication. The originator parameter in this primitive indicates an NS provider invoked NC release. The sequence of events is defined in the time sequence diagram in Figure 11.

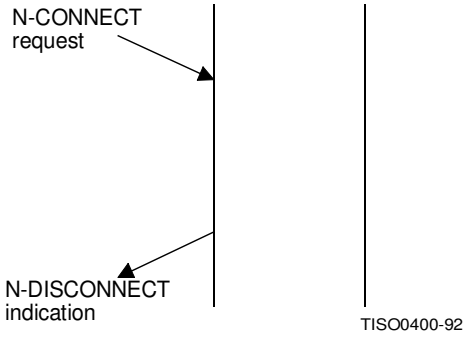

Figure 11 – Sequence of primitives in NS provider rejection of an NC establishment attempt

14 Data transfer phase

14.1 Data transfer

14.1.1 Function

The data transfer service primitives provide for an exchange of NS-user-data called Network-Service-data-units (NSDUs), in either direction or in both directions simultaneously on an NC. The Network Service preserves both the sequence and the boundaries of the NSDUs.

NOTE – Designers of higher layer protocols using the Network Service should realize that the requested QOS applies to complete NSDUs, and that divisions of available NS-user-data into small NSDUs may have cost implications because of the impact on cost optimization mechanisms operated by the NS provider.

14.1.2 Types of primitives and parameters

Table 14 indicates the types of primitives and the parameters needed for data transfer.

Table 14 – Data transfer primitives and parameters

Primitive / Parameter	N-DATA request	N-DATA indication
NS-user-data	X	X(=)
Confirmation request	X(C)	X(C=)

14.1.2.1 NS-user-data parameter

The NS-user-data parameter allows the transfer of an NSDU between NS users, without modification by the NS provider. The NS user may send any integer number of octets, one or greater, of NS-user-data that forms the NSDU.

14.1.2.2 Confirmation request parameter

The receipt confirmation of an NSDU transferred by means of an N-DATA primitive can be requested by setting the confirmation request parameter on the N-DATA request. The confirmation of receipt (COR) is provided by the N-DATA-ACKNOWLEDGE primitives (see 14.2). The value of the confirmation request parameter may indicate either that a COR is requested or that it is not requested. The parameter may only be present if use of the receipt confirmation service was agreed by both NS users and the NS provider during the establishment of the NC.

14.1.3 Sequence of primitives

The operation of the Network Service in transferring NSDUs can be modeled as a queue of unknown size within the NS provider (see clause 9). The ability of an NS user to issue an N-DATA request or of the NS provider to issue an N-DATA indication depends on the behaviour of the receiving NS user and the resulting state of the queue.

The sequence of primitives in a successful data transfer is defined in the time sequence diagram in Figure 12.

The sequence of primitives in Figure 12 may remain uncompleted if an N-RESET or an N-DISCONNECT primitive occurs.

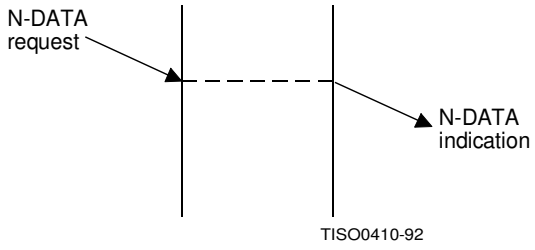

Figure 12 – Sequence of primitives in data transfer

14.2 Receipt confirmation service

14.2.1 Function

The receipt confirmation service is requested by the confirmation request parameter on the N-DATA primitives. For each and every NSDU transferred with the confirmation request parameter set, the receiving NS user should return a confirmation of receipt (COR) by issuing an N-DATA-ACKNOWLEDGE request. Such CORs should be issued in the same sequence as the corresponding N-DATA indications were received, and will be conveyed by the NS provider so as to preserve them distinct from any previous or subsequent CORs. The NS user may thus correlate them with the original N-DATA primitives (with "confirmation requests" set) by counting.

N-DATA-ACKNOWLEDGE requests will not be subject to the flow control affecting N-DATA requests at the same NC endpoint; N-DATA-ACKNOWLEDGE indications will not be subject to the flow control affecting N-DATA indications at the same NC endpoint.

The use of the receipt confirmation service shall be agreed by the two NS users of the NC and the NS provider during NC establishment by use of the receipt confirmation selection parameter on the N-CONNECT primitives. The service need not be provided by all NS providers.

14.2.2 Types of primitives and parameters

The receipt confirmation service involves two primitives:

- N-DATA-ACKNOWLEDGE request;
- N-DATA-ACKNOWLEDGE indication.

These primitives do not convey any parameters.

14.2.3 Sequence of primitives

The sequence of primitives in a successful data transfer with receipt confirmation is defined in the time sequence diagram in Figure 13.

The sequence of primitives in Figure 13 may remain uncompleted if an N-RESET or an N-DISCONNECT primitive occurs.

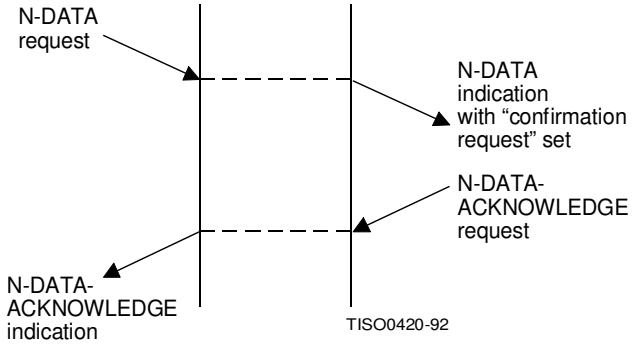

Figure 13 – Sequence of primitives in successful data transfer with receipt confirmation

An NS user must not issue an N-DATA-ACKNOWLEDGE request if no N-DATA indication with "confirmation request" set has been received or if a COR has already been issued for all such N-DATA indications. Following a reset procedure, signaled by means of an N-RESET indication or N-RESET confirm, an NS user must not issue an N-DATA-ACKNOWLEDGE request in response to an N-DATA indication (with "confirmation request" set) which was received before the reset procedure was signaled.

NOTES

1 The withholding of COR by an NS user can have an effect on the throughput attainable on the NC.

2 The use of receipt confirmation on an NC may have an effect on the flow control of normal data on the NC. For example, the issuing of a COR may result in the relaxation of flow control on NS-user-data flowing in the opposite direction from the COR.

3 Receipt confirmation is included in the Network Service only to support existing features of Recommendation X.25.

14.3 Expedited data transfer service

14.3.1 Function

The expedited data transfer service provides a further means of information exchange on an NC in both directions simultaneously. The transfer of expedited Network-Service-data-units (ENSDUs) is subject to different QOS and separate flow control from that applying to NS-user-data of the data transfer service. It is not intended to provide a qualified data transfer facility.

The NS preserves both the sequence and boundaries of the ENSDUs. The NS provider guarantees that an ENSDU will not be delivered after any subsequently issued NSDU or ENSDU on that NC.

The relationship between normal and expedited NS-user-data is modeled by the operation of changing of order within queues as described in 9.2.3. In particular, expedited NS-user-data can still be delivered when the receiving NS user is not accepting normal NS-user-data. However, the amount of normal NS-user-data bypassed by such changing of order cannot be predicted or guaranteed. Expedited data transfer cannot be guaranteed to bypass blockages in normal data flow where these blockages are occurring in lower layers.

The expedited data transfer service is a provider-option which may not be available in the Network Service. Its use must be agreed to by the two NS users of the NC and the NS provider during NC establishment by means of the expedited data selection parameter on the N-CONNECT primitives (see 12.2.6).

14.3.2 Types of primitives and parameters

Table 15 indicates the types of primitives and the parameters needed for expedited data transfer.

Table 15 – Expedited data transfer primitives and parameters

Primitive / Parameter	N-EXPEDITED-DATA request	N-EXPEDITED-DATA indication
NS-user-data	X	X(=)

14.3.2.1 NS-user-data parameter

The NS-user-data parameter allows the transfer of expedited NS-user-data between NS users, without modification by the NS provider. The NS user may send any integer number of octets of expedited NS-user-data between 1 and 32 inclusive.

14.3.3 Sequence of primitives

The sequence of primitives in a successful expedited data transfer is defined in the time sequence diagram in Figure 14.

The sequence of primitives in Figure 14 may remain uncompleted if an N-RESET or an N-DISCONNECT primitive occurs.

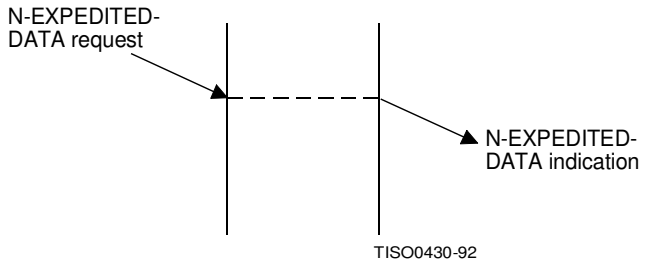

Figure 14 – Sequence of primitives in expedited data transfer

14.4 Reset service

14.4.1 Function

The reset service may be used:

a) by the NS user to resynchronize the use of the NC; or

b) by the NS provider to report detected loss of NS-user-data unrecoverable within the NS provider. All loss of NS-user-data which does not involve loss of the NC is reported in this way.

Invocation of the reset service will unblock the flow of NSDUs and ENSDUs in case of congestion of the NC; it will cause the NS provider to discard NSDUs, ENSDUs, or CORs associated with the NC, and to notify any NS user or users that did not invoke reset that a reset has occurred. The service will be completed in a finite time, irrespective of the acceptance of NSDUs, ENSDUs, and CORs by the NS users. Any NSDUs, ENSDUs, or CORs not delivered to the NS users before completion of the service will be discarded by the NS provider.

14.4.2 Types of primitives and parameters

Table 16 indicates the types of primitives and the parameters needed for the reset service.

Table 16 – Reset primitives and parameters

Primitive / Parameter	N-RESET request	N-RESET indication	N-RESET response	N-RESET confirm
Originator		X		
Reason	X	X		

14.4.2.1 Originator parameter

The originator parameter indicates the source of the reset. Its value indicates either the “NS user”, “NS provider”, or “undefined”.

14.4.2.2 Reason parameter

The reason parameter gives information indicating the cause of the reset. The value conveyed in this parameter will be as follows:

a) When the originator parameter indicates an NS provider invoked reset, the value is one of:

i) “congestion”;

ii) “reason unspecified”.

b) When the originator parameter indicates an NS user invoked reset, the value is “user resynchronization”.

c) When the originator parameter has the value “undefined”, then the value of the reason parameter is also “undefined”.

14.4.3 Sequence of primitives

The interactions between each NS user and the NS provider will be an exchange of these primitives, namely either:

a) an N-RESET request from the NS user, followed by an N-RESET confirm from the NS provider; or

b) an N-RESET indication from the NS provider, followed by an N-RESET response from the NS user.

The N-RESET request acts as a synchronization mark in the stream of NSDUs, ENSDUs, and CORs transmitted by the issuing NS user. The N-RESET indication likewise acts as a synchronization mark in the stream of NSDUs, ENSDUs, and CORs received by the receiving NS user. Similarly, the N-RESET response acts as a synchronizing mark in the stream of NSDUs, ENSDUs, and CORs transmitted by the responding NS user, while the N-RESET confirm acts as a synchronization mark in the stream of NSDUs, ENSDUs, and CORs received by the NS user which originally invoked the Reset.

The resynchronizing properties of the reset service are that

1) No NSDU, ENSDU, or COR transmitted by the NS user *before* the synchronization mark in that transmitted stream will be delivered to the other NS user *after* the synchronization mark in that received stream.

   The NS provider will discard all NSDUs, ENSDUs, and CORs submitted before the issuing of the N-RESET request which have not been delivered to the receiving NS user when the NS provider issues the N-RESET indication.

   Also, the NS provider will discard all NSDUs, ENSDUs, and CORs submitted before the issuing of the N-RESET response which have not been delivered to the initiator of the N-RESET when the NS provider issues the N-RESET confirm.

2) No NSDU, ENSDU, or COR transmitted by an NS user *after* the synchronization mark in that transmitted stream will be delivered to the other NS user *before* the synchronization mark in that received stream.

The N-RESET confirm may be issued to the initiator of the reset before the N-RESET indication is issued to the other NS user. The complete sequence of primitive depends upon the origin of the reset action and the occurrence or otherwise of resets with conflicting origins. Thus the reset service may be

i) invoked by one NS user, lading to interaction a) with that NS user and interaction b) with the peer NS user;

ii) invoked by both NS users, leading to interaction a) with both NS users;

iii) invoked by the NS provider, leading to interaction b) with both NS users;

iv) invoked by one NS user and the NS provider, leading to interaction a) with the originating NS user and b) with the peer NS user.

The sequence of primitives in these four cases is defined in the time sequence diagrams in Figures 15 to 18.

Further, there may be circumstances of reset “collision” which result in the number of reset procedures observed at one NC endpoint being different from the number of reset procedures observed at the other NC endpoint. Such circumstances result in two additional cases which may occur, where the reset service may be

v) invoked by one NS user while a previous reset procedure is still incomplete at the other NS user, leading to additional interaction a) with the NS user invoking the subsequent reset, only;

vi) invoked by the NS provider at one NC endpoint, while a previous reset procedure is still incomplete at the other, leading to additional interaction b) with the NS user at the first NC endpoint, only.

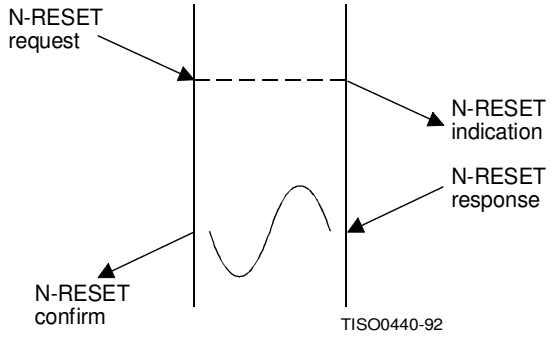

Figure 15 – Sequence of primitives in NS user invoked reset

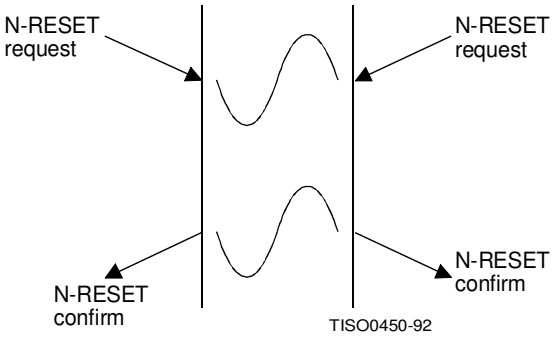

Figure 16 – Sequence of primitives in simultaneous NS user invoked reset

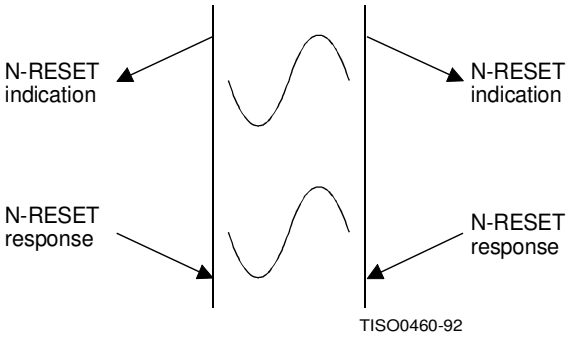

Figure 17 – Sequence of primitives in NS provider invoked reset

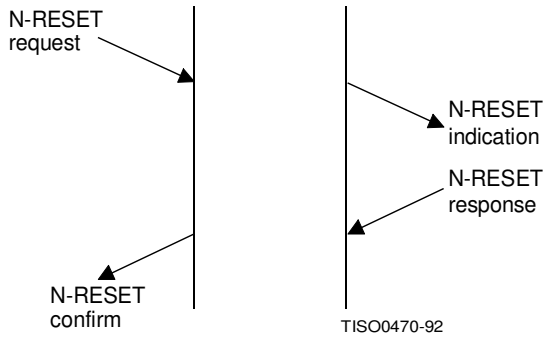

Figure 18 – Sequence of primitives in simultaneous NS user and NS provider invoked reset

There are many possible sequences of reset primitives for the two NC endpoints which may occur for cases v) and vi). These are not illustrated here by time sequence diagrams, but may be derived using the constraints on the allowed sequence of primitives for each NC endpoint, and the reset sequences illustrated in Figures 15 to 18. The synchronizing properties associated with the issuance of the N-RESET primitives are the same for all of the six cases outlined.

NOTE – Situations in which the number of reset procedures at the two ends of a NC which are not the same are not described by the operation of the queue model in 9.2.

Any sequence of reset primitives may remain uncompleted if an N-DISCONNECT primitive occurs. Once a reset procedure has been invoked at an NC endpoint (by means of an N-RESET request or N-RESET indication primitive), no further N-DATA, N-EXPEDITED-DATA, or N-DATA-ACKNOWLEDGE primitive can be issued by either the NS user or the NS provider until the reset procedure has completed (by means of an N-RESET confirm or N-RESET response).

SECTION 3 — DEFINITION OF THE CONNECTIONLESS-MODE SERVICE

15 Features of the connectionless-mode Network Service

The NS provides the following features to the NS user:

a) The means by which network-service-data-units are delimited and transparently transmitted from a source NSAP to a destination NSAP in a single network-connectionless-mode service access, without first establishing or later releasing a network-connection. The maximum size of a connectionless-mode NSDU is 64#512 octets.

b) Associated with each instance of network-connectionless-mode transmission, certain measures of quality which are agreed between the NS provider and the sending NS user when a network-connectionless-mode transmission is initiated.

16 Model of the connectionless-mode Network Service

16.1 Model of the connectionless-mode Network Layer Service

This Service Definition uses the abstract model for a layer service defined in clause 4 of CCITT Rec. X.210 | ISO/TR 8509. The model defines the interactions between the NS users and the NS provider which take place at the two NSAPs. Information is passed between the NS user and the NS provider by service primitives, which may convey parameters.

16.2 Model of a network connectionless-mode transmission

A defining characteristic of network-connectionless-mode transmission is the independent nature of each invocation of the connectionless-mode network service.

In practice, however, it is often possible to relate to NS users certain characteristics of the service for an association existing between a given pair of NSAPs. Such characteristics, when present, enhance the basic connectionless-mode Network Service, and can be used by NS users to effectively correlate the choice of transport protocol with the Network Service provided.

NOTE 1 — Such information is typically made available to the NS user through some control or management facility (or set of facilities). Annex C provides a description of the facilities considered to be necessary for the coordination of transport protocol selection with use of the connectionless-mode Network Service.

As a descriptive aid, the connectionless-mode Network Service, as provided between any two NSAPs, can be modelled in the abstract as an *a priori* association between the two NSAPs.

NOTE 2 — This model is intended solely to describe the appearance of the connectionless-mode Network Service to the NS users. It is not intended to be a model of the internal operation of the NS provider in providing the connectionless-mode Network Service.

Only one type of object, the Unitdata object, can be exchanged between the NS users. In Figure 19, User X represents the NS user that passes objects to the NS provider. User Y represents the NS user that accepts objects from the NS provider.

In general, the NS provider may perform any or all of the following actions:

a) discard objects;

b) duplicate objects; and

c) change the order of the objects.

However, with respect to a given association between a pair of NSAPs, the following additional characteristics might be observed by the NS user through the operation of some control or management facility that allows the NS provider to support service characteristics beyond those attributed to the basic connectionless-mode Network Service:

1) objects will be discarded only after a stated time;

2) objects must be discarded no later than a stated time;

3) objects will be discarded only if more than a certain number of objects are in the queue;

4) objects will not be discarded;

5) the order of the objects in the queue will not be changed;

6) objects will not be duplicated.

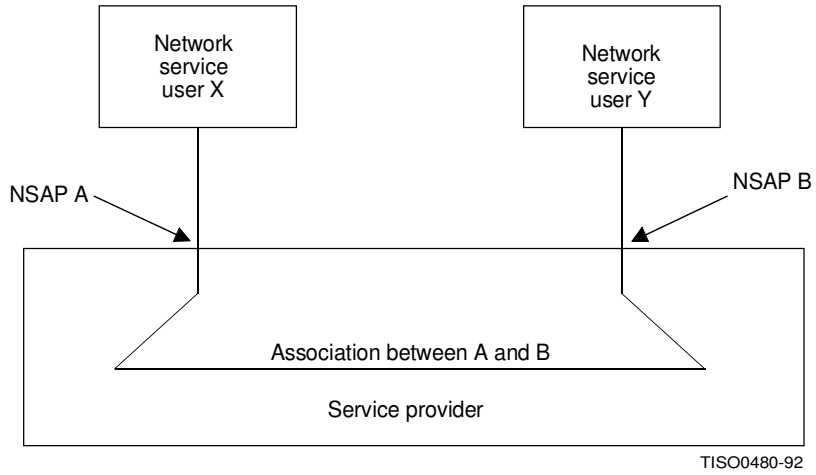

Figure 19 – Model of a network-connectionless-mode-transmission

Where such information is made known to the NS user prior to the invocation of the connectionless-mode Network Service, the NS user may make use of such knowledge to select a transport protocol.

The existence and the properties of the association, and the operations which are performed by the NS provider for a particular association, do not depend on the behaviour of the NS users. Awareness of the characteristics of the association is part of the NS users' *a priori* knowledge of the OSI environment.

17 Quality of the connectionless-mode Network Service

The term "Quality of Service" (QOS) also refers to certain characteristics of a network-connectionless-mode transmission as observed between a pair of NSAPs. QOS describes aspects of a network-connectionless-mode transmission which are attributable solely to the NS provider; it can only be properly determined in the absence of NS user behaviour (which is beyond the control of the NS provider) which specifically constrains or impedes the performance of the Network Service.

Whether the view of the QOS during each instance of the use of the network-connectionless-mode transmission is the same to each NS user associated with the service depends on the nature of their association and the type of information concerning the nature of the service made available to the NS user(s) by the NS provider prior to the invocation of the service.

17.1 Determination of QOS

A basic characteristic of a connectionless-mode service is that no peer-to-peer negotiation of the Quality of Service for a transmission takes place at the time the service is accessed. No dynamic association is set up between the parties involved as during a connection establishment; thus, characteristics of the service to be provided during the transfer are not negotiated on a peer-to-peer basis. An *a priori* agreement is assumed to exist between the NS users and the NS provider concerning those parameters, formats, and options that affect the transfer of data. (Such an *a priori* agreement might be established by local exchange of relevant information across the Network Layer boundary.) Thus, the process is one of local negotiation.

Associated with each network-connectionless-mode transmission, certain measures of Quality of Service are agreed between the NS provider and the sending NS user when the primitive action is initiated. The requested measures (or parameter values and options) are based on *a priori* knowledge by the NS user of the specific characteristics of the facilities which can be expected to be made available to it by the NS provider.

Knowledge of the characteristics of the connectionless-mode Network Service is made available to a sending NS user through a control or management facility prior to invocation of the connectionless-mode Network Service. The NSáuser not only has knowledge of the parties with which it may communicate, but also has explicit knowledge of the characteristics of the service with which it can expect to be provided upon (each) invocation of the service.

The NS provider may also provide information on the current Quality of Service, independent of access to the service by an NS user. This seemingly dynamic aspect of Quality of Service determination is not a negotiation but, again, a means by which the NS user is provided with knowledge of the characteristics of the service currently available outside of any instance of the invocation of the service itself.

The NS QOS-parameters associated with each network-connectionless-mode-transmission are defined in 17.2.

17.2 Definition of network connectionless-mode QOS-parameters

The QOS-parameters identified for the network-connectionless-mode transmission are defined below.

NOTE — Additional parameters that describe service characteristics beyond those defined in this subclause for the connectionless-mode Network Service are described in Annex C.

17.2.1 Transit delay

Transit delay is the elapsed time between an N-UNITDATA request and the corresponding N-UNITDATA indication. Elapsed time values are calculated only on NSDUs that are successfully transferred. Successful transmission of an NSDU is defined to occur when an NSDU transmitted from a sending NS user is delivered to the intended receiving NS user.

Transit delay is specified independently for each network-connectionless-mode transmission. Transit delay defines the value expected for the completion of the transmission of a particular NSDU. Its specification is based on an average NSDU size. It is determined by the NS provider and made known to the NS user prior to invocation of the service.

Transit delay for an individual NSDU may be greatly increased if local interface flow control is exercised at either the transmitting or receiving service provider to service user interface. Occurrences of local interface flow control initiated by the service user are excluded in calculating transit delay values.

17.2.2 Protection from unauthorized access

The extent to which an NS provider attempts to prevent unauthorized monitoring or manipulation of NS user-originated information is specified qualitatively by selecting one of four options:

a) no protection features;

b) protection against passive monitoring;

c) protection against modification, replay, addition, or deletion; and

d) both b) and c).

17.2.3 Cost determinants

A class of parameter values and options exist which provide a NS user with

a) the ability to indicate to the NS provider that it should choose, for example, the least expensive means available to it, even in situations where this may not be the most expedient means; or

b) the ability to specify maximum acceptable cost.

The cost may be specified in absolute or relative cost units. The cost of a network-connectionless-mode transmission is composed of communications and end-system resource costs.

17.2.4 Residual Error Probability

Residual Error Probability describes the likelihood that a particular NSDU will be lost, duplicated, or delivered incorrectly. The probability is estimated as the ratio of lost, duplicated, or incorrectly delivered NSDUs to total NSDUs transmitted by an NS provider during a measurement period.

An incorrectly delivered NSDU is one in which the user data are delivered in a corrupted condition, or the user data are delivered to an incorrect NSAP.

Lost data includes all NSDUs which are discarded by the NS provider due to congestion, transmission error, or some other error. NSDUs which are lost due to error by the NS user are not included.

17.2.5 Priority

This parameter allows the NS user to specify the relative priority of an NSDU in relation to any other NSDUs acted upon by the NS provider. An NSDU of higher priority is serviced by the NS provider before one of lower priority. The priority information is conveyed to the receiving NS user.

This parameter specifies the relative importance of network-connectionless-mode transmission with respect to

a) the order in which NSDUs are to have their Quality of Service degraded, if necessary; and

b) the order in which NSDUs are to be discarded to recover resources, if necessary.

This parameter has meaning only in the context of some management entity or structure able to judge relative importance. The number of priority levels is limited to 15.

17.3 Route selection considerations

The NS provider may use QOS information supplied by the NS user, in conjunction with management information, to derive details of the route that must be taken by the NSDU with which the information is associated.

The use of such information by the NS provider may result in the network-entity in the originating system being able to enumerate a complete list of the intermediate systems to be included in the path. This is referred to as “complete source routing”. Complete source routing requires that each intermediate system in the list (and only those intermediate systems in the list) must be visited; if this constraint cannot be satisfied, the NSDU shall be discarded.

Alternatively the network-entity in the originating system may be able to identify a list of some (but not all) of the intermediate systems to be included in the path. This is referred to as “partial source routeing”. With partial source routing, each intermediate system in the list must be visited in the order specified while en route to the destination. However, with this form of source routing the NSDU may take any path necessary to arrive at the next intermediate system in the list. The NSDU shall not be discarded (for source routing related causes) unless one of the intermediate systems in the list cannot be reached by any available path.

Alternatively, no source routing may take place at all. This is the case when the required QOS can be achieved without any need for the originating network-entity to specify particular intermediate systems which must be included in the path.

18 Sequence of primitives

The possible overall allowed sequences of primitives at an NSAP are defined in the state transition diagram in Figure 20. In Figure 20, the IDLE state represents the initial and final state of the primitive sequence.

The use of a state transition diagram to describe the allowable sequence of service primitives does not impose any requirements or constraints on the internal organization of any implementation of the Network Service.

19 Data transfer

19.1 Function

Network-connectionless-mode transmission service primitives can be used to transmit an independent, self-contained NSDU from one network-service-access-point to another network-service-access-point in a single service access. The NSDU is independent in the sense that the NS provider is not required to maintain any relationship between this and

any other NSDU transmitted through the invocation of the connectionless-mode Network Service or the connection-mode Network Service. It is self-contained in that all of the information required to deliver the NSDU is presented to the NS provider, together with the user data to be transmitted, in a single service access; thus no initial establishment or subsequent release of a network-connection is required.

The NS provider transfers individual NSDUs within the range of its specified QOS. Although the Network Service maintains the integrity of individual NSDUs, it does not necessarily deliver them to the receiving NS user in the order in which they are presented by the sending NS user.

The NS provider need not maintain any state information relative to any aspect of the flow of information between any specific pair of NSAPs. Flow control exerted by the NS provider upon the sending NS user can be described only in terms of interface flow control.

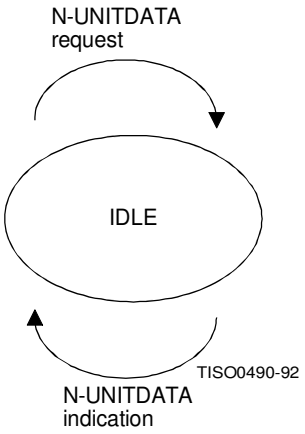

Figure 20 – State transition diagram for sequences of connectionless-mode primitives at one NSAP

19.2 Types of primitives and parameters

Table 17 identifies the types of primitives and the parameters needed for the connectionless-mode Network Service.

Table 17 — Connectionless-mode service primitives and parameters

Primitive / Parameter	N-UNITDATA request	N-UNITDATA indication
Source address	X	X(=)
Destination address	X	X(=)
Quality of Service	X	X
NS-USERDATA•	X	X(=)

19.2.1 Addresses

The addresses referred to in Table 17 are network-service-access-point addresses. Both the connection-mode and connectionless-mode Network Services use the same NSAP addresses, as described in Annex A.

19.2.2 Quality of Service

The value of the Quality of Service parameter is a list of sub-parameters. For each parameter, the values of the parameter in the two primitives are related so that:

a) in the request primitive, any defined value is allowed; however,

b) in the indication primitive, the Quality of Service indicated may be different from the QOS specified for the corresponding request primitive.

19.2.3 NS-User-data

This parameter allows the transmission of octets of NS user-supplied data between NS users, without modification by the NS provider. The NS user may transmit any integral number of octets greater than zero up to 64#512 octets.

19.3 Sequence of primitives

The sequence of primitives in a network-connectionless-mode transmission is defined in the Network Service primitive time sequence diagram, see Figure 21.

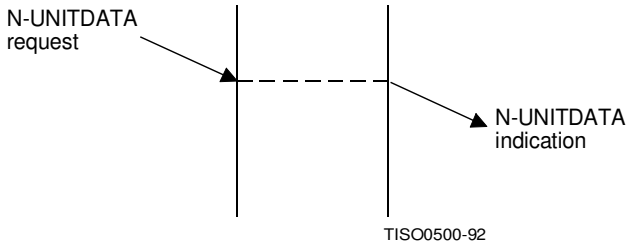

Figure 21 – Sequence of primitives in a network-connectionless-mode transmission

Annex A
Network Layer Addressing

(This annex forms an integral part of this Recommendation | International Standard)

A.1 General

This annex defines the abstract syntax and semantics of the Network address (Network service access point address). The Network address is the address that appears in the primitives of the connection-mode Network service as the *calling address*, *called address* and *responding address* parameters, and in the primitives of the connectionless-mode Network service as the *source address* and *destination address* parameters.

A.2 Scope

The scope of this annex is the definition of the abstract syntax and semantics of the Network address. This annex does not specify the way in which the semantics of the Network address are encoded in Network layer protocols.

A.3 Concepts and terminology

A.3.1 Network addresses

This annex defines the Network service access point (NSAP) address. Since the term "network address" is commonly used in different contexts to refer to different things, a more specific description of this concept is introduced below.

A.3.1.1 Subnetwork address

In one context, the term "network address" may be used to refer to the point at which a *real end system*, *real subnetwork*, or *interworking unit* is attached to a real subnetwork, or to the point at which the subnetwork service is offered within an end or intermediate system. In the case of attachment to a public data network, this point is called a DTE/DCE interface, and the term "DTE address" is used in reference to it.

The specific term *subnetwork address* (or *subnetwork point of attachment address*) is used in this case, as illustrated in Figure A.1.

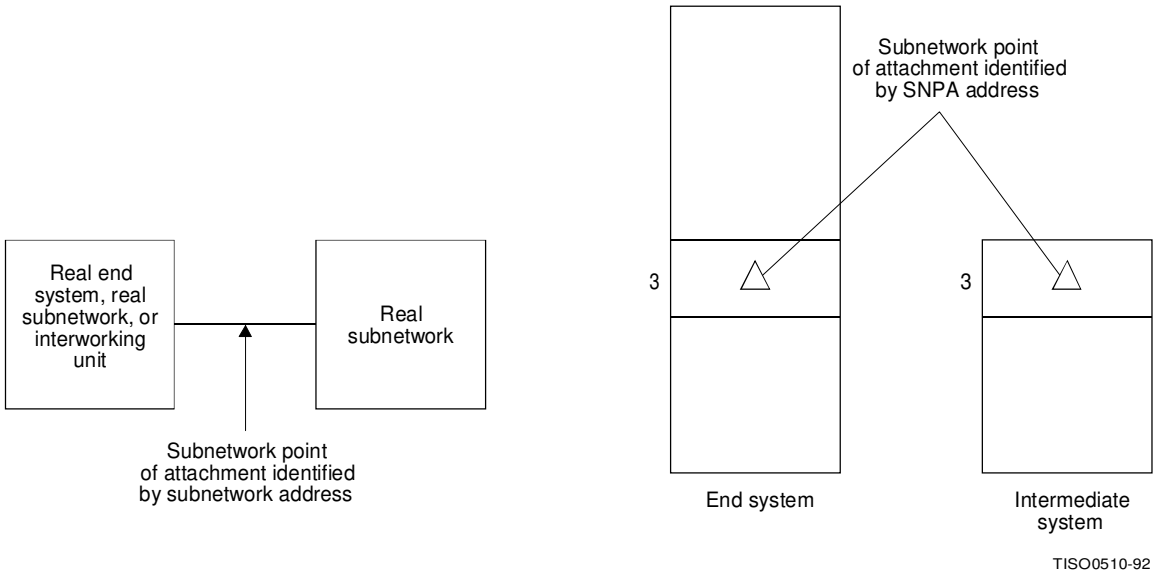

Figure A.1 – Subnetwork address

The subnetwork address is the information that a real subnetwork needs to identify a particular real end system, another real subnetwork, or an interworking unit, which is attached to that real subnetwork.

In the public network environment, the subnetwork address is what the public network operates on.

NOTE — The point identified by a subnetwork address is a point of interconnection between a real end system or interworking unit and a real subnetwork (in particular, in a public data network environment, a DTE/DCE interface), and is not a Network service access point.

A.3.1.2 NSAP address

In another context, the term "network address" is used to refer to the *Network service access point* (NSAP) at which the OSI Network service is made available to a Network service user by the Network service provider.

The specific term *NSAP address* is used in this case, as illustrated in Figure A.2

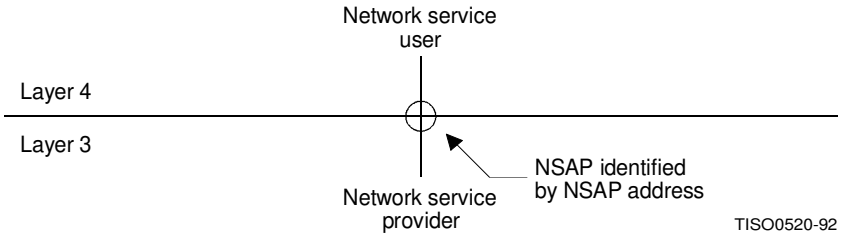

Figure A.2 – NSAP address

The NSAP address is the information that the OSI Network service provider needs to identify a particular Network service access point. The values of the *called address*, *calling address*, and *responding address* parameters of the N-CONNECT primitive, of the *responding address* parameter of the N-DISCONNECT primitive, and of the *source address* and *destination address* parameters of the N-UNITDATA primitive, are NSAP addresses.

It should be noted that since the Network service primitives are conceptual, no particular encoding of the NSAP address is specified by the Network service definition.

In both CCITT and ISO/IEC usage, the terms "Network Address" (with both the N and the A printed in capital letters) and "global network address" are synonymous with the term "NSAP address". Use of the term "NSAP address" is preferred when it is essential to avoid confusion, particularly in spoken references in which "capitalization" is not possible.

A.3.1.3 Network protocol address information

In a third context, the term "network address" is used to refer to an address that is carried as *Network protocol control information* in a Network protocol data unit (NPDU).

The specific term *Network protocol address information* (NPAI) is used in this case.

In the public network environment, NPAI is also known as an "address signal" or as the "coding of an address signal".

There is a relationship between the NSAP address that appears in Network service primitives and the NPAI that appears in a Network layer protocol, in that the semantics of the NSAP address are preserved by the NPAI. The precise encoding of NPAI is defined by Network layer protocol standards, which also specify the relationship between the NSAP address and the NPAI encoding employed by the protocol.

A.3.2 Domains

A.3.2.1 Global network addressing domain

The *global network addressing domain* is an addressing domain consisting of all of the NSAP addresses in the OSI environment.

A.3.2.2 Network addressing domain

A *network addressing domain* is a subset of the global network addressing domain consisting of all of the NSAP addresses that are assigned by one or more addressing authorities. Every NSAP address is part of a network addressing domain that is administered directly by one and only addressing authority. If that network addressing domain is part of a (hierarchically) higher addressing domain (which must wholly contain it), the authority for the (hierarchically) lower domain is authorized by the authority for the higher domain to assign NSAP addresses from the lower domain. All network addressing domains are in this way ultimately part of the global network addressing domain, for which the addressing authority is this annex.

The relationship of the concepts of A.3.2.1 and A.3.2.2 is illustrated by Figure A.3:

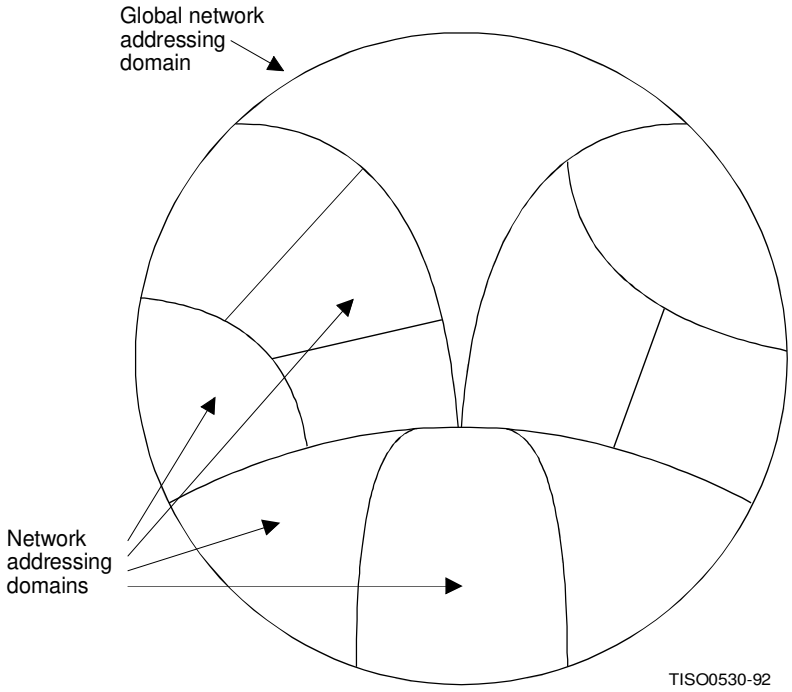

Figure A.3 – Network addressing domains

A.3.3 Authorities

The uniqueness of identifiers within a network addressing domain is ensured by an *authority* associated with that domain. The term “authority” does not necessarily refer to an organization or administration; it is intended to refer to whatever it is (in an abstract sense) that ensures the uniqueness of identifiers in the associated domain.

A network addressing domain is characterized by the addressing authority that administers the domain and by the rules that are established by that authority for specifying identifiers and identifying subdomains. The authority responsible for each domain determines how identifiers will be assigned and interpreted within that domain, and how any further subdomains will be created.

The operation of an authority is independent of that of other authorities at the same level of the hierarchy, subject only to any common rules imposed by the parent authority.

A.3.4 Network address allocation

An addressing authority shall either allocate complete NSAP addresses, or authorize one or more other authorities to allocate addresses. Each address allocated by an addressing authority shall include a domain identifier which identifies the allocating authority. An address shall not be allocated to identify a domain or NSAP if the address had previously been allocated in some other domain or NSAP, unless the authority can ensure that all use of the previous allocation has ceased.

The authority shall ensure that allocations are made in such a way that efficient use is made of the address space.

A.4 Principles for creating the OSI Network addressing scheme

A.4.1 Hierarchical structure

NSAP addresses are based on the concept of hierarchical addressing domains, as explained in A.3. Each domain may be further partitioned into subdomains. Accordingly, NSAP addresses have a hierarchical structure.

The conceptual structure of NSAP addresses follows the principle that, at any level of the hierarchy, an initial part of the address unambiguously identifies as subdomain, and the rest is allocated by the authority associated with the subdomain to unambiguously identify either a lower level subdomain or an NSAP within the subdomain. The part of the address that identifies the subdomain depends on the level at which the address is viewed.

NOTE — This conceptual structure should not be considered as implying any detailed administration of NSAP addresses.

Graphical representation of the hierarchical structure of NSAP addresses may be made according to an inverted tree diagram, as in Figure A.4a), or a domain diagram, as in Figure A.4b):

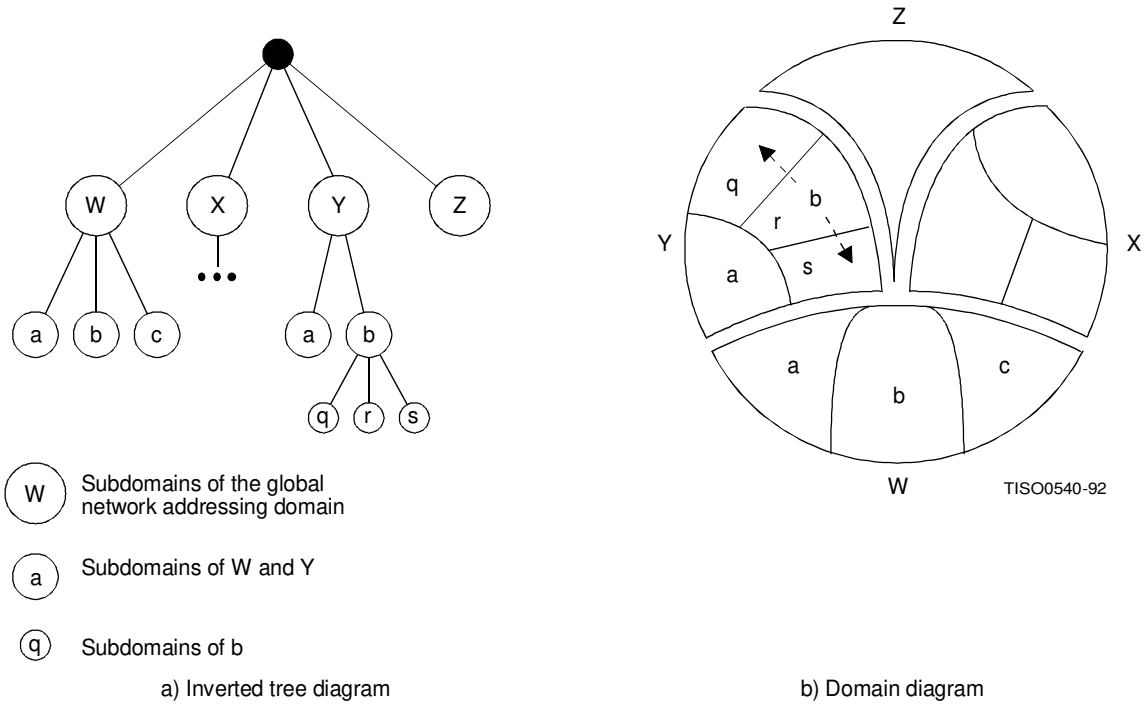

Figure A.4 – Hierarchical structure of NSAP address

A.4.2 Global identification of any NSAP

In the context of Open Systems Interconnection, it is possible to identify any NSAP within the global network addressing domain (see A.3.2.1). Consequently,

a) an NSAP address can be defined to unambiguously identify any NSAP;

b) at any NSAP, it is possible to identify any other NSAP, within any OSI end system;

c) the Network layer protocols established between correspondent Network entities convey the complete semantics of an NSAP address;

d) an NSAP address always identifies the same NSAP, regardless of which Network service user enunciates the address; and

e) a Network service user, when given an NSAP address by the Network service provider in an Indication service primitive, may subsequently use that NSAP address in another instance of communication with the corresponding NSAP.

NOTE — The global identification of NSAPs does not imply the universal availability of directory functions required to enable communication among all NSAPs to which NSAP addresses have been assigned, nor does it prevent the imposition of external restrictions on communication based on the technical feasibility of interconnection, security, charging, etc.

A.4.3 Route independence

Network service users cannot derive routing information from an NSAP address. They cannot influence the Network service provider's choice of route by means of the source and destination NSAP addresses. Similarly, they cannot determine the route that was used by the Network service provider by examining the source and destination NSAP addresses. This is not intended to exclude the possibility that an OSI end system may need to influence the route selected for a particular instance of communication with another OSI end system. (In particular, it may need to influence the selection of intermediate systems to be used, and the paths to be taken between them.) The means whereby such an influence may be exerted is, however, *not* the NSAP address. Elements of Network layer protocol may be required to control routing within intermediate systems; such elements of protocol are distinct from the network protocol address information (NPAI).

Notwithstanding the restrictions imposed on the use that a Network service *user* may make of an NSAP address, it is recognized that NSAP addresses should be constructed in such a way that routing through interconnected subnetworks is facilitated. That is, space the Network service *provider*, and relay-entities in particular, may take advantage of the address structure to achieve economical processing of routeing aspects.

A.5 Network address definition

The intent of this annex is best served by maintaining clear distinctions among three concepts: the *abstract semantics* of the NSAP address; the *abstract syntax* employed in this annex as a means of defining the abstract semantics of the NSAP address, and employed by network addressing authorities as a means of allocating and assigning NSAP addresses; and the *encoding* of the NSAP address semantics as NPAI in Network layer protocols. These distinctions are illustrated in Figure A.5:

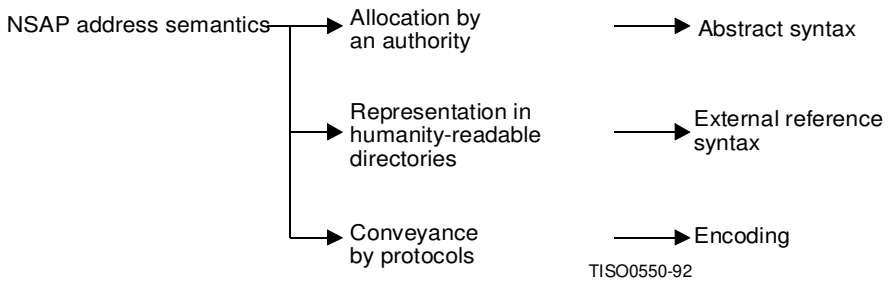

Figure A.5 – Relationship of NSAP address semantics and syntax

This annex does not specify the way in which the semantics of the NSAP address are encoded in Network layer protocols, although a preferred encoding is defined in A.5.3. Network layer protocol specifications define the way in which the NSAP address is encoded as NPAI (see A.3.1.3).

A.5.1 Network address semantics

The NSAP address consists of two basic semantic parts. The first part is the *initial domain part* (IDP). The second part is the *domain specific part* (DSP). This is illustrated by Figure A.6.

Following the conceptual structure of NSAP addresses described in A.4.1, the IDP is a network addressing domain identifier: it specifies a subdomain of the global network addressing domain (see Figure A.4), and identifies the network addressing authority responsible for assigning NSAP addresses in the specified subdomain. The DSP is the corresponding subdomain address. A further substructure of the DSP may or may not be defined by the authority identified by the IDP.

A.5.1.1 The IDP

The initial domain part of the NSAP address itself consists of two parts. The first part is the *authority and format identifier* (AFI). The second part is the *initial domain identifier* (IDI). This is illustrated by Figure A.6:

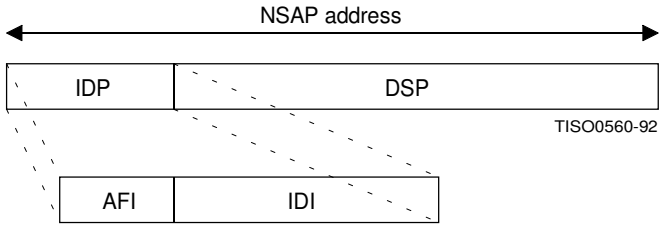

Figure A.6 – NSAP address structure

A.5.1.1.1 The AFI

The authority and format identifier specifies:

a) the format of the IDI (see A.5.2.1.2);

b) the network addressing authority responsible for allocating values of the IDI (see A.5.2.1.2);

c) whether or not leading zero digits in the IDI are significant (see A.5.3); and

d) the abstract syntax of the DSP (see A.5.2.2 and A.5.2.3).

A.5.1.1.2 The IDI

The initial domain identifier specifies:

a) the network addressing domain from which values of the DSP are allocated; and

b) the network addressing authority responsible for allocating values of the DSP from that domain.

A.5.1.2 The DSP

The semantics of the DSP are determined by the network addressing authority identified by the IDI (see A.5.1.1.2).

A.5.2 Network address abstract syntax

The Network address is defined in this annex in terms of an abstract syntax in which the semantics of the Network address can be expressed. The use of this abstract syntax as a descriptive device enables this annex to convey, in written form, a complete definition of the Network address without restricting it to the specific encoding of the NPAI. It also enables this annex to identify a preferred encoding of the Network address, to which reference may be made by Network layer protocol specification standards so as to unambiguously define the way in which the Network address is encoded as NPAI.

A.5.2.1 Abstract syntax and allocation of the IDP

This section defines the abstract syntax of the AFI, the currently allocated values of the AFI, and the IDI formats corresponding to the allocated AFI values. Among the currently allocated values of the AFI are values reserved for assignment to new IDI formats which may be identified by ISO or CCITT. Assignment of these AFI values to new IDI formats by either ISO or CCITT must be accompanied by appropriate modification of this annex. Allocation of new AFI values shall be by joint agreement between ISO and CCITT, and will require an appropriate modification of this annex.

The abstract syntax of the IDP is decimal digits. The allocation of the AFI ensures that the first decimal digit of the IDP can never be zero (see A.5.1.1). This provides an escape mechanism for use by protocols that expect to hold incomplete NSAP addresses in a field that normally carries a complete NSAP address. When the NSAP address is represented as binary octets, the representation of the IDP is as defined in A.5.3. The length of the IDP depends on the IDI format specified by the value of the AFI. The IDP length associated with each IDI format is given in A.5.2.1.2.

A.5.2.1.1 Abstract syntax and allocation of the AFI

The AFI consists of an integer with a value between 0 and 99 with an abstract syntax of two decimal digits. The values of the AFI are allocated or reserved as shown in Table A.1.

Table A.1 — AFI allocations

00-09	Reserved — will not be allocated
10-35	Reserved for future allocation by joint agreement of ISO/IEC and CCITT
36-59	Allocated and assigned to the IDI formats defined in A.5.2.1.2
60-69	Allocated for assignment to new IDI formats by ISO/IEC
70-79	Allocated for assignment to new IDI formats by CCITT
80-99	Reserved for future allocation by joint agreement of ISO and CCITT

A.5.2.1.2 Format and allocation of the IDI

A specific combination of IDI format and DSP abstract syntax is associated with each allocated AFI value, as summarized in Table A.2. Two AFI values are associated with each combination that involves a variable-length IDI format. In each case, both of the AFI values identify the same combination of IDI format and DSP abstract syntax. The numerically lower AFI value is used when the first significant digit in the IDI is non-zero. The numerically greater AFI value is used when the first significant digit in the IDI is zero.

In the following subsections, references are made to the number used in particular subnetwork numbering plans, and to the entity identified by such a number. These references refer to the entity located at the subnetwork point of attachment specified by the number, and not to some other entity (for example, a PTT Administration) whose identity might be inferred from inspection of some part of the number. The authority in such cases is therefore the authority associated with the entity at the SNPA, and is identified by the *full* number.

NOTES

1 The use of a particular IDI format as the basis for allocating an NSAP address does not constrain routing to that NSAP to go through any particular system or subnetwork. For example, the use of the E.163 IDI format as the basis for allocating an NSAP address does not mean that access to the NSAP with that address necessarily involves the user of the public telephone network (see A.4.3).

2 The IDI formats that are based on CCITT numbering plans may be affected by any changes to those plans. It should be understood that in identifying and describing these formats, this annex observes the current status of CCITT work on numbering plans, and does not establish any preference or position concerning the way in which CCITT may choose to modify the plans, or their relationships with one another, in the future. Changes to this annex may be necessary to take any such further work by CCITT into account. For example, the CCITT numbering plans in some cases may provide escape mechanisms (such as a zero, 8, or 9 prefix) from one numbering plan to another. This results in the possibility of a choice that must be made concerning which IDI format should be used for the allocation of NSAP addresses, and may also lead to suggestions that it is not necessary to include all of the IDI formats that are based on CCITT Recommendations in this annex. Such choices, however, are made within the context and responsibility of CCITT, and no preference for one choice or another is made or implied by this annex.

Table A.2 — Allocated AFI values

DSP syntax / IDI format	Decimal	Binary	ISO/IEC 646 character	National character
X.121	36, 52	37, 53	////////////////////////////	////////////////////////////
ISO DCC	38	39	////////////////////////////	////////////////////////////
F.69	40, 54	41, 55	////////////////////////////	////////////////////////////
E.163	42, 56	43, 57	////////////////////////////	////////////////////////////
E.164	44, 58	45, 59	////////////////////////////	////////////////////////////
ISO 6523-ICD	46	47	////////////////////////////	////////////////////////////
Local	48	49	50	51

NOTE — The Local IDI format is provided to accommodate the coexistence of OSI and non-OSI network addressing schemes, xparticularly in the context of a transition from non-OSI to OSI protocols. To provide the greatest flexibility in these environments, character and national character DSP syntaxes are defined for the Local IDI format.

A.5.2.1.2.1 X.121 IDI format

The IDI consists of a sequence of up to 14 digits allocated according to CCITT Recommendation X.121. The full X.121 number identifies an authority responsible for allocating and assigning values of the DSP.

IDP length: up to 16 digits.

A.5.2.1.2.2 ISO DCC IDI format

The IDI consists of a three-digit numeric code allocated according to ISO 3166. For countries with an ISO member body, the code is assigned to the ISO member body in the country identified by the code. For countries with no ISO member body, the code is assigned to an appropriately sponsored organization in the country identified by the code. The DSP is allocated and assigned by the ISO member body or sponsored organization to which the ISO DCC value has been assigned, or by an organization designated by the holder of the ISO DCC value to carry out this responsibility.

IDP length: 5 digits.

A.5.2.1.2.3 F.69 IDI format

The IDI consists of a telex number of up to 8 digits, allocated according to CCITT Recommendation F.69, commencing with a 2- or 3-digit destination code. The full telex number identifies an authority responsible for allocating and assigning values of the DSP.

IDP length: Up to 10 digits.

A.5.2.1.2.4 E.163 IDI format

The IDI consists of a public switched telephone network (PSTN) number of up to 12 digits allocated according to CCITT Recommendation E.163, commencing with the PSTN country code. The full PSTN number identifies an authority responsible for allocating and assigning values of the DSP.

IDP length: up to 14 digits.

A.5.2.1.2.5 E.164 IDI format

The IDI consists of an ISDN number of up to 15 digits allocated according to CCITT Recommendation E.164, commencing with the ISDN country code. The full ISDN number identifies an authority responsible for allocating and assigning values of the DSP.

IDP length: Up to 17 digits.

A.5.2.1.2.6 ISO 6523-ICD IDI format

The IDI consists of a 4-digit International Code Designator (ICD) allocated according to ISO 6523. The ICD identifies an organizational authority responsible for allocating and assigning values of the DSP.

IDP length: 6 digits.

NOTE — The use of an ICD in this context is in addition to, and does not affect the uses identified in ISO 6523. Of the things specified by ISO 6523, *only* the ICD is relevant to this annex.

A.5.2.1.2.7 Local IDI format

The IDI is null.

IDP length: 2 digits.

A.5.2.2 Abstract syntax and allocation of the DSP

Values of the DSP are allocated by the network addressing authority identified by the IDI in the syntax identified by the AFI (see A.5.1.1.2 and A.5.2.1.2). The allocating authority specifies the format and semantics of the DSP. If the authority identified by the IDI authorizes one or more authorities to allocate semantic parts of the DSP, then all of those authorities must allocate using the same abstract syntax used by the parent authority.

A network addressing authority may choose to allocate NSAP addresses with the DSP in a decimal or binary abstract syntax for all IDI formats. When the IDI format is “Local”, an authority may, in addition, choose to allocate NSAP addresses with the DSP in a character (ISO/IEC 646) or National Character abstract syntax (see Table A.2 and A.9). A network addressing authority may allocate NSAP addresses with no DSP (that is, addresses consisting of an IDP only) if and only if the value of the AFI specifies a *decimal* DSP syntax; a non-null DSP must be present for all other AFI values.

A.5.2.3 Abstract syntax of the DSP

The DSP may be allocated by the responsible authority in one of the following four syntaxes, depending on the value of the AFI:

a) *binary* — the DSP consists of one or more binary octets, up to the maximum specified in Table A.3,

b) *decimal* — the DSP, if present, consists of one or more decimal digits, up to the maximum specified in Table A.3,

c) *character* — the DSP consists of one or more of those ISO/IEC 646 graphic characters with no national variant, plus the space character, up to the maximum specified in Table A.3,

d) *national character* — The DSP consists of one or more characters from a national character set determined by the allocating authority, up to the maximum specified in Table A.3.

Table A.3 gives the maximum length of the DSP in its abstract syntax for each of the IDI formats defined in A.5.2.1.2. The corresponding total NSAP address lengths are given in A.5.4.

Table A.3 — Maximum DSP length

DSP syntax / IDI format	Decimal digits	Binary octets	ISO/IEC 646 characters	National characters
X.121	24	12	////////////////////////////	////////////////////////////
ISO DCC	35	17	////////////////////////////	////////////////////////////
F.69	30	15	////////////////////////////	////////////////////////////
E.163	26	13	////////////////////////////	////////////////////////////
E.164	23	11	////////////////////////////	////////////////////////////
ISO 6523-ICD	34	17	////////////////////////////	////////////////////////////
Local	38	19	19	9

NOTES
1 The values for the "Local" IDI format assume a National character representation of one character as two binary octets (see A.5.3).
2 These maximum values are dictated by the requirement that a the maximum length of an NSAP address in preferred encoding defined in A.5.3 be less than or equal to 20 binary octets.

A.5.3 Network address encodings

As described in A.5.1, the semantics of the NSAP address are represented by three fields in the following order:

a) the AFI, with an abstract syntax of two decimal digits;

b) the IDI, with an abstract syntax of a variable number of decimal digits; and

c) the DSP, with an abstract syntax of a variable number of one and only one of the following types; binary octets, decimal digits, characters, or national characters.

This annex does not specify the way in which the semantics of an NSAP address are encoded in Network layer protocols. These encodings are specified in Network layer protocol specifications.

Nevertheless, this annex identifies a "preferred" encoding of the Network address as a string of binary octets. Reference to this "preferred" binary encoding may be made by Network layer protocol standards. It is possible that the encoding used to convey the Network address semantics as Network protocol address information (NPAI) in a Network layer protocol may be chosen to be identical to this preferred encoding. However, it is not required that this be the case.

The entire NSAP address, taken as a whole, may be represented explicitly as a string of binary octets as defined below. Network layer protocol standards that specify the encoding of the Network address semantics by making reference to this annex must specify the way in which this binary encoding is used to convey the Network address semantics as NPAI (see A.3.1.3).

The preferred binary encoding defined in this clause requires that the IDI be padded with non-significant leading pad digits whenever: a) the AFI specifies a variable-length IDI format, and b) the value of the IDI is a string of decimal digits that is shorter than the maximum length of the IDI for that format (see A.5.2.1.2). This ensures that the end of the IDI (and thus of the IDP) can be determined; the preferred binary encoding does not reserve an explicit syntactic marker for this purpose. It is necessary, in these cases, to distinguish between significant and non-significant leading zero digits in the IDI, in order to ensure that non-significant pad digits are not confused with significant IDI digits. This distinction is provided, for each of the variable-length IDI formats. by the allocation of two AFI values for each combination of IDI format and the DSP abstract syntax (see A.5.2.1.1). In step b) below, the term "leading digits" therefore refers to leading zero (0) digits if the AFI value specifies that leading zero digits in the IDI are not significant; it refers to leading one (1) digits if the AFI value specifies that leading zero digits in the IDI are significant.

NOTE — The encoding defined in this clause requires that the IDI be padded to its maximum length, as described above, even when the value of the AFI specifies a decimal DSP syntax and the DSP is null.

The preferred binary encoding is generated by

a) using two semi-octets to represent the two digits of the AFI, yielding a value for each semi-octet in the range 0000-1001;

b) padding the IDI with leading digits if necessary to obtain the maximum IDI length (specified for each IDI format in A.5.2.1.2), then using a semi-octet to represent the value of each decimal digit (including leading pad digits, if present), yielding a value in the range 0000-1001; and, if the DSP syntax is not decimal digits, using the semi-octet value 1111 as a pad after the final semi-octet of IDI (if necessary) to obtain an integral number of octets;

c) representing a decimal syntax DSP by using a semi-octet to represent the value of each decimal digit, yielding a value in the range 0000-1001 for each digit, and using the semi-octet value 1111 as a pad after the final semi-octet of the DSP (if necessary) to obtain an integral number of octets;

d) representing a binary syntax DSP directly as binary octets;

e) when the IDI format is "Local", representing an ISO/IEC 646 character syntax DSP by converting each character to a number in the range 32-127 using the ISO/IEC 646 encoding, with zero parity and the parity bit in the most significant position, reducing the value by 32, giving a number in the range 0-95, encoding this result as a pair of decimal digits, and using a semi-octet to represent the value of each decimal digit, yielding a value in the range 0000-1001 for each digit, and

f) when the IDI format is "Local", representing a National Character syntax DSP by converting each national character to either one or two octets according to the rules specified by the authority responsible for allocating NSAP addresses including National Character DSP syntaxes.

A.5.4 Maximum Network address length

The maximum length of an NSAP address in the preferred binary encoding (A.5.3) is 20 octets. A Network layer protocol which is capable of conveying a string of variable length with a maximum length of 20 binary octets is therefore capable of encoding the full semantic content of any Network address.

A.6 Character based DSP allocation

A network addressing authority may choose to allocate NSAP addresses with the DSP in a National Character syntax. In such cases, the allocating authority must define and publish the mapping of the National Character syntax to a binary octet representation. The mapping must result in the representation of each national character as either one or two binary octets. The resulting DSP is considered to be based on a binary abstract syntax for the purposes of selecting AFI values from Table 2 and performing the preferred binary encoding defined in A.5.3.

NOTES

1 It is recommended that this mapping be done by reference to the *ISO Register of Character Sets*, which is maintained by the European Computer Manufacturers Association (ECMA) acting as a registration authority according to ISO 2375.

2 The ability to base DSP allocation on national character sets allows DSP allocation based on international character set standards such as ISO/IEC 646, and also allows DSP allocation based on specific nationally recognized character sets. This may simplify address assignment in some cases, and may facilitate representation of NSAP addresses in humanly-readable form. Nevertheless, NSAP addresses should not be confused with Application layer entity titles. NSAP addresses are not intended to provide the same degree of human-readable, user-friendly naming and addressing capabilities as may be expected in Application layer entity titles.

A.7 Reference publication formats

The reference publication format (RPF) is defined to facilitate unambiguous representation of NSAP addresses in both written and oral communication. It consists of a string of decimal digits which is the direct representation of the IDP, followed by the symbol "+", followed by a string of hexadecimal digits in which a pair of hexadecimal digits is used to represent the numeric value of each binary octet in the preferred binary encoding of the DSP (see A.5.3). In the case in which the DSP part of the NSAP address is zero length, the RPF consists only of the string of decimal digits representing the IDP; the symbol "+" is not present.

As an example, the written inscription of the RPF for an NSAP address with an AFI value of 39, an IDI value of 840, and a binary-syntax DSP value of 01001100 11100101 would be 39840+4CE5.

A.8 Network entity titles

In order to perform routing functions and to distribute Network layer management information concerning routing among Network entities, it is necessary to be able to unambiguously identify Network entities in end systems and intermediate systems. CCITT Rec. X.200 | ISO 7498 provides a definition of the concept of an (N)-entity title, which may be used to permanently and unambiguously identify a Network entity in an end system or intermediate system.

Any authority responsible for allocating addresses to NSAPs may choose also to allocate Network entity titles, following the same procedures and rules that it observes in the allocation of NSAP addresses. NSAP addresses and Network entity titles are syntactically indistinguishable; any value that the authority is permitted to allocate as an NSAP address may be allocated as a Network entity title.

Annex B
Rationales for the material in Annex A

(This annex does not form an integral part of this Recommendation | International Standard)

B.1 IDI formats (see A.5.2.1.2)

The rationale for the use of the specific IDI formats identified in A.5.2.1.2 is to allow the allocation and assignment of NSAP addresses to be based on existing, well-established network numbering plans and organization-identification standards.

The CCITT numbering plans are included so as to allow for the designation of the organization to which a number is assigned as an authority for the assignment of NSAP addresses. If the organization identified by a particular number from one of these plans chooses not to define any further subaddressing beyond that number, then the number itself constitutes an NSAP address when it is used in the OSI environment. This flexibility allows numbers allocated from the four CCITT numbering plans identified in A.5.2.1.2 to be used directly as NSAP addresses, with the addition of nothing more than the initial AFI digits that identify the plan.

The ISO DCC format is included so as to allow for the designation of the organization that represents a country in ISO (or an appropriately sponsored organization) as an authority for the assignment of geographically-based NSAP addresses. The way in which addresses are allocated and assigned in the ISO DCC format is determined by the designated organization, which might, for example, be the national standards body that represents a country in ISO.

The ISO 6523-ICD format is included so as to allow for the designation of an organization that may or may not be tied to a particular country as an authority for the assignment of NSAP addresses according to the hierarchy appropriate for the organization (which may not be based on geographical or national boundaries). The way in which addresses are allocated and assigned in the ISO 6523-ICD format is determined by the designated organization, which might, for example, be the United Nations World Health Organization. The ISO 6523-ICD format permits an organization already possessed of an ICD, for the purposes specified in ISO 6523, to use that ICD for the *additional* purpose of allocating Network addresses. This additional purpose is unrelated to the role of the ICD as the identifier of an Organization Code (OC) assignment scheme, which is the purpose for which ICDs are assigned. This does not change the criteria established in ISO 6523 for granting a request for an ICD allocation. Furthermore *only* the ICD is used in Network addresses that follow the ISO 6523-ICD IDI format. No part of any Network address corresponds to the OC defined in ISO 6523, and the OC is therefore not relevant to or affected by this.

The local format is included so as to allow for the coexistence of proprietary or other non-standard network addressing schemes with the standard OSI network addressing scheme. Use of the Local format for these non-standard addresses ensures that they cannot be confused with standard OSI network addresses. This capability will be useful in the evolution of existing networks to OSI, and for the accommodation of non-OSI addressing schemes that may be used in proprietary network architectures or for testing and other interim purposes. It should be emphasized that the Local format is not intended to give non-OSI schemes a permanent place in OSI, but rather to permit the OSI network addressing scheme to be used wherever possible without risk or conflict with other schemes (which can be encapsulated safely under the Local format).

B.2 Reservation of AFI values 00-09 (see Table A.2)

The reservation of AFI values beginning with the digit 0 is intended to allow for the use of an initial 0 handle special cases, such as

a) as an escape to some other addressing scheme;

b) as a technique for the optimization of NSAP address encoding in Network layer protocols, when different parts of NSAP address semantics are encoded in different fields of the protocol header;

c) as a way to indicate, in a protocol header, that a field that ordinarily contains a full NSAP address in fact contains something less than a full address (for example, a shorthand form that omits specification of the higher-order addressing domain(s), which might be used for communication within a particular subdomain environment).

There may be other cases in which the use of an initial zero digit is found to be useful. Annex A merely reserves the AFI values 00-09, and does not specify how they might be used; all such uses are outside the scope of Annex A.

B.3 Derivation of the preferred encodings (see A.5.3)

In describing the preferred binary encoding of the NSAP address, clause A.5.3 introduces two types of padding: padding with non-significant leading zero or one digits at the beginning of an IDI, and padding with a semi-octet with the value 1111 at the end of the binary encoding of an IDI with an odd number of decimal digits.

The first type of padding is necessary because in some formats the IDI consists of a variable number of digits. Since there is no explicit syntactic marker between the IDI and the DSP, the only way to find the boundary between them is to know how long the IDI is. The AFI, which identifies the IDI format, specifies only the *maximum* length of the IDI in that format. Rather than introduce either a specific syntactic marker or a new field containing the length of the IDI (either of which would complicate the encoding and parsing of NSAP addresses), Annex A specifies that for encoding purposes the IDI must first be padded out to its maximum length. Note that this does *not* apply to the DSP; only to the IDI.

The second type of padding is necessary to ensure that a binary encoding of the IDI consists of an integral number of binary octets.

Annex C
Facilities for conveying service characteristics in the connectionless-mode Network Service

(This annex does not form an integral part of this Recommendation | International Standard)

C.1 Introduction

In order to effectively coordinate the choice of transport protocol with the provision of the connectionless-mode Network Service, some mechanism must be made available to the Transport Layer for relating certain characteristics of the service which are

a) necessary for the correct operation of the transport protocol (for example, NSDU lifetime);

b) not clearly a part of the QOS which could be specified in a single invocation of the service (for example, sequence preservation); or

c) appropriately some aspect of Network Layer management (for example, congestion control).

Thus, a control flow independent of the data flow modeled in clause 16 can be modeled to describe the facility (or set of facilities) required to correlate transport protocol selection with the service made available between a pair of NSAPs. This control flow can be illustrated from the perspective of the (local) NS user as shown in Figure C.1.

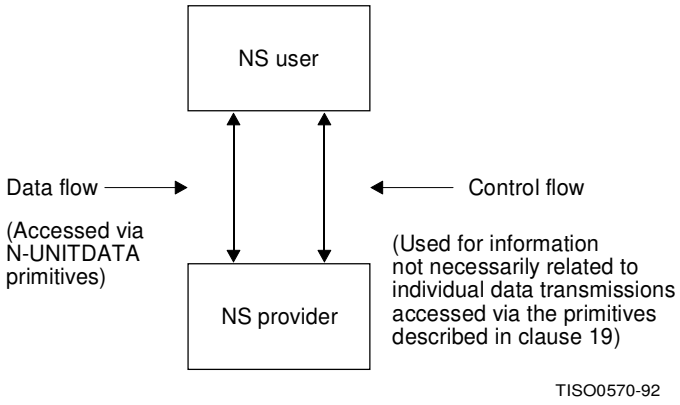

Figure C.1

A set of primitives describing the means by which control information is conveyed from the NS provider to the NS user has been provided in the following clauses. These primitives are defined in accordance with CCITT Rec. X.210 | ISO TR 8509.

C.2 Function

This annex describes primitives which convey information from the NS provider to the NS user concerning those characteristics of the service available for a given pair of NSAPs which are beyond the scope of those characteristics which can be expressed in a single invocation of the connectionless-mode Network Service.

The primitives are modeled as requests initiated by the NS user and indications initiated by the NS provider. Indications may be initiated by the NS provider independently from requests by the NS user. The circumstances which cause the initiation of these primitives are not discussed here.

NOTE — Primitives providing such information are required for effective application of the service. It is recognized that such primitives may equally be regarded as a layer management function.

C.3 Types of primitives and parameters

Three primitives are defined:

a) the N-FACILITY request;

b) the N-FACILITY indication; and

c) the N-REPORT indication.

The N-FACILITY request is issued by the NS user to request information about the characteristics of the service which may be expected to be available for an N-UNITDATA request(s) to a given destination NSAP.

The N-FACILITY indication conveys characteristics of the service to the NS user which may be expected to be available for N-UNITDATA request(s), given a particular destination NSAP. This primitive is issued by the NS provider in response to an NS user initiated N-FACILITY request, or it may be issued by the NS provider independent of any request by the NS user.

NOTE — Such information may be available, for example, from routing tables maintained by the NS provider, or from information supplied as an aspect of layer management, or by some other means.

The N-REPORT indication conveys information about a failure to provide the requested quality of service or service characteristic to the NS user relating to a failure to satisfy the constraints imposed upon the N-UNITDATA request, given a particular destination NSAP.

NOTE — This primitive will be issued by the NS provider in response to information pertaining to the failure of a particular N-UNITDATA request to be fulfilled. This information is typically obtained through the operation of the supporting Network Layer protocol. In addition, if the local subsystem determines immediately that the requested transmission cannot be fulfilled under the conditions imposed, it may abort the transmission and issue the primitive. The specific circumstances under which this primitive is generated are not discussed herein.

C.4 Service characteristics

C.4.1 Congestion control

The specification of the congestion control parameter is concerned with whether an NS user is willing to accept indications from the NS provider that transmission is refused. The specification of this parameter means that the NS user prefers refusal of transmission to the initiation of discard operations.

C.4.2 Sequence preservation probability

Sequence preservation probability is the ratio of sequence-preserved transmissions to total transmissions observed during a performance measurement.

A sequence-preserved transmission is one in which the following rule is observed: for a given pair of NSAPs, an N-UNITDATA indication will not occur at a destination NSAP after any other N-UNITDATA indications resulting form any subsequent N-UNITDATA requests issued from the source NSAP to the destination NSAP.

For a series of transmissions for which all other QOS parameters are the same, the specification of a sequence preservation probability of one (1) means that the NS provider shall ensure that change of order operations (see clause 16) are not initiated.

NOTES

1 According to the definition provided in ISO 7498/Add. 1, it is not a requirement that sequence preservation be provided by a connectionless-mode Network Service. While it is true that in some circumstances, NSDUs will be delivered in order at the destination NSAP if sequence is maintained by the underlying service, the fact that a series of data units handled one after the other by an NS user to the NS provider are delivered to the destination in the same order is not a basic characteristic of a connectionless-mode Network Service, as no relationship between the NSDUs, expressed or implied, is created by the connectionless-mode Network Service.

2 The maintenance of sequence from one invocation of the connectionless-mode Network Service to another, between the same pair of NS users may also be viewed as being part of the connection-mode Network Service; for example, when the network-connection is established at system initialization and exists permanently without an explicit connection establishment phase.

C.4.3 Maximum NSDU lifetime

Maximum NSDU lifetime is the maximum elapsed time between an N-UNITDATA request and any corresponding N-UNITDATA indication. Maximum NSDU lifetime is made available to the NS user in order that timer values associated with transport protocol operation might be correctly assigned.

NOTE — A connectionless-mode Network Service provider will, in general, continue to try to deliver an NSDU until the value of this parameter is reached. Typically, resources (such as buffers) are employed to achieve this. Where the NS user employs retransmission mechanisms, the maximum NSDU should not be set to a value much greater than the retransmission timer. Otherwise, the NS provider may fail to meet the other QOS requirements due to the need to retain resources for several NSDUs containing equivalent data.

C.5 Types of primitives and parameters

Table C.1 illustrates the types of primitives and the parameters needed for the provision of the management facilities previously described.

Table C.1 — Facility and Report parameters and primitives

Primitive / Parameter	N-FACILITY		N-REPORT
	Request	Indication	Indication
Destination address	X	X	X
Service characteristic QOS parameter	X (see Note)	X	X
Reason for report			X
NOTE — The parameter(s) may be implicitly associated with the primitive. There might be specific parameters for each individual service characteristic or quality of service parameter, or implicit requests for a group, or all of the parameters according to local implementation conventions.			

C.5.1 Destination address

The address referred to in Table C.1 is an NSAP address. Both the connection-mode and connectionless-mode Network Service use the same NSAP addresses, as described in Annex A.

C.5.2 Service characteristic/QOS parameter

The value of the service characteristic/QOS parameter is a list of sub-parameters. For each sub-parameter, any value permitted according to the particular sub-parameter may be indicated. QOS parameters are described in clause 17; service characteristics are described in C.4.

C.5.3 Reason for report

The value of the reason for report parameter indicates the reason that the N-REPORT indication was initiated by the NS provider. Reason codes might include:

a) no reason specified;

b) transit delay exceeded;

c) NS provider congestion;

d) other requested QOS/service characteristic unavailable;

e) NSDU lifetime exceeded; and

f) suitable route unavailable.

In the case where reason d) is indicated, the NS provider may return an indication of which quality of service parameter is not available.