TRANSPORT SERVICES AND PROTOCOLS FOR OSI.

L. LENZINI

Rapporto interno
C81-22

CNUCE - Aprile 1981

Progetto Finalizzato Informatica
Obiettivo: COMPUNET

A.I.C.A.
ASSOCIAZIONE ITALIANA
CALCOLO AUTOMATICO

20121 MILANO
Piazzale R.Morandi 2
Tel. (02) 784970

C.R.E.I.
CENTRO RETE EUROPEA INFORMATICA
20133 MILANO Piazza Leonardo da Vinci 7
Tel. (02) 296826

Giornata di studio su:

ARCHITETTURA OPEN SYSTEM INTERCONNECTION (O.S.I.) proposta dell'I.S.O. per la connessione di sistemi informativi in una rete

Milano, 15 Maggio 1981

Con l'adesione dell'UNIPREA (Ente di Unificazione Federato all'UNI)

Sede del Convegno:

Palazzo FAST
Federazione delle Associazioni
Scientifiche e Tecniche
20121 Milano - Piazzale R.Morandi 2

Transport Services and Protocols for CSI

L. Lenzini
CNR-Istituto CNUCE
Via S. Maria, 36
56100 Pisa
Tel. (050) 45245
Telex 500371

Abstract

A Transport Layer has been defined as Level 4 of the Open Systems Interconnection (OSI) Basic Reference Model. Its major characteristic is that it forms the boundary between data communication functions (the three lower layers of the Model) and data processing functions (Levels 5 through 7). The International Organization for Standardization (ISO) is currently defining a Standard Transport Service and Protocol for this Layer. Although both technical and political problems are presently hindering the achievement of this objective, it is clear that a firm and generally accepted international standard in this area is urgently needed by users and manufacturers. This paper presents rationale and justifications for the proposals under discussion, attempts to clarify the motivations for the differing viewpoints and suggests how they could be reconciled by application of a realistic and engineered approach to the problem.

1. Introduction

The transport layer is a key issue at the present time because:

- it is the termination of the functions which transfer data through the network;
- it has to isolate the upper layers from the different network technologies;
- it is the layer which offers its services for a variety of applications and should support their different requirements.

The main goals of the transport layer service (in a connection oriented environment) is to provide the transport users (session entities) with:

a) a reliable transport medium. That means that the user of layer 4 is free from the concern of reliable transmission. The transport layer ensures that units of data are received as they were sent i.e. without corruption, misordering, or duplication;

b) transparency. That means that messages presented for transportation by a transport user will be transferred and presented to the partner user of the transport layer without interpretation (by the Transport Service) of data contents;

c) cost effectiveness. That means that the transport layer contains mechanisms to optimize the transmission costs of the underlying network. The transport layer is required to optimize the use of the available communication resources to provide the performance required by each communicating transport user at minimum costs;

d) end to end significance. This means that it provides a service between users of different interconnected systems irrespective of their location and the network topology (number of intermediate nodes, network access mechanisms etc).

2. The transport Service

2.1 Model Used for Service Specification

The adopted method to describe the transport service in a non-ambigous manner is the transport service definition in terms of a set of implementation independent service primitives. Such a transport service definition is the first and necessary step toward the specification of the transport layer protocol /TS80/. The service primitives are conceptual and are not directly related to transport protocol elements, nor seen as "macro calls" of an access method to the transport service. Only those aspects which impact the communication between the users of the transport layer are considered in the service primitives. The mechanisms which are only related to local conventions between the transport service user and the transport service provider are not defined in the model. A service primitive is defined by:

- the unit of service which it provides,
- its 'direction',
- the set of parameters associated with it.

The direction of a service primitive is the relationship betweem the user of the transport layer and the transport layer. Service primitive request is a primitive directed from the user of the transport layer to the transport service prívider whereas a service primitive indication is directed vice-versa. The service primitive parameters are specified by their range of permissible values. The transport service is unambigously specified by the definition of the service primitives and the definition of the service primitives relationships.

2.2 Transport Service Primitives

In this chapter we summarize the Transport Service Primitives as proposed in documents provided by the main standardization groups. A short description of their functionality is given based on documents /TS80/. Only 'connection oriented' transport services are defined at the present time in the official documents, but a lot of work has to be mentioned for the 'connectionless' and Fast Connect/Disconnect services. The results may be expected to be integrated in future ISO documents.

3. Connection Oriented Services

This section describes informally the Services provided to the Session Layer from the Transport Layer. These Services are Connection Oriented, i.e. they consist of Services for establishing, and terminating Transport Connections. These Transport Connections are two-point Connections. The Transport Connection will be in the establishment, data transfer, or termination phase, respectively, depending on which of the three Service elements is currently active.

3.1 <u>Connection Establishment Phase/ Connection Termination Phase</u>

a) <u>Transport Connection Establishment Service</u>

This Service enables a Session Entity to establish a data transfer capability, known as a Transport Connection. The Quality of Service of the Transport Connection is negotiated with the Session Entity (Entities) and the Transport Service; i.e. the Transport Connection Establishment Service provides the Session Entity with the ability to request Transport Connection parameters and a Service Quality Class selected from a predefined set of classes. These classes would specify values for the attributes on a Transport Connection, such as throughput rate, response time, reliability of data transfer on a Transport Connection. (see 4.) A Session Entity has the option of accepting or rejecting a Connection Request. A quantity of Session Data may be transferred with the request to establish a Transport Connection. This data may be evaluated by the receiving Session Entity prior to deciding whether to accept or reject the Connection Request. This Session Data is transferred transparently. If the TS-provider is unable to provide this Service, a disconnect indication is given. The TS will announce the detection of an incoming Transport Connection Request to the addressed Session Entity. The established Transport Connection represents a two-way simultaneous data path between a pair of Transport Service Access Points (TSAPs).

<u>Establishment Service Primitives</u>

```
                            +--------------+   +--------------+
1: CONNECT request          | calling user |   | called user |
2: CONNECT indication       |  1:|     4:^  |   | 2:^     3:|  |
3: ACCEPT request          +----v------|-------------|-------v----+
4: ACCEPT indication       |      transport service provider      |
                           +--------------------------------------+
```

b) <u>Transport Connection Termination Service</u>

This Service provides the means by which either Session Entity associated with a Transport Connection can terminate the Transport Connection and inform the corresponding Session Entity of this request. This Service does not guarantee data delivery if the data is undelivered at the time the request was issued. If the TS is unable to maintain the Service Quality agreed upon for a particular Transport Connection then it terminates the Transport Connection and informs both Session Entities.

<u>Termination Service Primitives</u>

```
                                    +----------------+   +--------------+
1: DISCONNECT request*              | calling user   |   | called user  |
2: DISCONNECT indication            |  1:| or 2:^    |   | 2:^  or 1:|  |
                                  +---------------------------------------+
* user initiative                 |     transport service provider        |
                                  +---------------------------------------+
```

```
                                    +----------------+   +--------------+
2: DISCONNECT indication*           | calling user   |   | called user  |
                                    |        2:^     |   | 2:^          |
                                  +---------------------------------------+
* provider initiative             |     transport service provider        |
                                  +---------------------------------------+
```

<u>Parameters of the service primitives</u>

The establishment phase may be mainly understood as a three party negctiatioin phase between the two users of the transport service provider and the transport service provider itself. They negotiate the willingness to communicate and the quality of this communication. The information contained in the parameters may be classified in:

- parameters for the identificatioin of the two users and the connection;
- parameters defining the quality of service (QoS) required in the communication (e.g. throughput, transit delay, maximum data length etc.);
- reasons for the termination of a transport connection;
- transparent data.

3.2 <u>Data Transfer Phase</u>

a) <u>Normal Data Transfer Service</u>

This Service enables arbitrarily selected Transport Service Data Units (TSDUs) to be delimited and transparently transferred from the sending Transport Access Point to the receiving one, over a Transport Connection. The flow on the TSDUs across the Transport Service Access Point is controlled by the receiving party. The delay which can be caused by flow control is limited by the agreed upon class of Service. The effect of a "flow control not ready" of a receiving Session Entity is local, i.e. it is not specifically indicated to the sending Session Entity. The only mechanism of propagating it is by back pressure TSDUs to the submitting Transport Service Access Points which causes there a "flow control not ready" of the Transport

Service.

Data transfer service primitives

```
                           +-------------+   +-------------+
1: DATA request            | calling user|   | called user |
2: DATA indication         | 1:|a/o 2:^  |   | 2:^a/o 1:|  |
                         +---------------------------------+
(a/o means: and/or)      |    transport service provider   |
                         +---------------------------------+
```

b) Expedited Data Transfer Service

This Service enables arbitrarily selected Expedited Transport Service Data Units of limited size to be delimited and transparently transferred from one Transport Access Point to another as rapidly as when using the Normal Data Transfer Service. Such expedited data may bypass some normal data in their delivery to the receiving Transport Service Access Point. The amount of normal data bypassed cannot be predicted or guaranteed. However, no expedited data will be delivered to the receiving Transport Service Access Point later than normal data which have been sent subsequent to the expedited data.

Expedited Data Primitives

```
                                   +-------------+   +-------------+
3: EXPEDITED DATA request          | calling user|   | called user |
4: EXPEDITED DATA indication       | 3:|a/o 4:^  |   | 4:^a/o 3:|  |
                                 +---------------------------------+
(a/o means and/or)               |    transport service provider   |
                                 +---------------------------------+
```

c) Reset Service

This Service enables a Session Entity to reset a Transport Connection. This may result in loss of some or all of the data (normal and expedited) in the Transport Connection at that time. As a consequence, this Service requires recovery mechanisms to be effected at Transport User level. The Reset Service may be used by the Transport User to resynchronize the use of the Transport Connection and will unblock the flow of TSDUs if there is congestion of the Transport Connection.

Just recently the question has been raised as to whether the Reset Service is needed at the Transport Layer.

Parameters

The only parameter for the data transfer service primitive are data transmitted by the user of the transport service.

4. Quality of Service parameters

In the following we attempt to classify the parameters of QoS as given in /TS80/.

The set of characteristics identified below is mutually exclusive. However, a given parameter may have one, or more than one characteristic at any given layer boundary (e.g. a parameter may be of local significance at the session/transport layer boundary, and of protocol implication significance at the transport/network layer boundary).

- end user to end user significance

  A parameter with this significance passes across the local session/transport layer and the remote transport/session layer boundaries. This means that it is of interest to the remote session entity.

- local significance

  A parameter with this significance passes across the session/transport boundary and goes no further. It is dealt with by only the local transport entity, which usually relies on information from the local network (or link, or physical) entity.

- protocol implication significance

  A parameter is considered to have a protocol implication when it induces the selection of options in the local transport entity or when it causes some actions of the transport protocol to be performed with its peer entity. A parameter with this significance need not pass across the remote transport/session boundary because it is not of interest to the remote session entity.

4.1 Discussion of the parameters

a) Throughput

This parameter has end user to end user significance. In addition this parameter has protocol implication significance if:

- there are means (e.g. TPDU numbering and credit mechanism) in the protocol to control the throughput;
- the requested throughput is greater than that provided by a network connection. In this case, the transport entity may request the use of several network connections in parallel and distributes the TPDU indifferently onto the network connections (downward multiplexing) Adequate transport protocol elements are required.

b) Transit Delay

This parameter is of <u>local significance</u> if the Transit Delay reported by the network service (on a given network connection) is acceptable or better than the requirement. If this is not the case, the Transit Delay is of <u>protocol implication significance</u>. Adequate elements of protocol are used in this case. In fact, if the Transit Delay reported by the Network service on a given network connection is not acceptable, an improvement may be achieved by shortening the NSDUs or/and sending them onto several network connections (downward multiplexing).

c) Residual Error Rate (RER)

If the underlying network(s) has a stated RER which is higher than that specified, then a number of transport protocol mechanisms can be invoked to upgrade the RER using:

- TPDU sequence numbering;
- acknowledgement and time out;
- check sum;
- connection identification;
- retransmission;
- forward error correcting code.

This parameter is thus of <u>protocol implication significance</u>.

d) Connection Set-up Delay

This parameter is of <u>local significance</u> if the Set-up Delay reported by the network service is acceptable or better than the requirement. If this is not the case, the <u>Connection Set-up Delay</u> is of <u>protocol implication significance</u>. Adequate elements of protocol are then needed in this case. In fact, if the Connection Set-up Delay reported by the network service is longer than the delay required, the following protocol mechanisms may be used to shorten it:

- re-use of network connections for successive transport connections;
- multiplexing onto an existing network connectioin if unused capacity of it is available.

e) Resilience of Transport Connection (TC)

Failures of any underlying network connection may be masked from the transport service user if the transport service re-establishes the network connection (or uses an already existing network connection).

This parameter is of <u>protocol implication significance</u>

as it induces the selection of appropriate transport protocol mechanisms (e.g. network connection reconnect mechanisms).

f) Cost Saving

This parameter remains of local significance only if the achieved cost reported by the network service is less than or equal to the required cost. In this case no further action is required.

This parameter has protocol implication signigicance if the achieved cost reported by the network service is higher than that required. In this case a number of transport protocol mechanisms can be invoked to meet the cost saving transport user requirement:

- upward multiplexing;
- sequential reuse of a maintained network connection;
- use of low cost network connections in parallel;
- select TPDU size which best fits into the NSDU;
- use connectionless data transfer;
- avoid unnecessary costly functionalities.

g) Protection

This parameter is basically of end user to end user significance and the corresponding information must be passed from the local session entity to the remote session entity if there are no mechanisms in the transport layer (enchryption,...) to improve this QoS parameter.

However, the parameter becomes of protocol implication significance if there are suitable mechanisms in the transport protocol to improve this QoS parameter.

h) Connection Assurance

The implementation of this parameter implies that, idle traffic has to be generated (if necessary in both directions independently) by the transport layer on a transport connection during periods of quiescence. The idle traffic is not seen by the session layer. The Connection Assurance parameter is thus of protocol significance. It should be noted that this parameter will have severe cost implications on some networks.

i) Priority

This parameter has end user to end user significance. Obviously it is highly related to the Throughput parameter and both parameters may be combined into a single parameter which could be known as the "traffic class".

5. Transport Protocols

The protocols may be viewed as the mechanisms through which (physically) separated entities in a layer perform the functionality of that layer. Note that the functionality of the layer is defined by the specification of the services provided by the layer and the services offered by the lower layer. Thus the functionality of the transport layer and thus the transport protocols are determined by the services provided by the network layer and the services requested by the user of the transport layer.

Whereas the Transport Service (primitives) and Network Service (primitives) may differ mainly in level of 'quality' that can be requested and offered, the transport protocols may range from simple to complex.

Although the Network Layer has not yet been completely defined, an attempt can be made to identify different Network Service Classes by considering the wide variety of Networks already in operation or about to become operational. These include public and private, distributed as well as local, Networks.

On these Networks, techniques such as packet and circuit switching are used side by side with local wire broadcast Networks, satellite channels, multi-drop and point-to-point circuits. Each of these transmission media has its own, particular characteristics.

For reliability enhancement purposes, one of the major goals of the Transport Layer of each of the network types outlined above can be classified in accordance with their intrinsic reliability and their signalling of non-recovered errors. The purpose of such a classification is to provide a basis for decisions such as to install error detection or recovery mechanisms in the Transport Protocol to increase the reliability by recovering from Network failures.

We may recognize two ways of error signalling by a network:

a) environments in which errors (mis-ordering, loss, duplication, etc...) are not signalled by the Network to the upper Layer. Such situations exist with:

   - Datagram Networks (mis-ordering, loss, or duplication are possible);
   - several X.25 Virtual Circuits in parallel (mis-ordering);
   - X.25 Networks in which some errors are not signalled.

b) environments in which all errors are signalled by the Network to the upper Layer. It is perhaps interesting to distinguish between Networks in which signalled errors are relatively frequent and others in which such errors are relatively rare. Such situation may arise with:

- X.25 Networks;
- HDLC links;
- X.21 Networks together with HDLC data link procedures;
- telephone Networks with HDLC data link procedures.

Another major goal of the transport layer is to provide a cost effective data transfer. This may be realized by so called 'upward multiplexing' protocol mechanisms which means that a network connection may be used to support several transport connections.

Additional protocol elements such as flow control, connection identificatioin, downward multiplexing, etc. are needed to support these goals and goals expressed in other quality of service parameters.

These are only some of the rationale to explain the variety of the proposed transport layer protocols, classes of protocols and options resulting from the already existing documents. Analysing the different available documents one may distinguish in these specifications different classes of protocols.

The ISO document N537 clarifies the role of the Network and Transport Layers in the transport of data between Open Systems. In particular it states that the Transport Layer:

1) "performs any functions necessary to enhance the Quality of Service provided by the Network Layer";

2) "is principally concerned with the Quality of Service";

3) "is required to optimise the use of the available communications resources to provide the performance required by each communicating Transport User at minimum cost".

Keeping in mind the above considerations, in the following we point out two approaches to enhance the Quality of Service provided by the Network Layer.

- hierarchical protocol classes (/TS80/, /TC80/) in which the classes are ordered in sequence, each next class including the functions of the previous class and adding new functions;
- disjoint protocol classes based on the quality of service they enhance.

In section 7 we will try to analyse the pros. and cons. of the two approaches.

5.1 Hierarchical Protocol Classes

ECMA, the European Computer Manufacturer Association, which counts among its members most if not all af the major computer manufacturers, also including U.S. based ones, has recently standardized a Transport Protocol for OSI /ECM80/. This protocol is the one forming the current bases of the ISO work in this area /TP80/. The ECMA Transport Protocol proposal makes available two qualitative enhancements to the Network Service. These are a multiplexing enhancement which, depending on Network tariffs, can reduce costs, and a recovery enhancement which can increase reliability by recovering from Network failures.

ECMA proposes to organize all the Transport Protocol functions hierarchically into five Classes and options, in such a way that:

a) each Class consists of a predefined set of functions;

b) these Classes are structured so as to form a linear sequence, each Class including the functions of all previous ones;

c) options can define additional functions which may be associated with any Class and for which negotiation is necessary (the existence of options allows escape to hierarchy);

d) Class 0 (TELETEX /CCI80/ oriented) cannot have options;

e) Classes and options are negotiated during the Transport Connection Establishment phase (this is much simpler than negotiating each function separately);

ECMA, thus, proposes to organize the Classes as shown in the Tables 1 and 2 on the following page.

PROTOCOL ELEMENTS	CLASS 0	1	2	3	4
Multiplexing	N	Y	Y	Y	Y
Flow Control	N	N	Y	Y	Y
Error Recovery	N	N	N	Y	Y
Error Detection	N	N	N	N	Y

Table 1: ECMA Classes Vs Protocol Elements

SERVICES	CLASS 0	1	2	3	4
Connection Establishment	Y	Y	Y	Y	Y
Data	Y	Y	Y	Y	Y
Expedited Data (option in ECMA)	N	N	Y	Y	Y
Reset (option in ECMA)	N	N	Y	Y	Y
Connection Termination	N	Y	Y	Y	Y

Table 2; ECMA Classes Vs Transport services

The selection of the Protocol Class is made by the Transport Entities according to:

- the Quality of Service negotiated between Session Entities via the (CONNECT) request and (ACCEPT) request service primitives;
- the Quality of Service provided by the available Network.

With regard to the above mentioned ECMA classification, it is worthwile pointing out that, with the exception of Class 0, Transport Connections of different Classes may be multiplexed together onto the same Network Connection. A typical example might mix the flow controlled Classes 2, 3 and 4 although Class 1 would normally not be included as this Class does not provide a flow control mechanism.

Class 0-Simple Terminal Class

This Class has been defined by CCITT and is being used to support TELETEX terminals connected to switched Networks.

Class 1

This Class would typically be used for unsophisticated terminals, and when no multiplexing onto Network Connections is required. However, it can also be used for multiplexing in particular cases when the absence of individual Transport Connection flow control is acceptable; i.e. for Transport Connections with no critical response time requirements, or with infrequent short bursts of traffic and a predictable low combined level of Network Connection utilization.

Class 2-Flow Control Class
==========================

This Class was specified for heavy and continous traffic, or for intensive multiplexing. It provides flow control to avoid congestion at the end points of the network connection using a credit mechanism to allow the receiver to inform the sender of the amount of data he is willing to receive.

Class 3-Error Recovery Class
============================

This Class provides the functionality of Class 2 plus the ability to recover after a failure signalled by the Network Layer without involving the user of the Transport Service. The mechanisms used to achieve this functionality also allow the implementation of more flexible flow control.

Class 4 - Error Detection and Recovery Class
==

This Class provides the functionality of Class 3 plus the ability to detect (and recover as in Class 3) corrupted, lost, duplicated, or out of sequence TPDUs without involving the transport service user.

5.2 Protocol Classes based on Protocol Options

User needs can differ widely, and clearly the individual user may see no reason for an uniform Qualify of Service hierarchy. Some will need certain characteristics, others may desire others. Thus, judgements on the acceptability of different Quality of Service characteristics for specific Transport Users must be performed separately. In this way, the particular Quality of Service characteristics which need to be improved by the Transport Protocol can be identified. Enhancement mechanisms should then be invoked to meet these needs. The possibility to select each identified Quality of Service characteristic independently may not be necessary. A Class structure, in which each Class represents a predefined combination of parameters, may be desirable. The classes are intended to cover the Transport Service requirements of the various types of traffic generated by the Session Entities (e.g. batch type traffic, interactive type traffic, etc.). The negotiation of such a Class, at the time of Transport Connection Establishment, will certainly not impose a linear sequence in the resulting Protocols. At this moment no reference can be made to documents that propose Transport Protocols that are based on Quality of Service parameters, although experts within SC16 have announced to make such contributions.

6. Quality of Service enhancement:

- based on hierarchical protocol
- based on protocol options

This section is a further elaboration of section 4. It focuses in particular on the relationship between protocols and QoS parameters.

6.1 Based on Hierarchical Protocols

As has been pointed out in section 5.1, the ECMA transport protocol proposal is aimed at QoS enhancement of the network service in terms of residual error rate and cost saving. However, in document /TS80/ a list of QoS parameters are identified, defined and their units are given. The following discussion identifies the QoS parameters which can be enhanced for ech ECMA transport protocol class. A rationale given takes two important aspects into account:

- depending on the ECMA protocol class, QoS enhancement can be explicitely supported (indicated in Tab. 3 by *);
- QoS enhancement is possible if additional functionality is added to the supporting transport entity (indicated in Tab. 3 by o).

In addition, certain aspets of the ECMA proposal for the transport protocol now either under debate within ISO or still open and subject to further study are outlined.

Class 0
=======

No qualitative enhancements to the Network Service are explicitly provided by this class which can then be used in the case of a single underlying Network of high quality (e.g. certain X.25 Networks).

- Cost saving enhancement can be achieved if the supporting transport entity has the ability to select the TPDU size which best fits into the NSDU.

This holds when the tariffs of the network service are based on fixed length NSDUs irrespective of its utilization.

Class 1
=======

Transport Connection Identification and explicit Disconnection related to Class 1 enable the Transport Connection lifetime to be independent from the Network Connection lifetime. This implies that a Network Connection can be reused for another Transport Connection. Therefore, in addition to the cost saving enhancement already described for class 0, the following QoS enhancements can be achieved:

- Connection Set-up Delay, which can be achieved when the supporting transport entity provides a sequential reuse of a maintained (already existing) network connection. This possibility implies an additional cost saving enhancement when the tariffs of the network service are based on a significant network connectioin establishment figure.

- Resilience of TC, which can be achieved when the supporting transport entity provides recovery from a network-service-provider initiated disconnected either by establishing another network connection or by using an already existing but at that moment unused network connection.

  However, this QoS enhancement is possible if and only if data is not lost during the network-service-provider disconnect procedure.

NOTE

It is generally agreed that there is some redundancy between Class 0 and Class 1. Differences between the two classes are made apparent on examination of the functions within class 1. It is hoped by ISO people that in the future these two classes will be merged. Collaboration between ISO and CCITT (promoter of Class 0) could lead to one common Class.

Class 2
=======

The following QoS enhancement is explicitly supported in addition to those already reported for Class 1.

- Cost saving, which is achieved by multiplexing. This holds where the tariffs of the network service take the number of network connections to be supported into account.

Class 3
=======

The increased reliability offered by network failure recovery is only available when Class 3 of the ECMA protocol is used. Therefore, this protocol class explicitly supports the following QoS enhancement in addition to those already given for Class 2.

- Resilience of transport connection, which is achieved by recovering from network service provider notification.

NOTE

As it was underlined in section 5.1, Class 3 also includes a multiplexing function. It is strongly felt by some experts that a non multiplexed recovery class should be included in the protocol to provide the recovery enhancement without the overheads in processing the Network traffic processing which multiplexing imposes.

The reason which is often quoted to ustify this is that a destination reference is necessary for recovery and this will also allow multiplexing. However, nowhere in the proposal does it say that multiplexing is mandatory in any class.

One view is that use of the destination reference when no multiplexing is required would result result in unnecessary line overhead.

Another view is that the use of this reference can simplify algorithms in the system and thus facilitate the implementation. For example, it should be much easier for a Transport Entity to access to the control information related to the Transport Connection under way, if the destination reference is available.

Class 4
=======

The following QoS enhancements are explicitly supported by Class 4, in addition to those of Class 3:

- Throughput, which can be achieved when the supporting transport entity provides splitting (downward multiplexing);
- Transit delay, which can be achieved when the supporting transport entity provides splitting;
- Residual Error Rate, which is supported by recovering from loss, duplication and out of order TPDUs.

Tab. 3 summarizes the relationship between the ECMA protocol classes and the QoS parameters as identified in /TS80/ (Transport Service)

QoS parameter	ECMA Transport Protocol Class 0	Class 1	Class 2	Class 3	Class 4
Throughput	-	-	-	-	0
Transit Delay	-	-	-	-	0
Residual Error Rate	-	-	-	-	*
Connection Set up Delay	-	0	0	0	0
Resilience of TC	-	0	0	*	*
Cost Saving	0	0	*	*	*
Protection	-	-	-	-	-
TC Assurance	-	-	-	-	-
Priority	-	-	-	-	-

* QoS enhancement is explicitly supported

0 QoS enhancement is possible if additional functionality is added to the transport entity

Tab. 3 QoS enchancement by the ECMA transport protocol

6.2 Based on Protocol Options

A number of contributions show dissatisfaction with the current hierarchical structure proposed by ECMA /ECM80/ for the Transport Protocols.

One objection which is often made is that the scheme proposed by ECMA could well become untenable as further optional functions are identified. There is, for example, no single natural order of preference for:

- encryption;
- error control;
- multiplexing.

As it was mentioned in the NOTE of paragraph 6.1, there is also discussion as to why it should not be possible to have only error recovery without necessarily the adjunction of multiplexing.

Another difficulty with the use of Classes as proposed by ECMA arises when attempts are made to match Classes with underlying Network services. An approach which attempts to associate a particular Transport Protocol Class with a particular Network type seems crude.

For example, Class 4 of ECMA may use an error detection scheme more complex than necessary for Networks which introduces only loss of Network Service Data Units but not their mis-ordering and/or duplication.

In order to overcome these difficulties, another approach to the enhancement of Quality of the Service provided by the Network Layer has been proposed.

The basic idea behind this proposal is that while a Network Connection is not inherently acceptable or unacceptable to the Transport Layer, a Quality Of Service acceptable to one Transport User may be quite unacceptable to another.

In some cases, the service provided by the Network Connection will have both the properties and quality required by the Transport User. This can occur across the whole range of Transport User requirements. When this is true it is not necessary for the Transport Layer to implement extra functions. As an example if it is possible to use the service of the Network Layer flow control and the associated services of Reset and Expedited data than this should be done. The provision of functions for flow control in the Transport Layer itself must lead to the definition and use of TPDUs for acknowledgement. These will be carried as User data in the Network Layer, and thus will have charges associated with them. In the worst case this will lead to a dubling of communications costs. This conflicts with the requirements for optimal use of resources and minimum cost.

Therefore, an estimation of the Quality Of Service should be based on comparison between the needs of the Transport User and the facilities offered by the Network service.

A list of Quality of Service characteristics and their associated possible enhancement mechanisms follows (see section 4.1):

<u>Quality of Service characteristic</u>	<u>Mechanism</u>
Throughput	downward multiplexing TPDU numbering and credit mechanism
Transit-Delay	downward multiplexing + short NSDU's
Residual Error Rate - error detection: - loss - corruption	 acknowledgement and time-out checksum

- duplication/misordering	TPDU sequence numbering
- misdelivery	connection identification
- error recovery:	retransmission forward error correcting code
Connection Set Up Delay	sequential reuse of a maintained NC multiplexing if unused capacity on the NC
Resilience of Transport Connection	reconnect mechanisms on same or different NCs
Cost Saving	upward multiplexing sequential reuse of maintained NC use (N)-connectionless transfer non-maintained NC avoid costly functionalities use low cost NC's adapt TPDU + NSDU sizes
Protection	checksum on TPDU header exchange security information at Transport Connection Establ.
Connection Assurance	idle traffic
Priority	exchange of priority information

where:

TPDU ==> Transport Protocol Data Unit
NSDU ==> Network Service Data Unit
NC ==> Network Connection

Concluding it may be said that:

- unless the hierarchy/optionality issue is resolved in WG6 no clear picture can be given for the stability of the Transport Protocols
- unless the Transport Protocols are stable no clear picture can be given for the Transport Service.

7. Hierarchical Protocol Class vs. Protocol Classes based on QOS

By looking at the ECMA proposal for the Transport Protocol /ECM80/ it is obvious that the set of services and the qualitative enhancements supported by the protocol described in the proposal appear to be inter-related. It is the view of a number of experts that the requirements of applications are best met if the set of services provided and the use of qualitative enhancements are independent. In particular, in the ECMA proposal, it can happen that some services are not provided when no enhancements are required. For example, in

Class 0 and 1, the expedited service is not provided.

Practical experience in the implementation of Protocol and Networking Standards has demonstrated that any standards aiming at covering a wide spectrum of requirements through the medium of a shopping list of options is one that tends both to defeat the objectives of standardization and be costly to implement. If the list of options is large then the possible combinations of options is very large and the possibility of successful negotiation of a workable set for a given investment correspondingly small.

To achieve the objectives of interworking with previously defined standards covering a wide spectrum for applications, e.g. HDLC, ECMA has found it necessary to expend considerable post standardization effort on the definition of preferred or selected classes of protocol. To avoid this effort a basic requirement of the Transport Protocol was that it should have built into from the outset some kind of rational class structure.

Having made this decision it was also recognized from the outset, that any class structure aimed at achieving significant reduction of variance would necessarily need to be a compromise not satisfying all possible requirements in an optimum manner. The ECMA class structure, as it is embodied in document /TC80/ is one such compromise. It defines a structure based on a hierarchy for that set of functions of the Transport Layer clearly existing in a hierarchy and treats the remaining functions as options to be invoked within a class.

This scheme was chosen to facilitate evolutionary development of low cost devices initially implementing only the lower numbered classes. Consequently, it has an implementation bias. However, it is felt that a class structure based on any other criterion would be equally open to criticism given that it achieved a similar degree of reduction of variance.

References

/TS80/ : ISO/TC97/SC16 /TS80/: "Draft Transport Service Specification, November 80

/ECM80/ : Standard ECMA-72:"Transport Protocol", Geneve, December 80.

/TC80/ : ISO/TC97/SC16 N564: "Transport Protocol", November 80.

/CCI80/ : Recommendation S.72."Basic Network Independent Transport Service for TELETEX", Geneve, December 80.